



HAL
open science

Phase transition of the 3-majority opinion dynamics with noisy interactions

Francesco d'Amore, Isabella Ziccardi

► **To cite this version:**

Francesco d'Amore, Isabella Ziccardi. Phase transition of the 3-majority opinion dynamics with noisy interactions. *Theoretical Computer Science*, 2025, 1028, pp.115030. 10.1016/j.tcs.2024.115030 . hal-04860011

HAL Id: hal-04860011

<https://hal.science/hal-04860011v1>

Submitted on 31 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Phase Transition of the 3-Majority Dynamics with Uniform Communication Noise

Francesco d'Amore and Isabella Ziccardi

Bocconi University, BIDSa, Italy.
francesco.damore2@unibocconi.it
isabella.ziccardi@unibocconi.it

Abstract

Communication noise is a common feature in several real-world scenarios where systems of agents need to communicate in order to pursue some collective task. Indeed, many biologically inspired systems that try to achieve agreements on some opinion must implement *resilient* dynamics, i.e. are not strongly affected by noisy communications. In this work, we study the 3-MAJORITY dynamics, an opinion dynamics that has been shown to be an efficient protocol for the majority consensus problem, in which we introduce a simple feature of uniform communication noise, following D'Amore et al. (2022). We prove that, in the fully connected communication network of n agents and in the binary opinion case, the process induced by the 3-MAJORITY dynamics exhibits a phase transition. For a noise probability $p < 1/3$, the dynamics reaches in logarithmic time an almost-consensus metastable phase which lasts for a polynomial number of rounds with high probability. We characterize this phase by showing that there exists an attractive equilibrium value $s_{\text{eq}} \in [n]$ for the bias of the system, i.e. the difference between the majority community size and the minority one. Moreover, we show that the agreement opinion is the initial majority one if the bias towards it is of magnitude $\Omega(\sqrt{n \log n})$ in the initial configuration. If, instead, $p > 1/3$, we show that no form of consensus is possible, and any information regarding the initial majority opinion is lost in logarithmic time with high probability. Despite more communications per round being allowed, the 3-MAJORITY dynamics surprisingly turns out to be less resilient to noise than the UNDECIDED-STATE dynamics, whose noise threshold value is $p = 1/2$.

Keywords: opinion dynamics, consensus problem, randomized algorithms, distributed computing.

1 Introduction

The *consensus problem* is a fundamental problem in distributed computing [6] in which we have a system of agents supporting opinions that interact between each other by exchanging messages, with the goal of reaching an agreement on some *valid* opinion (i.e. an opinion initially present in the system). In particular, we focus on the *majority consensus problem* where the goal is to converge towards the initial majority opinion. The numerous theoretical studies in this area find justifications in many different application scenarios, ranging from social networks [2, 41], swarm robotics [5], cloud computing, communication networks [44], and distributed databases [20], to biological systems [25, 26]. As for the latter, the goal of the majority consensus problem is to model some real-world scenarios where biological entities need to communicate and agree in order to pursue some collective task. Many biological entities in different real situations perform this type of process, e.g. molecules [13], bacteria [4], flock of birds [9], school of fish [45], or social insects [27], such as honeybees [43].

In such scenarios, some form of noise often affects communication among agents. For this reason, one of the main goals of network information theory is to guarantee reliable communications in noisy networks [28]. In this context, error-correcting codes are very effective methods to reduce communication errors in computer systems [34, 40], and this is why many theoretical studies of the (majority) consensus problem assume that communication between entities occurs without error, and instead consider some adversarial behavior (e.g. byzantine fault [8]). Despite their effectiveness in computer applications, error-correcting codes are quite useless if we want to model consensus in biological systems. Indeed, they involve sending complicated codes through communication links, and it is reasonable to assume that biological type entities communicate between each other in a simpler way. For this reason, in recent years many works have been focusing on the study of opinion dynamics where the communication between entities is unreliable and subjected to uniform noise [17, 18, 25, 26].

Opinion dynamics do not have a formal definition, but we provide an informal one for the sake of the intuition, following [24].

Definition 1 (Opinion dynamics – informal). *An opinion dynamics is a synchronous distributed algorithm characterized by a very simple structure. In this structure, the state of a node at round t depends only on its own state and a symmetric function of the multiset of states of its neighbors at round $t - 1$. This function defines the update rule and can be randomized. Furthermore, the update rule is the same for every graph and every node and does not change over time.*

However, within the constraints of Definition 1, one can still devise very complex updating rules, which would violate the spirit of simplicity that characterizes the concept of opinion dynamics.

The first consensus dynamics (i.e. opinion dynamics for the consensus problem) that have been studied in the presence of noise communication are linear opinion dynamics, such as the VOTER dynamics and the AVERAGING dynamics. In these dynamics, the function describing a single agent’s state update rule is linear with respect to the states of its neighbors. Interestingly, they were studied in the presence of uniform noise communication [36] or of *stubborn agents* (i.e. agents that never change opinion) [38, 39, 49]. In these settings, only *metastable* forms of consensus can be achieved, where a large subset of the agents agree on an opinion while other opinions remain supported by smaller subsets of agents. However, the VOTER model has a slow convergence time even in fully connected networks and a large initial bias towards some majority opinion [31], and the AVERAGING dynamics requires agents to perform non-trivial computation and, more importantly, to have large local memory. For these reasons, linear opinion dynamics struggle explaining the observed metastable consensus in multi-agent systems [12, 16, 24], and many research papers have begun to investigate new, more plausible, non-linear opinion dynamics.

In this work, we consider the 3-MAJORITY dynamics, which is based on a majority rule, the latter being widely employed also in the biological research field [14, 22]. In particular, we introduce in the system a uniform communication noise feature, following the definition of [18]. Such dynamics, without communication noise, has a similar behavior to the UNDECIDED-STATE dynamics [6]. As we describe in the next section, the two dynamics behave similarly even in the presence of uniform noise, as both exhibit a phase transition, but with different thresholds. Indeed, although the 3-MAJORITY dynamics makes use of more per-round communications, it turns out to be less resilient to noise than the UNDECIDED-STATE dynamics.

1.1 Our result

In this work, we study the 3-MAJORITY dynamics over a complete network of n agents holding opinions from a binary set $\Sigma = \{\text{ALPHA}, \text{BETA}\}$. In each round, each agent pulls the opinions of three neighboring agents chosen independently uniformly at random and updates its opinion to the majority of the three. This dynamics is shown to be a fast, robust protocol for the majority consensus problem in different network topologies, ranging from complete graphs to

sparser graphs [6]. For further details on previous results about the 3-MAJORITY dynamics, we defer the reader to Section 1.2.

In this work, we introduce to the 3-MAJORITY dynamics an *uniform communication noise* feature, following the definition in [18], for which we give an equivalent formulation. Each communication with a sampled neighbor is noisy with probability $p \in (0, 1)$, i.e. with probability p the received opinion is sampled u.a.r. in the opinion set Σ . Instead, with probability $1 - p$ the communication is unaffected by noise.

Even if considering the complete graph is a strong assumption for such communication networks, we note that, in every round, an agent pulls an opinion from three neighbors, resulting in a round-by-round communication pattern that forms a dynamic graph with $\mathcal{O}(n)$ edges. Furthermore, this model can capture the behavior of bio-inspired multi-agent systems where mobile agents meet randomly at a relatively high rate. For more details about models for bio-inspired swarms of agents, we refer to [46].

To keep track of the progress of the dynamics towards consensus, we describe the dynamics via the bias at time t , denoted by s_t , which represents the difference between the sizes of the majority and minority opinion communities at time t . Note that the protocol, in the case of a complete graph of fixed size n with binary opinions, is completely described by $\{s_t\}_t$. We prove that, in the aforementioned setting, the 3-MAJORITY dynamics exhibit a phase transition. Our results are summarized in the following theorem.

Theorem. *Let $\{s_t\}_{t \geq 0}$ be the bias of the process induced by the 3-MAJORITY dynamics with uniform noise probability p , and let s_0 the initial value of the bias. The following statements hold.*

1. *If $p < 1/3$ and if $s_0 = \Omega(\sqrt{n \log n})$, consider $s_{eq} = \frac{n}{1-p} \sqrt{\frac{1-3p}{1-p}}$, and let $\varepsilon > 0$ be any sufficiently small constant. Then, there exists a time $\tau_1 = \mathcal{O}(\log n)$ such that, w.h.p.,¹ the process at time τ_1 reaches a metastable almost-consensus phase characterized by the equilibrium point s_{eq} , i.e.*

$$s_{\tau_1} \in [(1 - \varepsilon)s_{eq}, (1 + \varepsilon)s_{eq}].$$

Moreover, the bias remains confined in such interval for $n^{\Theta(1)}$ rounds w.h.p.

2. *If $p < 1/3$ and if $s_0 = \mathcal{O}(\sqrt{n \log n})$, there exists a time $\tau_2 = \mathcal{O}(\log n)$ such that, w.h.p., the system becomes unbalanced towards an opinion, i.e.*

$$s_{\tau_2} = \Omega(\sqrt{n \log n}).$$

3. *If $p > 1/3$ and if $s_0 = \Omega(\sqrt{n \log n})$, there exists a time $\tau_3 = \mathcal{O}(\log n)$ such that, w.h.p., at time τ_3 the majority opinion is lost, i.e. $s_{\tau_3} = \mathcal{O}(\sqrt{n})$. In addition, with constant probability, at time $\tau_3 + 1$ the majority opinion changes. Moreover, for $n^{\Theta(1)}$ additional rounds the absolute value of the bias is $\mathcal{O}(\sqrt{n \log n})$ w.h.p.*

Interestingly, when $p < 1/3$, the equilibrium point s_{eq} is the fixed point of the function determining the expectation of the bias at the next step, i.e. $\mathbb{E}[s_{t+1} | s_t = s_{eq}] = s_{eq}$.

Our results show that the 3-MAJORITY dynamics is less resilient to noise compared to the UNDECIDED-STATE dynamics, despite allowing more communication per-round. Indeed, in [18] (conference version [19]), it is demonstrated that the phase transition for the UNDECIDED-STATE dynamics has a threshold at $p = 1/2$.² In the UNDECIDED-STATE dynamics, at each round, each agent pulls a single neighboring opinion x uniformly at random. If the agent's former opinion y

¹An event holds *with high probability* (w.h.p. in short) with respect to n if the probability that it occurs is at least $1 - n^{-\Theta(1)}$.

²We remark that in [18], the threshold for the UNDECIDED-STATE dynamics is $p' = 1/6$ and corresponds to $p = 1/2$ in our model. In fact, in the cited work, each sampled opinion is non-noisy with probability $1 - 2p'$, otherwise it turns into the opposite opinion with probability p' and into the undecided state with probability p' . This corresponds to a framework where with probability $1 - 3p'$ the opinion is non-noisy, while with probability $p = 3p'$ the received opinion is sampled u.a.r. among the possible ones.

differs from x , the agent becomes *undecided*. Once undecided, the agent will adopt any subsequent opinion it encounters.

The behavior of the UNDECIDED-STATE dynamics with noise is similar to the one of the 3-MAJORITY dynamics stated in the above theorem. If p is below the threshold value, the dynamics w.h.p. rapidly breaks the symmetry and converges in logarithmic time to a metastable phase of almost-consensus. If otherwise p is above the threshold, no form of consensus is possible, since the bias keeps bounded by $\mathcal{O}(\sqrt{n \log n})$ for a polynomial number of rounds, w.h.p. In comparison to our work, [18] do not find the precise value at the equilibrium point, but we believe that also the UNDECIDED-STATE dynamics has an equilibrium point if p is below the threshold, as the authors show in the experimental section ([18, Figure 1E]).

In Section 6, we provide some experiments, which confirm our theoretical results. In addition, we simulate the 3-MAJORITY dynamics over communication graphs sampled as Erdős Rényi random graphs with different parameters. The latter results provide experimental evidence that the 3-MAJORITY dynamics is fast and resilient even in sparser communication graphs with strong connection properties, and phase transitions similar to the fully connected case are exhibited. For this reason, we believe that sparser topologies are worth to be analyzed. Furthermore, it would be interesting to see whether the 3-MAJORITY dynamics with an arbitrary number of possible opinions, with the same noise model, has the exact same phase transition at the noise threshold value $p = 1/3$. We remark that when $p = 1/3$, this corresponds to the fact that, for each node and at each round, exactly one communication among the three ones is noisy in expectation.

1.2 Related Works

Consensus problem. The study of the 3-MAJORITY dynamics arises on the ground of the results obtained for the MEDIAN dynamics in [21]. The MEDIAN dynamics considers a totally ordered opinion set, in which each agent pulls two neighbor opinions i, j u.a.r. and then updates its opinion k to the median between i, j , and k . The MEDIAN dynamics is a fault-tolerant, efficient dynamics for the majority consensus problem. However, as pointed out in [6], the MEDIAN dynamics may not guarantee with high probability convergence to a valid opinion in case of the presence of an adversary, which is needed for the consensus problem. Moreover, the opinion set must have an ordering, property that might not be met by applicative scenarios such as biological systems [6]. These facts naturally lead researchers to look for efficient dynamics that satisfy the above requirements.

To the best of our knowledge, [1] is the first work analyzing the h -MAJORITY dynamics. In detail, in the h -MAJORITY dynamics we have n nodes and, at every round, every node pulls the opinion from h random neighbors and sets his new opinion to the majority one (ties are broken arbitrarily). More extensive characterizations of the 3-MAJORITY dynamics over the complete graph are given in [7, 8, 10, 29].

In [7] it is shown that the 3-MAJORITY dynamics is a fast, fault-tolerant protocol for (valid) majority consensus in the case of $k \geq 2$ colors, provided that there is an initial bias towards some majority opinion. Furthermore, [7] shows an exponential time-gap between the 3-MAJORITY consensus process and the median process in [21], thus establishing its efficiency. In [8], the analysis is extended to any initial configuration in the many-color case, in the presence of a different kind of bounded adversaries. The authors of [8] emphasize how the absence of an initial majority opinion considerably complicates the analysis, in that it must be proved that the process breaks the initial symmetry despite the presence of the adversary. Indeed, before the symmetry breaking, the adversary is more likely to cause undesired behaviors. The strongest result about the convergence of the 3-MAJORITY is in [29]. The authors show that in the case of k opinions, the process converges in time $O(k \log n)$ rounds, and this result it is tight when $k = O(\sqrt{n})$. The 3-MAJORITY dynamics is also studied in different graph topologies: [32] analyzes the 3-MAJORITY process in graphs of minimum degree n^α , with $\alpha = \Omega((\log \log n)^{-1})$, starting from random biased binary configurations.

Another efficient opinion dynamics for the majority consensus problem is the 2-CHOICES dynamics, in which each agent samples two neighbors u.a.r. and updates its opinion to the majority opinion among its former opinion and the two sampled neighbor opinions if there is any, otherwise it keeps its opinion. We remark that the expected round-by-round behaviors of the 2-CHOICES dynamics and that of the 3-MAJORITY are the same, while the actual behaviors differ substantially in high probability [10]. For example, we have run simple experiments that suggest that our uniform noise model on the 2-CHOICES dynamics yields a threshold noise value $p = 1/2$, just like the UNDECIDED-STATE dynamics.

As the 2-CHOICES and the 3-MAJORITY dynamics, the UNDECIDED-STATE dynamics is an efficient majority consensus protocol, with the difference that it requires only one communication per-round for each agent. Further description was already given in the previous section. It is worth mentioning the more recent work [3], which analyzes a variant of the UNDECIDED-STATE dynamics in the many-color case starting from any initial configuration. For an overview on the state of the art about opinion dynamics we defer the reader to [6].

General consensus algorithms in the GOSSIP model. In [11,30], the authors address the majority consensus problem in the GOSSIP model where the underlying communication network is fully connected and the agents may support k different opinions. In the GOSSIP model, at each round every node samples a constant number of neighbors u.a.r. and exchanges B bits with it, where B usually is $\mathcal{O}(\log n)$. In [30], the authors propose an algorithm achieving majority consensus in $\mathcal{O}(\log k \log n)$ rounds, where the agents have local memory of size $\log k + \mathcal{O}(1)$ bits and can send messages of size $\log k + \mathcal{O}(1)$ bits. The majority consensus is reached provided that the initial bias is at least $\Omega(\sqrt{n \log n})$. In [11], the authors extend the result by [30] and provide a first algorithm achieving the majority consensus in $\mathcal{O}(\log k \log \log n)$ rounds, where each agent uses $\log k + \Theta(\log \log k)$ bit of memory, and a second algorithm with runtime $\mathcal{O}(\log n \log \log n)$ which uses $\log k + 4$ bit of memory for each agents. Both algorithms require an initial bias of $\Omega(\sqrt{n \log^2 n})$. Notice that the algorithms provided by [11,30] are more complex algorithms w.r.t. the consensus dynamics we mentioned before (e.g. 3-MAJORITY, 2-CHOICES, UNDECIDED-STATE), as they work in *phases*: the rules defining local computations and communications change over time and, hence, require higher local memory and communication bandwidth. Hence, such algorithms violate the constraints of Definition 1.

Consensus problem in the presence of noise or stubborn agents. The authors of [48] initiate the study of the consensus problem in the presence of communication noise. They consider the Vicsek model [47], in which they introduce a noise feature and a notion of robust consensus. Subsequently, dynamics for the consensus problem with noisy communications have received considerable attention. This direction is motivated, among many reasons, by the desire to find models for the consensus problem in natural phenomena [25].

The communication noise studied in this type of problem can be divided into two types: uniform (or unbiased) and non-uniform (or biased). The uniform noise wants to capture errors in communications between agents in real-world scenarios, in which communication noise affects all opinions in the same way. The non-uniform communication noise instead describes the case in which opinions are affected differently from one another.

The authors of [25] are the first to explicitly focus on the uniform noise model. In detail, they study the broadcast and the majority consensus problem when the opinion set is binary. In their model of noise, every bit in every exchanged message is flipped independently with some probability smaller than $1/2$. As a result, the authors give natural protocols that solve the aforementioned problems efficiently. The work [26] generalizes the above study to opinion sets of any cardinality.

As for the non-uniform communication noise case, in [17] it is considered the h -MAJORITY dynamics with a binary opinion set $\{\text{ALPHA}, \text{BETA}\}$, with a probability p that any received message is flipped towards a *fixed* preferred opinion BETA, while with probability $1 - p$ the former message

keeps intact. They suppose there is an initial majority agreeing on ALPHA, and they analyze the *time of disruption*, that is the time the initial majority is subverted. They prove there exists a threshold value $p^* = p^*(h)$ such that, if $p < p^*$ the time of disruption is at least polynomial w.h.p., and if $p > p^*$, the time of disruption is constant, w.h.p. Their result holds for any sufficiently dense graph.

The noise feature affecting opinion dynamics, when the underlying communication network is complete, has been shown to be equivalent to a model without noise, in which communities of stubborn agents (i.e. they never change opinion) are added to the network [18]. For this reason, we discuss some previous works that consider such a model.

In [49], the authors show that, in the voter model, the presence of stubborn agents with opposite opinions precludes the convergence to consensus. The work [42] studies the asynchronous voter rule and the asynchronous majority rule dynamics with Poisson clocks when the opinion set is binary. The authors use mean-field techniques, and focus on two different scenarios: In the first, some agents have a probability (which depends on their current opinion) not to update when the clock ticks. In the second, there are stubborn agents. In the second case, which directly relates to our work, they show that for the 3-MAJORITY dynamics, there are either one or two possible stable equilibria, depending on the sizes of the stubborn communities, which are reached in logarithmic time. If the two sizes are close to each other and not too large, then agreement on both opinions is possible. Otherwise, either no agreement is possible, or the process converges to an agreement towards a single opinion, which is that of the largest stubborn community. The case in which the two stubborn communities have equal size corresponds to the uniform communication noise model.

We remark, however, that in our work we consider a different setting: First, we consider the synchronous version of the 3-MAJORITY dynamics, which cannot be analyzed with the same tools of the asynchronous version. Furthermore, the mean-field arguments do not capture several aspects of the process, such as metastability, which in [42] is shown only through simulations. For example, the actual behavior of the process in the long term is oscillating in a very small interval around the equilibrium values, spending long times in those intervals, and eventually switching between the two. We characterize the width of the oscillation interval and show there is high probability of convergence, providing also a lower bound on the time the process spends in the equilibrium interval.

Structure of the paper

Section 2 contains the preliminaries for the analysis and the result statements. Section 3 contains the victory of majority case, i.e. where $p < 1/3$ and there is large enough initial bias. Section 4 contains the symmetry breaking result, i.e. where $p < 1/3$ but we have no condition on the initial bias. In Section 5 we analyze the victory of noise case, i.e. where $p > 1/3$. Finally, in Section 6 we show some simulation results.

2 Preliminaries

We consider the complete graph $G = K_n = (V, E)$ of n nodes (the agents), where each node is labelled uniquely with labels in $[n] := \{1, \dots, n\}$. Each node supports a binary opinion from a set of opinions $\Sigma = \{\text{ALPHA}, \text{BETA}\}$. The 3-MAJORITY dynamics defines a Markov Chain $\{M_t\}_{t \in \mathbb{N}}$ which is described by the opinion of the nodes at each time step, i.e. $M_t = (i_1(t), \dots, i_n(t)) \in \Sigma^n$ for every $t \geq 0$, where $i_j(t)$ is the opinion of node j at time t . The transition probabilities are characterized iteratively by the majority update rule as follows: given any time $t \geq 0$, let $M_t \in \Sigma^n$ be the state of the process at time t . Then, at time $t + 1$, each node $u \in V$ samples three agents independently uniformly at random (with replacement) and updates its opinion to the majority one among the sampled neighbor opinions. For the sake of clarity, we remark that when u samples a neighbor node twice, the corresponding opinion counts twice.

Communication noise. To this dynamics, we introduce a uniform communication noise feature. Let $0 < p < 1$ be a noise probability. When a node pulls a neighbor opinion, with probability p the received opinion is sampled u.a.r. in Σ and instead, with probability $1 - p$, the former opinion is kept intact as it is received.

Since the communication network is complete, knowledge of the graph's size implies that the process's state is fully characterized by the number of nodes supporting a given opinion. Hence, we can write $M_t = (a_t, b_t)$, where a_t is the number of the nodes supporting opinion ALPHA at time t , and b_t is the analogous for opinion BETA.

To keep track of the progress of the dynamic towards the consensus, we define the bias of the process at time t as the difference between the number of agents supporting the opinion BETA and the number of agents supporting ALPHA, i.e.

$$s_t = b_t - a_t = 2b_t - n, \quad (1)$$

which takes value in $\{-n, \dots, n\}$. We notice that, once n is fixed, the process can be fully characterized only by the values of the bias, i.e. $\{s_t\}_{t \geq 0}$. We remark that $s_t > 0$ if the majority opinion at time t is BETA and $s_t < 0$ if it is ALPHA. We say that configurations having bias $s_t \in \{n, -n\}$ are monochromatic, meaning that every node supports the same opinion, while a configuration with $s_t = 0$ is symmetric. We finally remark that the random variable b_t (and, analogously, a_t) is the sum of i.i.d. Bernoulli random variables. In detail, if $X_i^{(t)}$ is the r.v. yielding 1 if node i adopts opinion BETA at round $t + 1$, and 0 otherwise, then $b_t = \sum_{i \in [n]} X_i^{(t)}$. Therefore, for (1),

$$s_t = 2 \sum_{i \in [n]} X_i^{(t)} - n = \sum_{i \in [n]} (2X_i^{(t)} - 1), \quad (2)$$

where $(2X_i^{(t)} - 1)$ are i.i.d. taking values in $\{-1, 1\}$. For this reason, we can apply the Hoeffding bound (Lemma 17) to the bias.

Further notation. For any function $f(n)$, we make use of the standard Landau notation $\mathcal{O}(f(n)), \Omega(f(n)), \Theta(f(n))$. Furthermore, for a constant $c > 0$, we write $\mathcal{O}_c(f(n)), \Omega_c(f(n)),$ and $\Theta_c(f(n))$ if the hidden constant in the notation depends on c .

We first compute the expectation of the bias at time t , conditional on its value at time $t - 1$.

Lemma 1. *Let $\{s_t\}_{t \geq 0}$ be the process induced by the 3-MAJORITY dynamics with uniform noise probability $p \in (0, 1)$. The conditional expectation of the bias is*

$$\mathbb{E}[s_t \mid s_{t-1} = s] = \frac{s(1-p)}{2} \left(3 - \frac{s^2}{n^2}(1-p)^2 \right). \quad (3)$$

Proof. Let $b = b_t$ and $a = a_t$. Then $s = b - a$ and $n = a + b$, which implies $b = (n + s)/2$ and $a = (n - s)/2$. The probability that, when a node samples a neighbor, it receives opinion BETA is $b' = (b/n) \cdot (1 - p) + p/2$, where $(b/n) \cdot (1 - p)$ is the probability to receive a non-noisy message which contains opinion BETA, and $p/2$ is the contribution of the noise. Analogously, the probability that it receives opinion ALPHA is $a' = (a/n) \cdot (1 - p) + p/2$. Then, the probability the node updates its opinion to BETA is $(b')^3 + 3a'(b')^2$. For the sake of the calculations, we define $s' = b' - a' = (s/n)(1 - p)$. Since $1 = b' + a'$, we have $b' = (1 + s')/2$ and $a' = (1 - s')/2$. So, Eq. (1) implies that

$$\begin{aligned} \mathbb{E}[s_t \mid s_{t-1} = s] &= 2n \left((b')^3 + 3a'(b')^2 \right) - n \\ &= n \left(2(b')^2(b' + 3a') - 1 \right) \\ &= \frac{n}{2} \left((1 + s')^2(2 - s') - 2 \right) \end{aligned}$$

$$\begin{aligned}
&= \frac{n}{2} (3s' - (s')^3) \\
&= \frac{s(1-p)}{2} \left(3 - \frac{s^2}{n^2} (1-p)^2 \right).
\end{aligned}$$

□

By the lemma above, we deduce that there are three fixed points of the conditional expectation of the bias in the next step. The first one corresponds to $s = 0$, and the other (possible) equilibrium correspond to the condition

$$\frac{1-p}{2} \left(3 - \frac{s^2}{n^2} (1-p)^2 \right) = 1.$$

The latter condition results in

$$\begin{aligned}
s &= \pm \frac{n}{(1-p)} \cdot \sqrt{\frac{3(1-p) - 2}{(1-p)}} \\
&= \pm \frac{n}{(1-p)} \cdot \sqrt{\frac{1-3p}{1-p}},
\end{aligned}$$

which is well defined if only if $p \leq 1/3$. We will denote the absolute value of the latter two values by s_{eq} .

3 Victory of the majority

The aim of this section is to prove the following theorem, which shows how the dynamics solves the majority consensus problem when $p < 1/3$ in a weak form, since only an almost-consensus is reached.

Theorem 2 (Victory of the majority). *Let $\{s_t\}_{t \geq 0}$ be the process induced by the 3-MAJORITY dynamics with uniform noise probability $p < 1/3$. Let $\varepsilon > 0$ be any arbitrarily small constant such that $\varepsilon < 1/3$ and $\varepsilon \leq \frac{2(1-3p)}{3(1-p)}$, and let $\gamma > 0$ be any constant. Let $s_{\text{eq}} = \frac{n}{(1-p)} \sqrt{\frac{1-3p}{1-p}}$. Then, for any starting configuration s_0 such that $s_0 \geq \gamma \sqrt{n \log n}$ and for any sufficiently large n , the following holds w.h.p.:*

- (i) *there exists a time $\tau_1 = \mathcal{O}_{\gamma, \varepsilon, p}(\log n)$ such that $(1 - \varepsilon)s_{\text{eq}} \leq s_{\tau_1} \leq (1 + \varepsilon)s_{\text{eq}}$;*
- (ii) *there exists a value $c = \Theta_{\gamma, \varepsilon, p}(1)$ such that, for all $k \leq n^c$, $(1 - \varepsilon)s_{\text{eq}} \leq s_{\tau_1+k} \leq (1 + \varepsilon)s_{\text{eq}}$.*

In each of the following statements we assume that $\{s_t\}_{t \geq 0}$ is the bias of the process induced by the 3-MAJORITY dynamics with uniform noise probability $p < 1/3$.

We first show a lemma which states that, for any small constant $\varepsilon > 0$, whenever $s_{t-1} \notin [(1 - \varepsilon)s_{\text{eq}}, (1 + \varepsilon)s_{\text{eq}}]$, then s_t gets closer to the s_{eq} .

Lemma 3. *For any constant $0 \leq \varepsilon \leq \frac{2(1-3p)}{3(1-p)}$ and for any $\gamma > 0$, if $s \geq \gamma \sqrt{n \log n}$, the following statements hold*

- (i) *if $s \leq (1 - \varepsilon)s_{\text{eq}}$, then $\mathbb{P} [s_t \geq (1 + 3\varepsilon^2/4)s \mid s_{t-1} = s] \geq 1 - \frac{1}{n\gamma^2\varepsilon^4/32}$;*
- (ii) *if, $s \geq (1 + \varepsilon)s_{\text{eq}}$, then $\mathbb{P} [s_t \leq (1 - 3\varepsilon^2/4)s \mid s_{t-1} = s] \geq 1 - \frac{1}{n\gamma^2\varepsilon^4/32}$.*

Proof. We first notice that

$$(1 - \varepsilon)s_{\text{eq}} \leq \frac{n}{1-p} \sqrt{\frac{1-3p-2\varepsilon^2}{1-p}}. \quad (4)$$

Indeed,

$$1 - 3p - 2\varepsilon^2 - (1 - \varepsilon)^2(1 - 3p) = \varepsilon(2(1 - 3p) - 3\varepsilon(1 - p))$$

is non-negative if and only if $0 \leq \varepsilon \leq \frac{2(1-3p)}{3(1-p)}$.

From Lemma 1, if each $s \leq (1 - \varepsilon)s_{\text{eq}}$, then

$$\begin{aligned} \mathbb{E}[s_t \mid s_{t-1} = s] &= \frac{s(1-p)}{2} \left(3 - \frac{s^2}{n^2}(1-p)^2 \right) \\ &\geq s \left(\frac{3-3p}{2} - \frac{1-3p-2\varepsilon^2}{2} \right) \\ &= s(1 + \varepsilon^2). \end{aligned}$$

where the inequality follows by (4). By (2) and by the Hoeffding bound (Lemma 17), it holds that

$$\mathbb{P}[s_t \leq s(1 + \varepsilon^2) - s\varepsilon^2/4 \mid s_{t-1} = s] \leq e^{-s^2\varepsilon^4/(32n)} \leq e^{-\gamma^2\varepsilon^4 \log n/32} \leq \frac{1}{n\gamma^2\varepsilon^4/32}.$$

The second inequality in the lemma follows by a symmetric argument, observing that

$$(1 + \varepsilon)s_{\text{eq}} \geq \frac{n}{1-p} \sqrt{\frac{1-3p+2\varepsilon^2}{1-p}},$$

Indeed, the expression

$$(1 + \varepsilon)^2(1 - 3p) - (1 - 3p + 2\varepsilon^2) = \varepsilon(2(1 - 3p) - \varepsilon(1 + 3p))$$

is non-negative if and only if $0 \leq \varepsilon \leq \frac{2(1-3p)}{1+3p}$, which is guaranteed since $\varepsilon \leq \frac{2(1-3p)}{3(1-p)}$. \square

The following lemma serves to bound how far the bias can get from the interval $[(1 + \varepsilon)s_{\text{eq}}, (1 - \varepsilon)s_{\text{eq}}]$.

Lemma 4. *For any constants $\varepsilon > 0$ and $\gamma > 0$, if $s \geq \gamma\sqrt{n \log n}$, the following statements hold*

- (i) *if $s \leq (1 + \varepsilon)s_{\text{eq}}$, then $\mathbb{P}[s_t \geq (1 - \varepsilon - \varepsilon^2)s \mid s_{t-1} = s] \geq 1 - \frac{1}{n\gamma^2\varepsilon^2/16}$;*
- (ii) *if $s \geq (1 - \varepsilon)s_{\text{eq}}$ with $\varepsilon < 1$, then $\mathbb{P}[s_t \leq (1 + \varepsilon)s \mid s_{t-1} = s] \geq 1 - \frac{1}{n\gamma^2\varepsilon^2p^2}$.*

Proof. The proof is similar to that of the previous lemma. From Lemma 1, we get that

$$\mathbb{E}[s_t \mid s_{t-1} = s] \geq s \left(1 - \varepsilon - \frac{\varepsilon^2}{2} \right),$$

which follows since $s \leq (1 + \varepsilon)s_{\text{eq}}$ by simple calculations. By using the Hoeffding bound (Lemma 17), we get

$$\mathbb{P}\left[s_t \leq s \left(1 - \varepsilon - \frac{\varepsilon^2}{2} \right) - \frac{\varepsilon^2 \cdot s}{2} \mid s_{t-1} = s \right] \leq e^{-\frac{\gamma^2\varepsilon^4}{16}} = \frac{1}{n\frac{\gamma^2\varepsilon^2}{16}}.$$

The second claim follows symmetrically from Lemma 1 by observing that, since $s \geq (1 - \varepsilon)s_{\text{eq}}$

$$\mathbb{E}[s_t \mid s_{t-1} = s] \leq s(1 + (1 - 3p)\varepsilon).$$

The Hoeffding bound implies

$$\begin{aligned} \mathbb{P}[s_t \geq s(1 + \varepsilon) \mid s_{t-1} = s] &\leq \mathbb{P}[s_t \geq s(1 + (1 - 3p)\varepsilon) + 2p\varepsilon \cdot s \mid s_{t-1} = s] \\ &\leq e^{-\gamma^2\varepsilon^2p^2} = \frac{1}{n\gamma^2\varepsilon^2p^2}. \end{aligned}$$

\square

We provide another lemma to control the behavior of the bias. The proof consists in the application of simple concentration bounds.

Lemma 5. *For any constant $k > 0$, the following statements hold:*

(i) *if $s \geq s_{\text{eq}}$, then $\mathbb{P}[s_t \geq 2s_{\text{eq}}/3 \mid s_{t-1} = s] \geq 1 - 1/n^k$.*

(ii) *if $0 \leq s \leq 2s_{\text{eq}}/3$, then $\mathbb{P}[s_t \leq s_{\text{eq}} \mid s_{t-1} = s] \geq 1 - 1/n^k$.*

Proof. Let k be any arbitrarily large constant. As for (i), Lemma 1 gives that

$$\mathbb{E}[s_t \mid s_{t-1} = s] \geq s(1-p) \geq s_{\text{eq}}(1-p),$$

since $s_{\text{eq}} \leq s \leq n$. Then, let $\delta = (1-3p)/3 > 0$. By using the Hoeffding bound, it holds that

$$\mathbb{P}[s_t \leq s_{\text{eq}}(1-p) - \delta \cdot s_{\text{eq}} \mid s_{t-1} = s] \leq e^{-\frac{\delta^2 s_{\text{eq}}^2}{4}} \leq \frac{1}{n^k},$$

where the latter inequality holds since $s_{\text{eq}} = \Theta(n)$ and $s_{\text{eq}} > (2k/\delta) \log n$ for a sufficiently large n . As for (ii), Lemma 1 implies that

$$\mathbb{E}[s_t \mid s_{t-1} = s] \leq \frac{3s(1-p)}{2} \leq s_{\text{eq}}(1-p),$$

which is true since $0 \leq s \leq 2s_{\text{eq}}/3$. The Hoeffding bound then gives

$$\mathbb{P}[s_t \geq s_{\text{eq}}(1-p) + ps_{\text{eq}} \mid s_{t-1} = s] \leq e^{-\frac{p^2 s_{\text{eq}}^2}{4}} \leq \frac{1}{n^k},$$

where the latter inequality holds since $s_{\text{eq}} = \Theta(n)$ and so $s_{\text{eq}} > (2k/p) \log n$ for a sufficiently large n . □

We can piece together the above lemmas, which imply the following corollary.

Corollary 6. *For any constant $\varepsilon > 0$ such that $\varepsilon < \frac{2(1-3p)}{3(1-p)}$,*

(i) *If $|s_{\text{eq}} - s| \leq (\varepsilon/4)s_{\text{eq}}$, then*

$$\mathbb{P}[|s_{\text{eq}} - s_t| \leq \varepsilon s_{\text{eq}} \mid s_{t-1} = s] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^2 p^2 / 2^5}};$$

(ii) *If $(\varepsilon/4)s_{\text{eq}} \leq |s_{\text{eq}} - s| \leq s_{\text{eq}}/3$, then*

$$\mathbb{P}\left[|s_{\text{eq}} - s_t| \leq |s_{\text{eq}} - s| \cdot \left(1 - \frac{3\varepsilon^2}{2^5}\right) \mid s_{t-1} = s\right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{18} 3^2)}}.$$

Proof. First, we prove (i). By using Lemma 4 and the union bound, we have that

$$\mathbb{P}\left[\left(1 - \frac{\varepsilon}{4} - \frac{\varepsilon^2}{16}\right) \cdot \left(1 - \frac{\varepsilon}{4}\right) s_{\text{eq}} \leq s_t \leq \left(1 + \frac{\varepsilon}{4}\right) \cdot \left(1 + \frac{\varepsilon}{4}\right) s_{\text{eq}} \mid s_{t-1} = s\right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^2 p^2 / 2^5}}.$$

The claim follows by observing that

$$\left[\left(1 - \frac{\varepsilon}{4} - \frac{\varepsilon^2}{16}\right) \cdot \left(1 - \frac{\varepsilon}{4}\right) s_{\text{eq}}, \left(1 + \frac{\varepsilon}{4}\right) \cdot \left(1 + \frac{\varepsilon}{4}\right) s_{\text{eq}}\right] \subseteq [(1-\varepsilon)s_{\text{eq}}, (1+\varepsilon)s_{\text{eq}}].$$

As for claim (ii), we divide the proof in two different cases. Suppose, first, that $2s_{\text{eq}}/3 \leq s \leq (1-\varepsilon/4)s_{\text{eq}}$. A constant $\varepsilon/4 \leq \delta \leq 1/3$ exists such that $s = (1-\delta)s_{\text{eq}}$. Then, from Lemmas 3 and 4, we have that

$$\mathbb{P}\left[(1-\delta) \left(1 + \frac{3\varepsilon^2}{2^6}\right) s_{\text{eq}} \leq s_t \leq s_{\text{eq}} \mid s_{t-1} = s\right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / 2^{14}}}.$$

Notice that

$$\begin{aligned} & \left| s_{\text{eq}} - (1 - \delta) \left(1 + \frac{3\varepsilon^2}{26} \right) s_{\text{eq}} \right| = s_{\text{eq}} - (1 - \delta) \left(1 + \frac{3\varepsilon^2}{26} \right) s_{\text{eq}} \\ & = (s_{\text{eq}} - (1 - \delta)s_{\text{eq}}) \cdot \left(1 - \frac{(1 - \delta) \cdot \frac{3\varepsilon^2}{26} \cdot s_{\text{eq}}}{\delta \cdot s_{\text{eq}}} \right) \leq (s_{\text{eq}} - s) \cdot \left(1 - \frac{3\varepsilon^2}{25} \right), \end{aligned}$$

where in the last inequality we used that $\delta \leq 1/3$. Hence,

$$\mathbb{P} \left[|s_{\text{eq}} - s_t| \leq |s_{\text{eq}} - s| \cdot \left(1 - \frac{3\varepsilon^2}{25} \right) \mid s_{t-1} = s \right] \geq 1 - \frac{1}{n^{\gamma^2 p^2 \varepsilon^4 / 2^{14}}}. \quad (5)$$

Second, suppose $(1 + \varepsilon/4)s_{\text{eq}} \leq s \leq 3s_{\text{eq}}/2$. A constant $\varepsilon/4 \leq \delta \leq 1/3$ exists such that $s = (1 + \delta)s_{\text{eq}}$. From Lemmas 3 and 4, it holds that

$$\mathbb{P} \left[(1 + \delta)(1 - \delta - \delta^2) s_{\text{eq}} \leq s_t \leq (1 + \delta) \left(1 - \frac{3\varepsilon^2}{26} \right) s_{\text{eq}} \mid s_{t-1} = s \right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{18} 3^2)}},$$

where we used the union bound and the fact that $\delta \leq 1/3$. Notice that

$$\begin{aligned} |s_{\text{eq}} - (1 + \delta)(1 - \delta - \delta^2) s_{\text{eq}}| & = s_{\text{eq}} - (1 + \delta)(1 - \delta - \delta^2) s_{\text{eq}} \\ & = ((1 + \delta)s_{\text{eq}} - s_{\text{eq}}) \cdot \left(\frac{(1 + \delta)(\delta + \delta^2)s_{\text{eq}}}{\delta s_{\text{eq}}} - 1 \right) \\ & \leq ((1 + \delta)s_{\text{eq}} - s_{\text{eq}}) \cdot \left(\frac{16}{9} - 1 \right) \\ & = ((1 + \delta)s_{\text{eq}} - s_{\text{eq}}) \cdot \left(1 - \frac{2}{9} \right), \end{aligned}$$

where the inequality holds since $\delta \leq 1/3$. By simple calculations, it can be seen that $(1 + \delta) \left(1 - \frac{3\varepsilon^2}{26} \right) \geq 1$. Then, we have also that

$$\begin{aligned} \left| s_{\text{eq}} - (1 + \delta) \left(1 - \frac{3\varepsilon^2}{26} \right) s_{\text{eq}} \right| & = (1 + \delta) \left(1 - \frac{3\varepsilon^2}{26} \right) s_{\text{eq}} - s_{\text{eq}} \\ & = ((1 + \delta)s_{\text{eq}} - s_{\text{eq}}) \cdot \left[1 - \frac{(1 + \delta) \cdot \frac{3\varepsilon^2}{26} s_{\text{eq}}}{\delta s_{\text{eq}}} \right] \\ & \stackrel{(a)}{\leq} ((1 + \delta)s_{\text{eq}} - s_{\text{eq}}) \cdot \left[1 - 3 \left(1 + \frac{\varepsilon}{4} \right) \cdot \frac{3\varepsilon^2}{26} \right] \\ & \stackrel{(b)}{\leq} ((1 + \delta)s_{\text{eq}} - s_{\text{eq}}) \cdot \left[1 - \frac{9\varepsilon^2}{26} \right], \end{aligned}$$

where (a) holds since $\varepsilon/4 \leq \delta \leq 1/3$, and (b) holds since $\varepsilon > 0$. Thus,

$$\mathbb{P} \left[|s_{\text{eq}} - s_t| \leq |s_{\text{eq}} - s| \cdot \left[1 - \frac{9\varepsilon^2}{26} \right] \mid s_{t-1} = s \right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{18} 3^2)}}. \quad (6)$$

Combining Eqs. (5) and (6), we get that, whenever $(\varepsilon/4)s_{\text{eq}} \leq |s_{\text{eq}} - s| \leq s_{\text{eq}}/3$, then

$$\mathbb{P} \left[|s_{\text{eq}} - s_t| \leq |s_{\text{eq}} - s| \cdot \left[1 - \frac{3\varepsilon^2}{25} \right] \mid s_{t-1} = s \right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{18} 3^2)}}.$$

□

We are finally ready to prove the theorem.

Proof of Theorem 2. We divide the proof in different cases. First, suppose that $(\varepsilon/4)s_{\text{eq}} \leq |s_{\text{eq}} - s| \leq \varepsilon s_{\text{eq}}$. Let $T_1 = n^{\gamma^2 \varepsilon^4 p^2 / (2^{19} 3^2)}$. Then, from Corollary 6.(i) and (ii), and by using the chain rule, we have that

$$\mathbb{P} \left[\bigcap_{k=1}^{T_1} \{|s_{\text{eq}} - s_{t+k}| \leq \varepsilon s_{\text{eq}}\} \mid s_t = s \right] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{20} 3^2)}}.$$

This proves statement (ii) of the theorem.

Second, suppose that $\varepsilon s_{\text{eq}} \leq |s_{\text{eq}} - s| \leq s_{\text{eq}}/3$. Then, from Corollary 6.(ii) and by using the chain rule, a time T_2 exists, such that

$$T_2 = \mathcal{O} \left(-\frac{\log n}{\log \left(1 - \frac{3\varepsilon^2}{2^5} \right)} \right) = \mathcal{O}(\log n / \varepsilon^2)$$

such that

$$\mathbb{P} [|s_{\text{eq}} - s_{t+T_2}| \leq \varepsilon s_{\text{eq}} \mid s_t = s] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{20} 3^2)}}.$$

Third, suppose that $s \leq 2s_{\text{eq}}/3$. From Lemma 3.(i) and Lemma 5.(ii), by the chain rule and the union bound, there is a time

$$T_3 = \mathcal{O} \left(\frac{\log n}{\log \left(1 + \frac{3\varepsilon^2}{4} \right)} \right) = \mathcal{O}(\log n / \varepsilon^2)$$

such that

$$\mathbb{P} [2s_{\text{eq}}/3 \leq s_{t+T_3} \leq s_{\text{eq}} \mid s_t = s] \geq 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 / 2^6}}.$$

Then, we are in one of the first two cases, and we conclude by using the chain rule.

Fourth, suppose that $s \geq (1 + \frac{1}{3})s_{\text{eq}}$. From Lemma 3.(ii) and Lemma 5.(i), and by using the chain rule, a time T_4 exists, with $T_4 = \mathcal{O}(\log n)$, such that

$$\mathbb{P} [|s_{\text{eq}} - s_{T_4}| \leq s_{\text{eq}}/3 \mid s_t = s] \geq 1 - \frac{1}{n^{\gamma^2 / (3^4 2^6)}}.$$

Statement (i) of the theorem follows by setting $\tau_1 = T_2 + T_3 + T_4$. □

4 Symmetry breaking

The following theorem shows how the dynamics quickly break the initial symmetry. Combining this result with Theorem 2, it shows that the consensus problem is solved.

Theorem 7 (Symmetry breaking). *Let $\{s_t\}_{t \geq 0}$ be the process induced by the 3-MAJORITY dynamics with uniform noise probability $p < 1/3$, and let $\gamma > 0$ be any constant. Then, for any starting configuration s_0 such that $|s_0| \leq \gamma \sqrt{n \log n}$ and for any sufficiently large n , w.h.p. there exists a time $\tau_2 = \mathcal{O}_{\gamma, p}(\log n)$ such that $|s_{\tau_2}| \geq \gamma \sqrt{n \log n}$.*

In each of the following statements, we assume that $\{s_t\}_{t \geq 0}$ is the process induced by the 3-MAJORITY dynamics with uniform noise probability $p < 1/3$. The symmetry breaking analysis essentially relies on the following lemma, which has been proved in [15].

Lemma 8. *Let $\{X_t\}_{t \in \mathbb{N}}$ be a Markov Chain with finite-state space Ω and let $f : \Omega \mapsto [0, n]$ be a function that maps states to integer values. Let c_3 be any positive constant and let $m = c_3 \sqrt{n \log n}$ be a target value. Assume the following properties hold:*

(i) for any positive constant h , a positive constant $c_1 < 1$ (which depends only on h) exists, such that for any $x \in \Omega : f(x) < m$,

$$\mathbb{P} [f(X_t) < h\sqrt{n} \mid X_{t-1} = x] < c_1;$$

(ii) there exist two positive constants δ and c_2 such that for any $x \in \Omega : h\sqrt{n} \leq f(x) < m$,

$$\mathbb{P} [f(X_t) < (1 + \delta)f(X_{t-1}) \mid X_{t-1} = x] < e^{-c_2 f(x)^2/n}.$$

Then the process reaches a state x such that $f(x) \geq m$ within $\mathcal{O}_{c_2, \delta, c_3}(\log n)$ rounds with probability at least $1 - 2/n$.

Our goal is to apply the above lemma to the 3-MAJORITY process, which defines a Markov chain. In particular, we claim the hypothesis of Lemma 8 are satisfied when the bias of the system is $o(\sqrt{n \log n})$, with $f(\mathbf{x}) = s(\mathbf{x})$, $m = \gamma\sqrt{n} \log n$ for any constant $\gamma > 0$. Then, Lemma 8 implies the process reaches a configuration with bias greater than $\Omega(\sqrt{n \log n})$ within time $\mathcal{O}(\log n)$, w.h.p. We need to prove that the two hypotheses hold.

Lemma 9. For any constant $c_3 > 0$, let s be a value such that $|s| < c_3\sqrt{n} \log n$. Then,

(i) For any constant $h > 0$, there exists a positive constant $c_1 < 1$ depending only on h , such that

$$\mathbb{P} [s_t < h\sqrt{n} \mid s_{t-1} = s] < c_1;$$

(ii) Two positive constants δ, c_2 exist (depending only on p), such that if $|s| \geq h\sqrt{n}$, then

$$\mathbb{P} [s_t < (1 + \delta)s \mid s_{t-1} = s] < e^{-\frac{c_2 s^2}{n}}.$$

Proof. As for the first claim, a simple domination argument implies that

$$\mathbb{P} [|s_t| < h\sqrt{n} \mid s_{t-1} = s] \leq \mathbb{P} [|s_t| < h\sqrt{n} \mid s_{t-1} = 0]. \quad (7)$$

As shown in Section 2, s_t is a sum of n i.i.d. Rademacher r.v.s with zero mean and unitary variance. We can hence make use of the Lemma 18 (Berry-Essen inequality). In particular, let $\Phi(x)$ be the cumulative function of a standard normal distribution. A constant $C > 0$ exists such that

$$|\mathbb{P} [s_t \leq h\sqrt{n} \mid s_{t-1} = 0] - \Phi(h)| \leq \frac{C}{\sqrt{n}}.$$

Since $\Phi(h) = c$ for some constant $c > 0$ which depends only on h , we have that

$$c - \frac{C}{\sqrt{n}} \leq \mathbb{P} [s_t \leq h\sqrt{n} \mid s_{t-1} = 0] \leq c + \frac{C}{\sqrt{n}}.$$

Since $\mathbb{P} [|s_t| < h\sqrt{n} \mid s_{t-1} = 0] \leq \mathbb{P} [s_t \leq h\sqrt{n} \mid s_{t-1} = 0]$, for n large enough we get

$$\mathbb{P} [|s_t| < h\sqrt{n} \mid s_{t-1} = 0] < 2c.$$

By setting $c_1 = c/2$, we get claim (i) from Eq. (7).

As for the second claim, assume $s > 0$ and $h\sqrt{n} \leq s \leq h\sqrt{n} \log n$. By Lemma 1 and the fact that $h\sqrt{n} \leq s \leq h\sqrt{n} \log n \leq (1 - \sqrt{\varepsilon})s_{eq}$, we have (as in Lemma 3)

$$\begin{aligned} \mathbb{E} [s_t \mid s_{t-1} = s] &= \frac{s(1 - 2p)}{2} \left(3 - \frac{s^2}{n^2} (1 - 2p)^2 \right) \\ &\geq s \left(\frac{3}{2} - 3p - \frac{1 - 6p - 2\varepsilon}{2} \right) \end{aligned}$$

$$= s(1 + \varepsilon).$$

By the Hoeffding bound (Lemma 17), we get that

$$\mathbb{P}[s_t \leq s(1 + \varepsilon) - s\varepsilon/4 \mid s_{t-1} = s] \leq e^{-s^2\varepsilon^2/(32n)}.$$

Observe that $\mathbb{P}[|s_t| \leq s(1 + 3\varepsilon/4) \mid s_{t-1} = s] \leq \mathbb{P}[s_t \leq s(1 + 3\varepsilon/4) \mid s_{t-1} = s]$. Thus, we have the claim by setting $\delta = 3\varepsilon/4$ and $c_2 = \varepsilon^2/32$. \square

The symmetry breaking is then a simple consequence of the above Lemma.

Proof of Theorem 7. Apply Lemmas 8 and 9 with $h = c_3 = \gamma$. \square

5 Victory of noise

The following theorem shows that no form of consensus is possible when $p > 1/3$.

Theorem 10 (Victory of noise). *Let $\{s_t\}_{t \geq 0}$ be the process induced by the 3-MAJORITY dynamics with uniform noise probability $p > 1/3$. Let $\varepsilon > 0$ be any arbitrarily small constant such that $\varepsilon < \min\{1/4, (1-p), (3p-1)/2\}$ and let $\gamma > 0$ be any positive constant. Then, for any starting configuration s_0 such that $|s_0| \geq \gamma\sqrt{n \log n}$ and for any sufficiently large n , the following holds w.h.p.:*

- (i) *there exists a time $\tau_3 = \mathcal{O}_{\varepsilon,p}(\log n)$ such that $s_{\tau_3} = \mathcal{O}_{\varepsilon}(\sqrt{n})$ and, moreover, the majority opinion switches at the next round with probability $\Theta_{\varepsilon}(1)$;*
- (ii) *there exists a value $c = \Theta_{\gamma,\varepsilon}(1)$ such that, for all $k \leq n^c$, it holds that $|s_{\tau_3+k}| \leq \gamma\sqrt{n \log n}$.*

In each of the following statements, we assume that $\{s_t\}_{t \geq 0}$ is the process induced by the 3-MAJORITY dynamics with uniform noise probability $p > 1/3$.

We apply tools from drift analysis (Lemma 15) to the absolute value of the bias of the process, showing that it reaches magnitude $\mathcal{O}(\sqrt{n})$ quickly. Then, since the standard deviation of the bias is $\Theta(\sqrt{n})$, we have that the majority opinion switches with constant probability (Lemma 13). Finally, with Lemma 14, we show that the bias keeps bounded in absolute value by $\mathcal{O}(\sqrt{n \log n})$.

Lemma 11. *For any constant $\varepsilon > 0$ such that $\varepsilon < (1-p)$, if $s \geq 2\sqrt{n}/(\varepsilon^2)$, the following holds*

$$\mathbb{E}[|s_t| \mid s_{t-1} = s] \leq \mathbb{E}[s_t \mid s_{t-1} = s] \cdot \left(1 + \frac{\varepsilon}{2}\right).$$

Proof. It holds that

$$|s_t| \leq |s_t - \mathbb{E}[s_t \mid s_{t-1} = s]| + |\mathbb{E}[s_t \mid s_{t-1} = s]|.$$

Furthermore, from Lemma 1, it follows that $\mathbb{E}[s_t \mid s_{t-1} = s] \geq 0$ as long as $s \geq 0$. By writing

$$|s_t - \mathbb{E}[s_t \mid s_{t-1} = s]| = \sqrt{(s_t - \mathbb{E}[s_t \mid s_{t-1} = s])^2},$$

and by using the Jensen's inequality, it follows that

$$\begin{aligned} \mathbb{E}[|s_t| \mid s_{t-1} = s] &\leq \sqrt{\mathbb{E}\left[(s_t - \mathbb{E}[s_t \mid s_{t-1} = s])^2 \mid s_{t-1} = s\right]} + \mathbb{E}[s_t \mid s_{t-1} = s] \\ &= \sigma(|s_t| \mid s_{t-1} = s) + \mathbb{E}[s_t \mid s_{t-1} = s], \end{aligned} \tag{8}$$

where $\sigma(x)$ represents the standard deviation of a r.v. x . As pointed out in Section 2, the bias can be written as the sum of i.i.d. random variables $Y_i^{(t)}$ taking values in $\{-1, +1\}$. For such sum of variables, the variance is linear:

$$\sigma(|s_t| \mid s_{t-1} = s)^2 = \sum_{i=1}^n \sigma\left(Y_i^{(t)} \mid s_{t-1} = s\right)^2 \leq n,$$

where the latter inequality holds since $\sigma\left(Y_i^{(t)} \mid s_{t-1} = s\right)^2 \leq 1$ for every i . From Lemma 1, we deduce that

$$\mathbb{E}[s_t \mid s_{t-1} = s] \geq \frac{s(1-p)(3-(1-p)^2)}{2} \geq s(1-p).$$

Since $s \geq \frac{2\sqrt{n}}{\varepsilon^2} \geq \frac{2\sqrt{n}}{\varepsilon(1-p)}$, we get that $\mathbb{E}[s_t \mid s_{t-1} = s] \geq \frac{2\sqrt{n}}{\varepsilon}$. Combining the latter facts with (8), we obtain

$$\begin{aligned} \mathbb{E}[|s_t| \mid s_{t-1}] &\leq \mathbb{E}[s_t \mid s_{t-1} = s] \cdot \left(1 + \frac{\sigma(|s_t| \mid s_{t-1} = s)}{\mathbb{E}[s_t \mid s_{t-1} = s]}\right) \\ &\leq \mathbb{E}[s_t \mid s_{t-1} = s] \cdot \left(1 + \frac{\sqrt{n}}{\frac{2\sqrt{n}}{\varepsilon}}\right) \\ &\leq \mathbb{E}[s_t \mid s_{t-1} = s] \cdot \left(1 + \frac{\varepsilon}{2}\right). \end{aligned}$$

□

With next lemma, we show that the absolute value of the process quickly becomes of magnitude $\mathcal{O}(\sqrt{n})$.

Lemma 12. *For any constant $\varepsilon > 0$ such that $\varepsilon < \min\{(1-p), (3p-1)/2\}$ we define $s_{\min} = \sqrt{n}/\varepsilon^2$. Then, for any starting configuration s_0 such that $s_0 \geq s_{\min}$, with probability at least $1 - 1/n$ there exists a time $\tau = \mathcal{O}_\varepsilon(\log n)$ such that $|s_\tau| \leq s_{\min}$.*

Proof. Let $h(x) = \frac{\varepsilon \cdot x}{2}$ be a function. Let $X_t = |s_t|$ if $s_t \geq s_{\min}$, otherwise $X_t = 0$. We now estimate $\mathbb{E}[X_t - X_{t-1} \mid X_{t-1} \geq s_{\min}, \mathcal{F}_{t-1}]$, where \mathcal{F}_t is the natural filtration of the process X_t . We have that

$$\begin{aligned} &\mathbb{E}[X_t - X_{t-1} \mid X_{t-1} \geq s_{\min}, \mathcal{F}_{t-1}] \\ &= \mathbb{E}[X_t \mid X_{t-1} \geq s_{\min}, \mathcal{F}_{t-1}] - X_{t-1} \\ &\stackrel{(a)}{\leq} \mathbb{E}[|s_t| \mid s_{t-1} \geq s_{\min}, \mathcal{F}_{t-1}] - s_{t-1} \\ &\stackrel{(b)}{\leq} \mathbb{E}[s_t \mid s_{t-1} \geq s_{\min}, \mathcal{F}_{t-1}] \cdot \left(1 + \frac{\varepsilon}{2}\right) - s_{t-1} \\ &\stackrel{(c)}{\leq} s_{t-1}(1 - \varepsilon) \left(1 + \frac{\varepsilon}{2}\right) - s_{t-1} \\ &\leq -\frac{\varepsilon \cdot s_{t-1}}{2}, \end{aligned}$$

where (a) holds because $X_t \leq |s_t|$, (b) follows from Lemma 11, and (c) from Lemma 1. Thus,

$$\mathbb{E}[X_{t-1} - X_t \mid X_{t-1} \geq s_{\min}, \mathcal{F}_{t-1}] \geq h(X_{t-1}).$$

Since $h'(x) = \varepsilon/2 > 0$, we can apply Lemma 15.(iii). Let τ be the first time $X_t = 0$ or, equivalently, $|s_t| < s_{\min}$. Then

$$\begin{aligned} \mathbb{P}[\tau > t \mid s_0] &< \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \int_{s_{\min}}^{s_0} \frac{2}{\varepsilon \cdot y} dy\right)\right] \\ &\leq \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \int_{s_{\min}}^n \frac{2}{\varepsilon \cdot y} dy\right)\right] \\ &= \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \frac{2}{\varepsilon}(\log n - \log s_{\min})\right)\right] \\ &= \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \frac{2}{\varepsilon}((\log n)/2 + 2 \log \varepsilon)\right)\right] \end{aligned}$$

$$\leq \exp \left[-\frac{\varepsilon \cdot t}{2} + 1 + \frac{\log n}{2} \right].$$

If $t = 4(\log n)/\varepsilon$, then we get that $\mathbb{P}[\tau > t \mid s_0] < e^{-3(\log n)/2+1} < 1/n$. \square

Next lemma states that, whenever the absolute value of the bias is of order of $\mathcal{O}(\sqrt{n})$, then the majority opinion switches at the next round with constant probability.

Lemma 13. *For any constant $\varepsilon > 0$ such that $\varepsilon < 1/4$, and let s_{t-1} be a configuration such that $|s_{t-1}| = s \leq \sqrt{n}/\varepsilon$. Then, the majority opinion switches at the next round with constant probability.*

Proof. Without loss of generality, we assume that $s_{t-1} > 0$. Now, $s_{t-1} = b_{t-1} - a_{t-1}$, with $n/2 < b_{t-1} \leq n/2 + \sqrt{n}/(2\varepsilon)$ and $n/2 - \sqrt{n}/(2\varepsilon) \leq a_{t-1} < n/2$. Both b_{t-1} and a_{t-1} can be expressed as the sum of i.i.d. Bernoulli r.v.s. Since $\mathbb{E}[a_t \mid n/2 - \sqrt{n}/(2\varepsilon) \leq a_{t-1} < n/2] \leq n/2$, we have

$$\mathbb{P} \left[a_t \geq \frac{n}{2} + \frac{\sqrt{n}}{2\varepsilon} \mid s_{t-1} = s \right] = \mathbb{P} \left[a_t \geq \frac{n}{2} \cdot \left(1 + \frac{1}{\varepsilon\sqrt{n}} \right) \mid s_{t-1} = s \right] \geq e^{-\frac{9}{2\varepsilon^2}},$$

where the latter inequality holds by using the reverse Chernoff bound (Lemma 19), whose hypothesis is satisfied since $\varepsilon < 1/4$. Thus, there is at least constant probability that the majority opinion switches. \square

Next lemma shows that the signed bias decreases each round.

Lemma 14. *For any constant $\varepsilon > 0$ such that $\varepsilon \leq (3p-1)/2$, the following statements hold*

- (i) *if $s \geq \frac{\gamma}{2}\sqrt{n \log n}$, then $\mathbb{P}[s_t \leq (1-3\varepsilon/4)s \mid s_{t-1} = s] \geq 1 - \frac{1}{n^{\gamma^2\varepsilon^2/2^7}}$;*
- (ii) *if $s \geq 0$, then $\mathbb{P}[-\frac{\gamma}{2}\sqrt{n \log n} \leq s_t \leq s + \frac{\gamma}{2}\sqrt{n \log n} \mid s_{t-1} = s] \geq 1 - \frac{2}{n^{\gamma^2/8}}$.*

Proof. From Lemma 1, for each $s \geq 0$ it holds that

$$\mathbb{E}[s_t \mid s_{t-1} = s] \leq \frac{3s(1-p)}{2} \leq (1-\varepsilon)s, \quad (9)$$

where the second inequality is true since $\varepsilon \leq (3p-1)/2$. We now apply the Hoeffding bound (Lemma 17) to s_t :

$$\mathbb{P}[s_t \geq (1-\varepsilon)s + \varepsilon \cdot s/4] \leq e^{-s^2\varepsilon^2/(32n)} \leq e^{-\gamma^2\varepsilon^2 \log n/2^7} \leq \frac{1}{n^{\frac{\gamma^2\varepsilon^2}{2^7}}}.$$

As for the second claim, we notice that, from Eq. (9), $\mathbb{E}[s_t \mid s_{t-1} = s] \leq s$. The Hoeffding bound (Lemma 17) now implies that

$$\mathbb{P} \left[s_t \geq s + \frac{\gamma}{2}\sqrt{n \log n} \right] \leq e^{-\gamma^2 \log n/8} \leq \frac{1}{n^{\gamma^2/8}}.$$

Moreover, from Lemma 1, for any $0 \leq s \leq n$, $\mathbb{E}[s_t \mid s_{t-1} = s] \geq 0$. Applying again the Hoeffding bound, we get that

$$\mathbb{P} \left[s_t \geq -\frac{\gamma}{2}\sqrt{n \log n} \mid s_{t-1} = s \right] \leq e^{-\gamma^2 \log n/8} \leq \frac{1}{n^{\gamma^2/8}},$$

By the union bound, we get the second claim. \square

We are ready to prove Theorem 10.

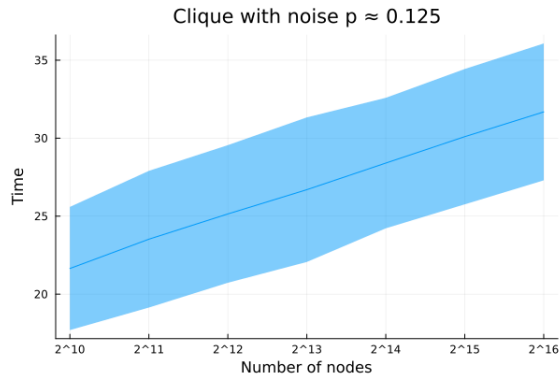
Proof of Theorem 10. Claim (i) follows directly from Lemmas 12 and 13. As for claim (ii), whenever the bias at some round $t = \tau + k$ becomes $|s_t| \geq (\gamma/2)\sqrt{n \log n}$, from Lemma 14.(ii) (and its symmetric statement), we have that $|s_t| \leq \gamma\sqrt{n \log n}$ with probability $1 - 2/n^{\frac{\gamma^2}{8}}$. Then, from Lemma 14.(i) it follows that the bias starts decreasing each round with probability $1 - 1/n^{\gamma^2\varepsilon^2/2^7}$ until reaching $(\gamma/2)\sqrt{n \log n}$. This phase in which the absolute value of the bias keeps bounded by $|\gamma\sqrt{n \log n}|$ lasts for at least $n^{\gamma^2\varepsilon^2/2^8}$ with probability at least $1 - 1/(2n^{\gamma^2\varepsilon^2/2^8})$ by using the chain rule. \square

6 Experiments

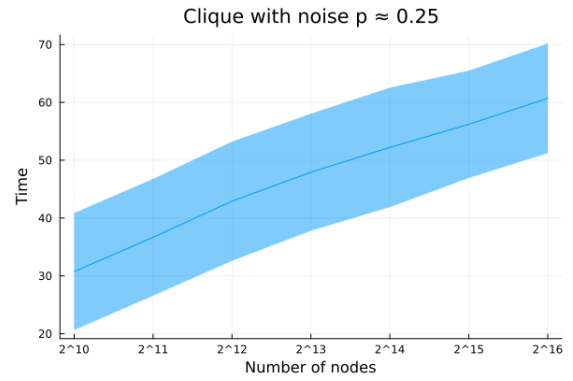
In this section, we describe the experiments we conducted on the 3-MAJORITY dynamics to show, in practice, the behavior of the bias for different input sizes and noise parameters.

In Fig. 1 we show the average convergence time of the 3-MAJORITY dynamics to almost-consensus in s and also Erdős-Rényi graphs $G_{n,q}$ (with $q = 1.25 \cdot \ln(n)/n$ to ensure connectivity with high probability). Interestingly, even though our theorems give asymptotic bounds, the convergence time behaves as expected even for small input sizes (e.g. $n = 2^{10}$ nodes). Furthermore, the asymptotic convergence time to almost-consensus seems to be logarithmic even in the Erdős-Rényi graph.

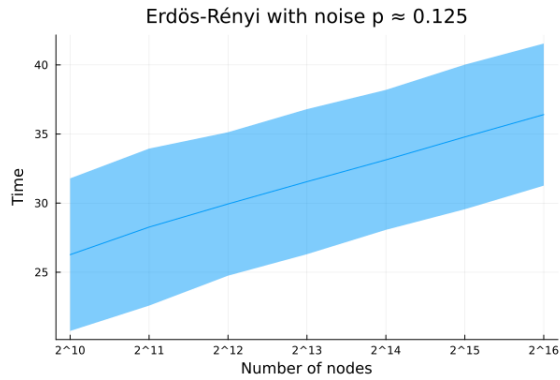
We also test the dynamics against expanders that have low degree: random regular graphs with degrees $d \in \{3, 5\}$ (Fig. 2). It seems that the convergence time is still logarithmic but quantitatively very different from the case of s : e.g. in Fig. 2b the convergence time for the random regular graph of degree $d = 5$ with 2^{16} nodes is around 700, whereas for s and Erdős-Rényi graphs it is between 60 and 80 (Figs. 1b and 1d). Notably, we also found out that the dynamics over the random regular graph with $d = 3$ doesn't seem to converge when the noise parameter is any value $p \geq 1/5$ (Fig. 3d). This phenomenon suggests that the noise values determining the phase-transition can depend on the expansion and/or sparsity of the underlying graph: these values seem to be reduced whenever the expansion of the graph decreases and its sparsity increases, as noticeable in Fig. 3. For future work, it would be interesting to study the relation between expansion/sparsity of the graph and the behavior of the 3-MAJORITY dynamics.



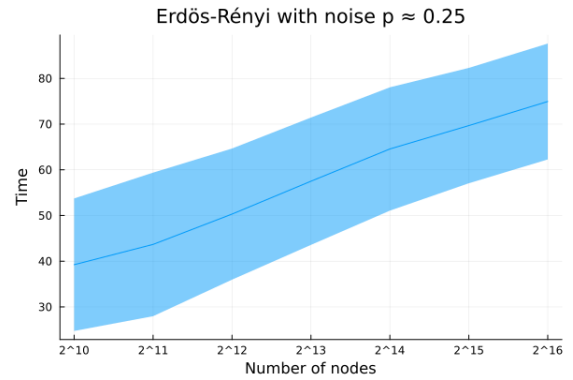
(a) Average convergence time to almost-consensus in a with noise parameter $p = 1/8$.



(b) Average convergence time to almost-consensus in a with noise parameter $p = 1/4$.

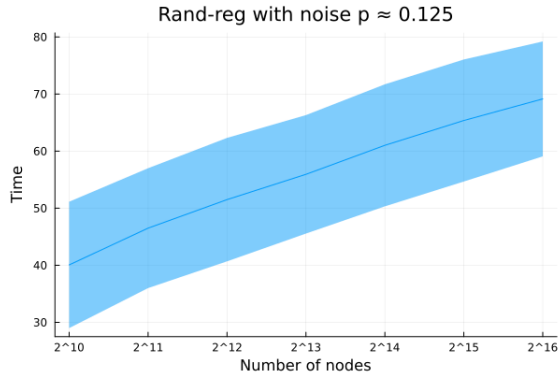


(c) Average convergence time to almost-consensus in an Erdős-Rényi graph $G_{n,q}$ with $q = 1.25 \cdot \ln(n)/n$ and noise parameter $p = 1/8$.

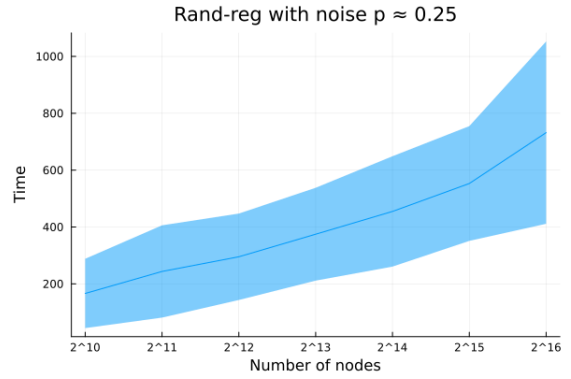


(d) Average convergence time to almost-consensus in an Erdős-Rényi graph $G_{n,q}$ with $q = 1.25 \cdot \ln(n)/n$ and noise parameter $p = 1/4$.

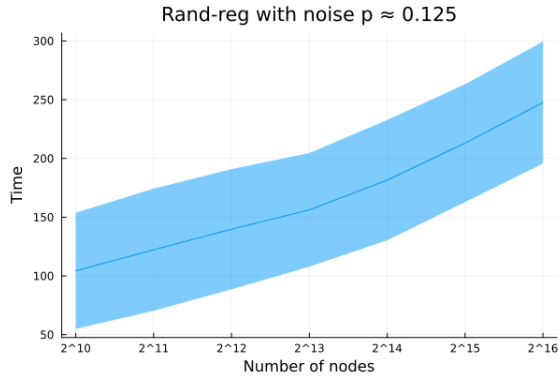
Figure 1: Average convergence times to almost-consensus in s and Erdos-Rényi graphs. The average is computed over 1000 different runs, and the random underlying graph is sampled at each run. The shaded areas represent the sample standard deviations.



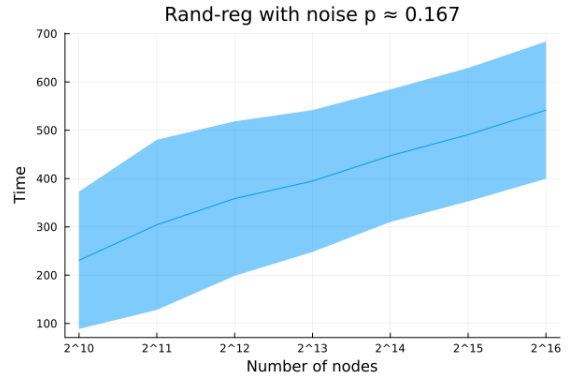
(a) Average convergence time to almost-consensus in a configuration graph with degree $d = 5$ with noise parameter $p = 1/8$.



(b) Average convergence time to almost-consensus in a configuration graph with degree $d = 5$ with noise parameter $p = 1/4$.

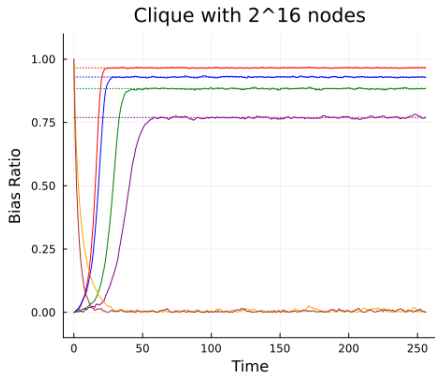


(c) Average convergence time to almost-consensus in a configuration graph with degree $d = 3$ and noise parameter $p = 1/8$.

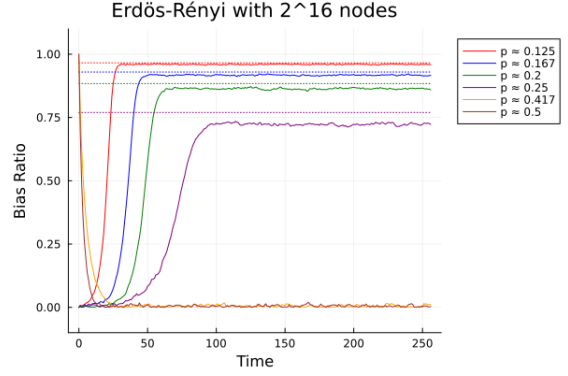


(d) Average convergence time to almost-consensus in a configuration graph with degree $d = 3$ and noise parameter $p = 1/6$.

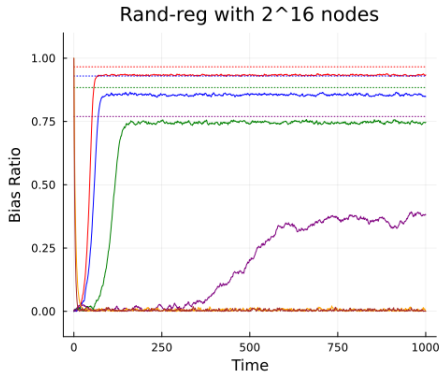
Figure 2: Average convergence times to almost-consensus in random regular graphs with degrees 5 and 3. The average is computed over 1000 different runs, and the random underlying graph is sampled at each run. The shaded areas represent the sample standard deviations.



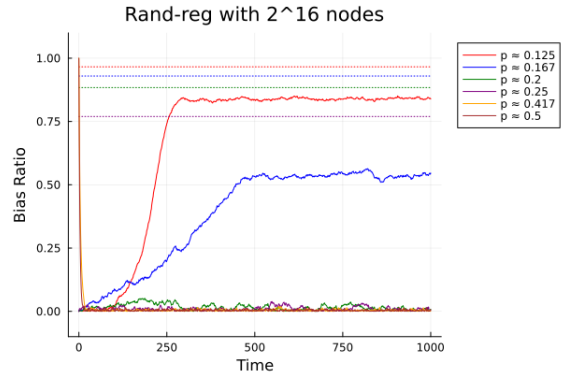
(a) Behavior of the bias in a clique for different noise values.



(b) Behavior of the bias in an Erdős-Rényi graph $G_{n,q}$ with $q = 1.25 \cdot \ln(n)/n$ for different noise values.



(c) Behavior of the bias in random regular graph with degree $d = 5$ for different noise values.



(d) Behavior of the bias in a random regular graph with degree $d = 3$ for different noise values.

Figure 3: Behaviors of the bias ratio $\left| \frac{\text{bias}}{\text{size}} \right|$ in different topologies and for different noise values over a single run of the dynamics. The noise parameter p is chosen from set $\{1/8, 1/6, 1/5, 1/4, 5/12, 1/2\}$. The dotted colored lines represent the bias ratio's equilibria in cliques, depending on the corresponding noise value. All graphs have 2^{16} nodes. For the sake of readability, we report here the values of the noise parameters for each colored line: red = $1/8$, blue = $1/6$, green = $1/5$, purple = $1/4$, gold = $5/12$, brown = $1/2$.

References

- [1] M. A. Abdullah and M. Draief. Global majority consensus by local majority polling on graphs of a given degree sequence. *Discret. Appl. Math.*, 180:1–10, 2015.
- [2] D. Acemoglu, G. Como, F. Fagnani, and A. E. Ozdaglar. Opinion fluctuations and disagreement in social networks. *Math. Oper. Res.*, 38(1):1–27, 2013.
- [3] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling. Fast consensus via the unconstrained undecided state dynamics. In J. S. Naor and N. Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. SIAM, 2022.
- [4] B. L. Bassler. Small talk: Cell-to-cell communication in bacteria. *Cell*, 109(4), 2002.
- [5] L. Bayindir. A review of swarm robotics tasks. *Neurocomputing*, 172:292–321, 2016.
- [6] L. Becchetti, A. E. F. Clementi, and E. Natale. Consensus dynamics: An overview. *SIGACT News*, 51(1), 2020.
- [7] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, R. Silvestri, and L. Trevisan. Simple dynamics for plurality consensus. *Distributed Comput.*, 30(4):293–306, 2017.
- [8] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, and L. Trevisan. Stabilizing consensus with many opinions. In R. Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 620–635. SIAM, 2016.
- [9] O. Ben-Shahar, S. Dolev, A. Dolgin, and M. Segal. Direction election in flocking swarms. *Ad Hoc Networks*, 12:250–258, 2014.
- [10] P. Berenbrink, A. E. F. Clementi, R. Elsässer, P. Kling, F. Mallmann-Trenn, and E. Natale. Ignore or comply?: On breaking symmetry in consensus. In E. M. Schiller and A. A. Schwarzmann, editors, *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC*. ACM, 2017.
- [11] P. Berenbrink, T. Friedetzky, G. Giakkoupis, and P. Kling. Efficient plurality consensus, or: the benefits of cleaning up from time to time. In I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPICs*, pages 136:1–136:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [12] L. Boczkowski, A. Korman, and E. Natale. Minimizing message size in stochastic communication patterns: fast self-stabilizing protocols with 3 bits. *Distributed Computing*, June 2019.
- [13] M. C. Carroll. The complement system in regulation of adaptive immunity. *Nature Immunology*, 5:981–986, 10 2004.
- [14] C. Chaouiya, O. Ourrad, and R. Lima. Majority rules with random tie-breaking in boolean gene regulatory networks. *PLOS ONE*, 8(7):1–14, 07 2013.
- [15] A. E. F. Clementi, M. Ghaffari, L. Gualà, E. Natale, F. Pasquale, and G. Scornavacca. A tight analysis of the parallel undecided-state dynamics with two colors. In I. Potapov, P. G. Spirakis, and J. Worrell, editors, *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, Liverpool, UK*, volume 117 of *LIPICs*, pages 28:1–28:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

- [16] A. Condon, M. Hajiaghayi, D. Kirkpatrick, and J. Manuch. Approximate majority analyses using tri-molecular chemical reaction networks. *Natural Computing*, 2019.
- [17] E. Cruciani, H. A. Mimun, M. Quattropiani, and S. Rizzo. Phase transitions of the k-majority dynamics in a biased communication model. In *ICDCN '21: International Conference on Distributed Computing and Networking*. ACM, 2021.
- [18] F. d’Amore, A. Clementi, and E. Natale. Phase transition of a nonlinear opinion dynamics with noisy interactions. *Swarm Intelligence*, Nov 2022.
- [19] F. D’Amore, A. E. F. Clementi, and E. Natale. Phase transition of a non-linear opinion dynamics with noisy interactions - (extended abstract). In A. W. Richa and C. Scheideler, editors, *Structural Information and Communication Complexity - 27th International Colloquium, SIROCCO 2020*, Lecture Notes in Computer Science. Springer, 2020.
- [20] M. Dietzfelbinger, A. Goerdt, M. Mitzenmacher, A. Montanari, R. Pagh, and M. Rink. Tight thresholds for cuckoo hashing via XORSAT. In S. Abramsky, C. Gavaille, C. Kirchner, F. M. auf der Heide, and P. G. Spirakis, editors, *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010*, Lecture Notes in Computer Science. Springer, 2010.
- [21] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. Stabilizing consensus with the power of two choices. In R. Rajaraman and F. M. auf der Heide, editors, *SPAA 2011: Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures*. ACM, 2011.
- [22] J. Dong, D. Fernández-Baca, F. McMorris, and R. C. Powers. Majority-rule (+) consensus trees. *Mathematical Biosciences*, 228(1):10–15, 2010.
- [23] D. P. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [24] Emanuele Natale. *On the Computational Power of Simple Dynamics*. PhD Thesis, Sapienza University of Rome, 2017.
- [25] O. Feinerman, B. Haeupler, and A. Korman. Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication. *Distributed Comput.*, 2017.
- [26] P. Fraigniaud and E. Natale. Noisy rumor spreading and plurality consensus. *Distributed Comput.*, 2019.
- [27] N. Franks, S. Pratt, E. Mallon, N. Britton, and D. 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*, 357:1567–83, 12 2002.
- [28] A. E. Gamal and Y. Kim. *Network Information Theory*. Cambridge University Press, 2011.
- [29] M. Ghaffari and J. Lengler. Nearly-tight analysis for 2-choice and 3-majority consensus dynamics. In C. Newport and I. Keidar, editors, *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018*. ACM, 2018.
- [30] M. Ghaffari and M. Parter. A polylogarithmic gossip algorithm for plurality consensus. In G. Giakkoupis, editor, *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 117–126. ACM, 2016.
- [31] Y. Hassin and D. Peleg. Distributed Probabilistic Polling and Applications to Proportionate Agreement. *Information and Computation*, 171(2):248–268, Dec. 2001.

- [32] N. Kang and N. Rivera. Best-of-three voting on dense graphs. In C. Scheideler and P. Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA*. ACM, 2019.
- [33] P. N. Klein and N. Young. On the number of iterations for dantzig-wolfe optimization and packing-covering approximation algorithms. *ArXiv*, cs.DS/0205046, 1999.
- [34] R. Koetter and F. R. Kschischang. Coding for errors and erasures in random network coding. *IEEE Trans. Inf. Theory*, 2008.
- [35] P. K. Lehre and C. Witt. Concentrated hitting times of randomized search heuristics with variable drift. In H. Ahn and C. Shin, editors, *Algorithms and Computation - 25th International Symposium, ISAAC 2014, Jeonju, Korea, Proceedings*, volume 8889 of *Lecture Notes in Computer Science*, pages 686–697. Springer, 2014.
- [36] W. Lin, L. Zhixin, and G. Lei. Robust consensus of multi-agent systems with noise. In *2007 Chinese Control Conference*, 2007.
- [37] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
- [38] M. Mobilia. Does a Single Zealot Affect an Infinite Group of Voters? *Physical Review Letters*, 91(2):028701, July 2003.
- [39] M. Mobilia, A. Petersen, and S. Redner. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08029, 2007.
- [40] T. K. Moon. *Error Correction Coding: Mathematical Methods and Algorithms*. Wiley-Interscience, USA, 2005.
- [41] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Auton. Agents Multi Agent Syst.*, 28(3):408–429, 2014.
- [42] A. Mukhopadhyay, R. R. Mazumdar, and R. Roy. Opinion dynamics under voter and majority rule models with biased and stubborn agents. *CoRR*, abs/2003.02885, 2020.
- [43] A. Reina, J. A. R. Marshall, V. Trianni, and T. Bose. Model of the best-of- n nest-site selection process in honeybees. *Phys. Rev. E*, 95:052411, May 2017.
- [44] Y. Ruan and Y. Mostofi. Binary consensus with soft information processing in cooperative networks. In *Proceedings of the 47th IEEE Conference on Decision and Control, CDC 2008*. IEEE, 2008.
- [45] D. J. Sumpter, J. Krause, R. James, I. D. Couzin, and A. J. Ward. Consensus decision making by fish. *Current Biology*, 18(22):1773–1777, 2008.
- [46] G. Valentini, E. Ferrante, and M. Dorigo. The best-of- n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9, 2017.
- [47] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.*, 75:1226–1229, Aug 1995.
- [48] L. Wang and Z. Liu. Robust consensus of multi-agent systems with noise. *Sci. China Ser. F Inf. Sci.*, 2009.
- [49] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. Binary Opinion Dynamics with Stubborn Agents. *ACM Trans. Econ. Comput.*, 1(4), 2013.

A Tools

We make use of the following general result on (super/sub)-martingales, which can be found in [35].

Lemma 15. *Let $\{X_t\}_{t \in \mathbb{N}}$ be a stochastic process adapted to a filtration $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$, over some state space $S \subseteq \{0\} \cup [x_{\min}, x_{\max}]$, where $x_{\min} \geq 0$. Let $h : [x_{\min}, x_{\max}] \rightarrow \mathbb{R}^+$ be a function such that $1/h(x)$ is integrable and $h(x)$ differentiable on $[x_{\min}, x_{\max}]$. Define $T := \min\{t \mid X_t = 0\}$. Then, the followings hold.*

(i) *If $\mathbb{E}[X_t - X_{t+1} \mid X_t \geq x_{\min}, \mathcal{F}_t] \geq h(X_t)$ and $\frac{d}{dx}h(x) \geq 0$, then*

$$\mathbb{E}[T \mid X_0] \leq \frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^{X_0} \frac{1}{h(y)} dy.$$

(ii) *If $\mathbb{E}[X_t - X_{t+1} \mid X_t \geq x_{\min}, \mathcal{F}_t] \leq h(X_t)$ and $\frac{d}{dx}h(x) \leq 0$, then*

$$\mathbb{E}[T \mid X_0] \geq \frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^{X_0} \frac{1}{h(y)} dy.$$

(iii) *If $\mathbb{E}[X_t - X_{t+1} \mid X_t \geq x_{\min}, \mathcal{F}_t] \geq h(X_t)$ and $\frac{d}{dx}h(x) \geq \lambda$ for some $\lambda > 0$, then*

$$\mathbb{P}[T > t \mid X_0] < \exp\left(-\lambda\left(t - \frac{x_{\min}}{h(x_{\min})} - \int_{x_{\min}}^{X_0} \frac{1}{h(y)} dy\right)\right).$$

(iv) *If $\mathbb{E}[X_t - X_{t+1} \mid X_t \geq x_{\min}, \mathcal{F}_t] \leq h(X_t)$ and $\frac{d}{dx}h(x) \leq -\lambda$ for some $\lambda > 0$, then*

$$\mathbb{P}[T < t \mid X_0] < \frac{e^{\lambda t} - e^\lambda}{e^\lambda - 1} \exp\left(-\lambda\left(\frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^{X_0} \frac{1}{h(y)} dy\right)\right).$$

For an overview on the forms of Chernoff bounds see [23].

Lemma 16 (Multiplicative forms of Chernoff bounds). *Let X_1, X_2, \dots, X_n be independent $\{0, 1\}$ random variables. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. Then:*

(i) *for any $\delta \in (0, 1)$ and $\mu \leq \mu_+ \leq n$, it holds that*

$$\mathbb{P}[X \geq (1 + \delta)\mu_+] \leq e^{-\frac{1}{3}\delta^2\mu_+}; \quad (10)$$

(ii) *for any $\delta \in (0, 1)$ and $0 \leq \mu_- \leq \mu$, it holds that*

$$\mathbb{P}[X \leq (1 - \delta)\mu_-] \leq e^{-\frac{1}{2}\delta^2\mu_-}. \quad (11)$$

We also make use of the Hoeffding bounds [37].

Lemma 17 (Hoeffding bounds). *Let $0 < a < b$ be two constants. Let X_1, X_2, \dots, X_n be independent random variables such that $\mathbb{P}[a \leq X_i \leq b] = 1$ for all $i \in [n]$. Let $X = \sum_{i=1}^n X_i$ and $\mathbb{E}[X] = \mu$. Then:*

(i) *for any $t > 0$ and $\mu \leq \mu_+$, it holds that*

$$\mathbb{P}[X \geq \mu_+ + t] \leq \exp\left(-\frac{2t^2}{n(b-a)^2}\right); \quad (12)$$

(ii) for any $t > 0$ and $0 \leq \mu_- \leq \mu$, it holds that

$$\mathbb{P}[X \leq \mu_- - t] \leq \exp\left(-\frac{2t^2}{n(b-a)^2}\right). \quad (13)$$

We make use of the following result which makes explicit the convergence “speed” in the central-limit theorem.

Lemma 18 (Berry-Esseen). *Let X_1, \dots, X_n be n i.i.d. (either discrete or continuous) random variables with zero mean, variance $\sigma^2 > 0$, and finite third moment. Let Z the standard normal random variable, with zero mean and variance equal to 1. Let $F_n(x)$ be the cumulative function of $\frac{S_n}{\sigma\sqrt{n}}$, where $S_n = \sum_{i=1}^n X_i$, and $\Phi(x)$ that of Z . Then, there exists a positive constant $C > 0$ such that*

$$\sup_{x \in \mathbb{R}} |F_n(x) - \Phi(x)| \leq \frac{C}{\sqrt{n}}$$

for all $n \geq 1$.

Finally, we use some anti-concentration inequalities know as reverse Chernoff bounds. The proof can be found in the appendix of [33].

Lemma 19 (Reverse Chernoff bounds). *Let X_1, X_2, \dots, X_n be i.i.d. $\{0, 1\}$ random variables. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$, with $\mu \leq n/2$. Furthermore, let $\delta \in (0, 1/2]$ be a constant. If $\delta^2\mu \geq 3$, then:*

(i) for any $\mu \leq \mu_+ \leq n$, it holds that

$$\mathbb{P}[X \geq (1 + \delta)\mu_+] \geq e^{-9\delta^2\mu_+}; \quad (14)$$

(ii) for any $0 \leq \mu_- \leq \mu$, it holds that

$$\mathbb{P}[X \leq (1 - \delta)\mu_-] \geq e^{-9\delta^2\mu_-}. \quad (15)$$