

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

Francesco d'Amore, Isabella Ziccardi

▶ To cite this version:

Francesco d'Amore, Isabella Ziccardi. Phase Transition of the 3-Majority Dynamics with Uniform Communication Noise. [Research Report] Inria & Université Cote d'Azur, CNRS, I3S, Sophia Antipolis, France; Università degli Studi dell'Aquila. 2021. hal-03467403v1

HAL Id: hal-03467403 https://hal.science/hal-03467403v1

Submitted on 6 Dec 2021 (v1), last revised 9 Dec 2021 (v2)

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

Eligible for the best student paper award*

Francesco d'Amore¹ and Isabella Ziccardi²

¹Université Côte d'Azur, Inria, CNRS, I3S, France. francesco.d-amore@inria.fr ²Università degli Studi dell'Aquila, L'Aquila, Italy. isabella.ziccardi@graduate.univaq.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. In particular, many biologically inspired systems that try to achieve agreements on some opinion must implement resilient dynamics that are not strongly affected by noisy communications. In this work, we study the popular 3-MAJORITY dynamics, an opinion dynamics which has been proved 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. 2020). 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. Furthermore, departing from previous analyses, we further characterize this phase by showing that there exists an attractive equilibrium value $s_{eq} \in [n]$ for the bias of the system, i.e. the difference between the majority community size and the minority one. Moreover, the agreement opinion turns out to be the initial majority one if the bias towards it is of magnitude $\Omega\left(\sqrt{n\log n}\right)$ in the initial configuration. If, instead, p > 1/3, 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 are allowed, the 3-MAJORITY dynamics surprisingly turns out to be less resilient to noise than the UNDECIDED-STATE dynamics (d'Amore et al. 2020), whose noise threshold value is p = 1/2.

1 Introduction

The consensus problem is a fundamental problem in distributed computing [6] in which we have a system of agents supporting some 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, many research papers focus on the majority consensus problem where the goal is to converge towards the initial majority opinion. The numerous theoretical studies in this area are justified by many different application scenarios, ranging from social networks [2,38], swarm robotics [5], cloud computing, communication networks [41], and distributed databases [18], to biological systems [23,24]. 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

^{*}Both authors are PhD students.

different real situations perform this type of process, e.g. molecules [12], bacteria [4], flock of birds [9], school of fish [42], or social insects [25], such as honeybees [40].

In such applicative scenarios, communication among agents is often affected by some form of noise. For this reason, one of the main goal in network information theory is to guarantee reliable communications in noisy networks [26]. In this context, error-correcting codes are very effective methods to reduce communication errors in computer systems [31, 37], 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 the (majority) consensus problem where the communication between entities is unreliable and subjected to uniform noise [16, 17, 23, 24].

The first consensus dynamics that have been studied in the presence of noise communication are linear opinion dynamics, such as the VOTER dynamics and the AVERAGING dynamics. In particular, they were studied in the presence of uniform noise communication [33] or in the presence of some communities of *stubborn agents* (i.e. agents that never change opinion) [35, 36, 46]. In these settings, only *metastable* forms of consensus can be achieved, where a large subset of the agents agree on some opinion while other opinions remain supported by smaller subsets of agents, and this setting lasts for a relatively-long time. However, the VOTER model has a slow convergence time even in fully connected networks and a large initial bias towards some majority opinion [28], and the AVERAGING dynamics requires agents to perform nontrivial 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 [11, 15, 22], and many research papers have begun to investigate new, more plausible, non-linear opinion dynamics.

To the best of our knowledge, the UNDECIDED-STATE dynamics is the first non-linear opinion dynamics analyzed in the presence of uniform communication noise [17]. The aforementioned dynamics exhibits a phase-transition which depends on the noise parameter, and a metastable phase of almost-consensus is quickly reached and kept for long time when the noise isn't too high. The UNDECIDED-STATE dynamics turns out to be a fast, very resilient dynamics, and this may explain why this type of process is adopted in some biological systems [40].

In this work, we consider the popular 3-MAJORITY dynamics, which is based on majority update-rules, the latter being widely employed also in the biological research field [13, 20]. In particular, we introduce in the system an uniform communication noise feature, following the definition of [17]. It has been proven that such dynamics, without communication noise, has a very similar behaviour to that of the UNDECIDED-STATE dynamics [6]. As we describe in the next section, the two dynamics behaves similarly (even if with crucial differences) even in presence of uniform noise, as both exhibit a phase transition. However, although 3-MAJORITY dynamics makes use of more per-round communications, it turns out to be less resilient to noise than UNDECIDED-STATE dynamics.

1.1 Our results and their consequences

In this work, we study the 3-MAJORITY dynamics over a network of n agents, which induces a process that works as follows: at the beginning, each agent holds an opinion from a set Σ ; at each subsequent discrete round, each agent pulls the opinions of three neighbor agents chosen independently uniformly at random and updates its opinion to the majority one, if there is any; otherwise, the agent adopts a random opinion among the sampled ones. This dynamics is a fast, robust protocol for the majority consensus problem in different network topologies (raging from

complete graphs to sparser graphs) [6]. For a discussion about the origin and previous results of the 3-MAJORITY dynamics we defer the reader to Section 1.2.

We consider the dynamics in the binary opinion case over the fully connected network. We introduce in the process an *uniform communication noise* feature, following the definition in [17] and for which we give an equivalent formulation: for each communication with a sampled neighbor, there is probability $p \in (0, 1)$ that it is noisy, i.e. the received opinion is sampled u.a.r. between the possible opinions. Instead, with probability 1 - p the communication is unaffected by noise. As shown in [17], this noise model (over the complete network) is equivalent to a model without any communication noise and where two communities of stubborn agents (that is, they never change opinion) of equal size $\frac{pn}{2(1-p)}$ are present in the network, where each of the two community holds a different opinion. Even though the complete graph is a strong assumption for such communication networks, we remark that, at every round, an agent pulls an opinion from three neighbors: therefore, the round-per-round communication pattern results is a dynamic graph with $\mathcal{O}(n)$ edges. Furthermore, such a model can be used to capture the behavior of bio-inspired multi-agent systems in which mobile agents meet randomly at a relatively high rate. For more details about models for bio-inspired swarms of agents, we refer to [43].

In the aforementioned setting, we prove that the process induced by the 3-MAJORITY dynamics exhibits a phase-transition. Our results are summarized in the following theorem.

Theorem. Let $\{s_t\}_{t\geq 0}$ be the bias of the process¹ induced by the 3-MAJORITY dynamics with uniform noise probability p. We prove the followings.

• If p < 1/3, let $s_0 = \Omega(\sqrt{n \log n})$ be the bias at the beginning of the process, $s_{eq} = \frac{n}{1-p}\sqrt{\frac{1-3p}{1-p}}$, 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.², 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 oscillates in such interval for $n^{\Theta(1)}$ rounds w.h.p.

• If p < 1/3, let $s_0 = \mathcal{O}(\sqrt{n \log n})$ be the bias at the beginning of the process. Then, there exists a time $\tau_2 = \mathcal{O}(\log n)$ such that, w.h.p., the system becomes unbalanced towards an opinion, i.e.

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

• If p > 1/3, let $s_0 = \Omega(\sqrt{n \log n})$ be the bias at the beginning of the process. Then, 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.

Our result shows that 3-MAJORITY dynamics is less resilient to noise than the UNDECIDED-STATE dynamics, despite in the 3-MAJORITY dynamics more communication per-round are allowed. Indeed, the phase transition for the UNDECIDED-STATE dynamics turns out to be at the threshold p = 1/2 [17],³ in the same setting as ours: since the threshold is higher than 1/3,

¹The bias s_t is the difference between the majority opinion community size and the minority opinion one at time t.

²An event holds with high probability (w.h.p. in short) with respect to n if the probability it occurs is at least $1 - n^{-\Theta(1)}$.

³In the cited work, an equivalent definition of noise model is given, and their formulation yields the threshold p = 1/6.

the dynamics is able to solve the consensus problem even in the presence of more noise than the 3-MAJORITY dynamics.

We briefly recall the UNDECIDED-STATE dynamics: at each round, each agent pulls a single neighbor opinion x u.a.r. If the agent former opinion y is different from x, the agent becomes undecided. If the agent is undecided, then it simply adopts any opinion it sees. This two-phases update-rule turns out to be more resilient to noise and, hence, a swarm of agents would benefit from it. In [17], the authors prove that the dynamics exhibits a similar phase transition for the noise probability p = 1/2. Below the threshold, the dynamics w.h.p. rapidly breaks the symmetry and converges in logarithmic time to a metastable phase of almost-consensus that lasts for polynomial time, in which the majority opinion exceeds the minority one by a bias of order of $\Theta(n)$. 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.

Nevertheless, we remark that our work shows technical novelties compared to [17]. A first difference lies in the fact that we find a precise equilibrium value s_{eq} that is attractive for the bias. Secondly, we characterize in detail what happens in the metastable almost-consensus phase: for every arbitrary small value $\varepsilon > 0$, we prove that the bias oscillates in the interval $[(1 - \varepsilon)s_{eq}, (1 + \varepsilon)s_{eq}]$ for polynomial time w.h.p. Instead, in [17] no precise equilibrium value is found, and in the metastable-almost consensus phase the bias lies in an interval of width $\Theta(n)$, without arbitrarily approaching an equilibrium state; nevertheless, we remark that we think the UNDECIDED-STATE process should behave in such a way.

On the other hand, when the noise probability is above the threshold 1/3, we prove than no form of consensus is possible w.h.p. as in [17], but we also show that the majority opinion switches every $\mathcal{O}(\log n)$ rounds with constant probability. In order to prove this, some drift analysis results with super-martingale arguments are used [32].

As future directions, sparser topologies are worth to be investigated. We believe that, as long as the communication graph shows strong connection properties, similar phase transitions will be exhibited. 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: in general, 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

Origin of the 3-Majority dynamics. The study of the 3-MAJORITY dynamics arises on the ground of the results obtained for the MEDIAN dynamics in [19]. 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 dynamics turns out to be a fault-taulerant, 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, 27].

In [7] it is shown that the 3-MAJORITY dynamics is a fast, fault-tolerant protocol for (valid) majority consensus in the case of $k \ge 2$ 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 [19], thus establishing its efficiency. In [8], the analysis is extended to any (even balanced) initial configuration in the many-color case, in the presence of an 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 that in [27]. The authors show that in the case of k opinions, the process converges in time $O(k \log n)$ rounds, and it is tight when $k = O(\sqrt{n})$. The 3-MAJORITY dynamics is also studied in different topologies: [29] 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.

Other popular non-linear opinion dynamics. Other important and efficient opinion dynamics for the majority consensus problem are the 2-CHOICES and the UNDECIDED-STATE dynamics. For an overview on the state of the art about opinion dynamics we defer the reader to [6]. We just want to quickly give the definitions of the 2-CHOICES dynamics (the UNDECIDED-STATE was already defined in the previous subsection). In the 2-CHOICES, 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 just want to remark that the expected per-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]. This is why mean-field arguments are sometimes not sufficient to analyze such processes. For example, we have ran 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 turns out to be an efficient majority consensus protocol, with the difference that it requires only one communication per round for each agent. Further description is 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.

Consensus dynamics in the presence of noise or stubborn agents. The authors of [45] initiate the study of the consensus problem in the presence of communication noise. They consider the Vicsek model [44], 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. In particular, as mentioned in the introduction, this direction is motivated, among many reasons, by the desire to find models for the consensus problem in natural phenomena [23].

The communication noise studied in this type of problem can be devided in two types: uniform (or unbiased) and non-uniform (or biased). The uniform case wants to capture errors in communications between agents in real-world scenarios. The non-uniform communication noise instead describes the case in which agents have a preferred opinion. The authors of [23] 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 [24] generalizes the above study to opinion sets of any cardinality.

As for the non-uniform communication noise case, in [16] it is considered the *h*-MAJORITY dynamics with a binary opinion set {ALPHA, BETA}, with a probability p that any received message is flipped towards a *fixed* preferred opinion, say 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^* (which depends on h), such that 1) if $p < p^*$, the time of disruption is at least polynomial, w.h.p., and 2) if $p > p^*$, the time of disruption is constant, w.h.p. Their result holds for any sufficiently dense graph. We remark that our work differs from [16] in that there is no preferred opinion, and the noise affecting communications may result in any possible opinion.

The noise feature affecting opinion dynamics 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 [17]. Hence, we discuss some previous works that consider such a model. In [46], the authors focus on the voter model, and show that the presence stubborn agents with opposite opinions precludes the convergence to consensus. The work [39] studies the asynchronous voter rule and the asynchronous majority rule dynamics with Poisson clocks, when the opinion set is binary using mean-field techniques, in presence of agents that either have a probability (which depends on their current opinion) not to update when the clock ticks, or are stubborn. In the second case, which directly relates with our work, they show that for the 3-MAJORITY dynamics there are either one or two possible stable equilibrium, depending on the sizes of the stubborn communities, which are reached in logarithmic time. If the two sizes are close between each other and not too large, then agreement on both opinions is possible in the steady state. Otherwise, either no agreement is possible, or the process converges to an agreement towards a single opinion (that of the largest stubborn community). This work includes the case in which the two stubborn communities have equal size, which corresponds to the uniform communication noise model. Nevertheless, we have some crucial differences: first of all, our work consider the synchronous version of the 3-MAJORITY dynamics, which cannot be analyzed with the same tools. Indeed, in each update round, the synchronous model has nonzero probability to reach any of the two monochromatic configurations. This feature is absent in the asynchronous version, since at each update, with probability equal to 1, at most one agent can change opinion. Furthermore, we want to remark that mean-field arguments do not capture important aspects of the process, such as metastability, which in [39] is shown only through simulations. For example, the actual behavior of the process in the long-term is oscillation 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.

1.3 Structure of the paper

Next section contains the preliminaries for the analysis and the result statements. Section 3 is devoted to the statements of the main theorems. In Section 4 we prove the theorems. Finally, in Appendix A we state some probabilistic tools we use throughout the analysis, and in Appendix B add some proofs due to space limitations.

2 Preliminaries

The 3-majority dynamics. Let G = (V, E) be a finite graph of n nodes (the agents), where each node is labelled uniquely with labels in $[n] := \{1, \ldots, n\}$. Furthermore, each node supports an opinion from a set of opinions Σ . The 3-MAJORITY dynamics defines a stochastic process $\{M_t\}_{t\in\mathbb{N}}$ which is described by the opinion of the nodes at each time step, i.e. $M_t = (i_1(t), \ldots, i_n(t)) \in \Sigma^n$ for every $t \ge 0$, 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 $t \ge 0$, 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 neighbors in G independently uniformly at random (with repetition) and updates its opinion to the majority one among the sampled neighbor opinions, if there is any. Otherwise, it adopts a random opinion among the sampled ones. For the sake of clarity, we remark that when u samples a neighbor node twice, the corresponding opinion counts twice.

Since M_t depends only on M_{t-1} , it follows that the process is a Markov chain. In the following, we will call the state of the process also by *configuration of the graph*.

The communication noise. We introduce an uniform communication noise feature in the dynamics, which is equivalent to that in [17]. Let 0 be a constant. When a node pulls a neighbor opinion, there is probability <math>p that the received opinion is sampled u.a.r. in Σ ; instead, with probability 1 - p, the former opinion keeps intact and is received.

3-Majority dynamics in the binary opinion case. The communication network we focus on is the complete graph $G = K_n$ with self loops in the binary opinion case, i.e. $\Sigma = \{\text{ALPHA}, \text{BETA}\}$. For the symmetry of the network, the state of the process is fully characterized by the number of nodes supporting a given opinion, which implies that the nodes do not require unique IDs. 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. Moreover, since at each time t, $a_t + b_t = n$, it suffices to know $\{b_t\}_{t>0}$ to fully describe the process.

We define the bias of the process at time t by

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

which takes value in $\{-n, \ldots, n\}$, and we notice that the process can also be characterized by the values of the bias alone, i.e. $\{s_t\}_{t\geq 0}$. We will use the latter sequence to refer to the process. 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. In the introduction, we took the bias to be $|s_t|$ but, for the sake of the analysis, we consider its *signed* version here. We finally remark that the random variable b_t (and, analogously, a_t) is the sum of i.i.d. Bernoulli r.v.s, which allows us to make use of the popular Chernoff bounds (Lemmas 16 and 19). In detail, if $X_i^{(t)}$ is the r.v. yielding 1 if nodes i adopts opinion BETA at round t + 1, and 0 otherwise, then $b_t = \sum_{i \in [n]} X_i^{(t)}$. Therefore, for (1),

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

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

Some 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.

3 Results

We here show our three main theorems. The first one shows how the dynamics solves the majority consensus problem when p < 1/3, even if in a "weak" form (since only an almost-consensus is reached). Section 4.1 is devoted to the proof of this theorem.

Theorem 1 (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 $\varepsilon^2 \le (1-3p)/2$) and let $\gamma > 0$ be any constant. Let $s_{eq} = \frac{n}{(1-p)}\sqrt{\frac{1-3p}{1-p}}$. Then, for any starting configuration s_0 such that $s_0 \ge \gamma \sqrt{n \log n}$ 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 $(1-\varepsilon)s_{eq} \leq s_{\tau_1} \leq (1+\varepsilon)s_{eq}$;
- (ii) there exists a value $c = \Theta_{\gamma,\varepsilon,p}(1)$ such that, for all $k \leq n^c$, $(1-\varepsilon)s_{eq} \leq s_{\tau_1+k} \leq (1+\varepsilon)s_{eq}$.

Our second theorem shows how the dynamics is capable of quickly breaking the initial symmetry. By applying also Theorem 1, it shows that the consensus problem is solved. The proof of the theorem is shown in Section 4.2.

Theorem 2 (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 positive 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 $|s_{\tau_2}| \geq \gamma \sqrt{n \log n}$.

Our last theorem shows that no form of consensus is possible when p > 1/3, and it is proved in Section 4.3.

Theorem 3 (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 $|s_0| \ge \gamma \sqrt{n \log n}$ 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 $|s_{\tau_3+k}| \leq \gamma \sqrt{n \log n}$.

4 Analysis

In this section we analyze the process. We first give some preliminary results. Afterwards, in Section 4.1 we prove Theorem 1, in Section 4.2 we prove Theorem 2, while Section 4.3 is devoted to the proof Theorem 3.

We now give the expectation of the bias at time t, conditional on its value at time t-1. Its proof can be found in Appendix B, and it is based on simple calculations.

Lemma 4. 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

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

By the lemma above, we deduce that there are up to three equilibrium configurations in expectation. 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

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

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} .

4.1 Victory of the majority

The aim of this subsection is to prove Theorem 1: so, in each statement 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 first show a lemma which states that, for any small constant $\varepsilon > 0$, whenever $s_{t-1} \notin [(1-\varepsilon)s_{eq}, (1+\varepsilon)s_{eq}]$, then s_t gets closer to the interval.

Lemma 5. For any constant $\varepsilon > 0$ such that $\varepsilon^2 < (1-3p)/2$ and for any $\gamma > 0$, if $s \ge \gamma \sqrt{n \log n}$, the followings hold

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

Proof. We first notice that

$$(1-\varepsilon)s_{\rm eq} \le \frac{n}{1-p}\sqrt{\frac{1-3p-2\varepsilon^2}{1-p}},\tag{4}$$

which holds since $\varepsilon^2 \leq (1-3p)/2$ and can be proved with simple calculations.

For Lemma 4, if each $s \leq (1 - \varepsilon)s_{eq}$, then

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

where the inequality follows from (4). Since (2), for the Hoeffding bound (Lemma 17), it holds that

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

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

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

for ε such that $\varepsilon^2 < (1 - 3p)/2$.

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

Lemma 6. For any constants $\varepsilon > 0$ and $\gamma > 0$, if $s \ge \gamma \sqrt{n \log n}$, the followings hold

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

Proof. The proof is similar to that of the previous lemma. From Lemma 4, we get that

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

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

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

The second claim comes symmetrically from Lemma 4 by observing that, since $s \ge (1 - \varepsilon)s_{eq}$

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

The Hoeffding bound implies

$$\mathbb{P}\left[s_t \ge s\left(1+\varepsilon\right) \mid s_{t-1}=s\right] \le \mathbb{P}\left[s_t \ge s\left(1+(1-3p)\varepsilon\right)+2p\varepsilon \cdot s \mid s_{t-1}=s\right] \le e^{-\gamma^2\varepsilon^2p^2} = \frac{1}{n^{\gamma^2\varepsilon^2p^2}}.$$

We provide another lemma to control the behavior of the bias. The proof is again deferred to Appendix \mathbf{B} , and consists in the application of simple concentration bounds.

Lemma 7. For any constant k > 0, the followings hold:

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

We can piece together the above lemmas, which imply the following corollary, whose proof consists in many calculations and is thus deferred to Appendix B.

Corollary 8. For any constant $\varepsilon > 0$ such that $\varepsilon < 1/3$ and $\varepsilon^2 < (1 - 3p)/2$, the followings hold:

(i) if $|s_{eq} - s| \le (\varepsilon/4)s_{eq}$, then

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

(ii) if $(\varepsilon/4)s_{eq} \le |s_{eq} - s| \le s_{eq}/3$, then

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

We are finally ready to prove the theorem.

Proof of Theorem 1. We divide the proof in different cases. First, suppose that $(\varepsilon/4)s_{eq} \leq |s_{eq} - s| \leq \varepsilon s_{eq}$. Let $T_1 = n^{\gamma^2 \varepsilon^4 p^2/(2^{19}3^2)}$. Then, from Corollary 8.(i) and (ii), for the chain rule, we have that

$$\mathbb{P}\left[\bigcap_{k=1}^{T}\left\{\left|s_{\mathrm{eq}} - s_{t+k}\right| \le \varepsilon s_{\mathrm{eq}}\right\} \ \middle| \ s_t = s\right] \ge 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{20} 3^2)}}.$$

Second, suppose that $\varepsilon s_{eq} \leq |s_{eq} - s| \leq s_{eq}/3$. Then, from Corollary 8.(ii), for the chain rule, a time T_2 exists, with

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

such that

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

Third, suppose that $s \leq 2s_{eq}/3$. From Lemma 5.(i) and Lemma 7.(ii), for the chain rule and the union bound, there is a time

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

such that

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

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

Fourth, suppose that $s \ge (1+\frac{1}{3})s_{eq}$. From Lemma 5.(ii) and Lemma 7.(i), for the chain rule, a time T_4 exists, with $T_4 = \mathcal{O}(\log n)$, such that

$$\mathbb{P}\left[|s_{\text{eq}} - s_{T_4}| \le s_{\text{eq}}/3 \mid s_t = s\right] \ge 1 - \frac{1}{n^{\gamma^2 3^4/2^6}}.$$

The theorem follows with $\tau_1 = \mathcal{O}(T_2 + T_3 + T_4)$

4.2 Symmetry breaking

The aim of this section is to prove Theorem 2: so, in each statement 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 [14]. We report the proof in Appendix B.

Lemma 9. 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,

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

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

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

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

Our goal is yo apply the above lemma to the 3-MAJORITY process, which defines a Markov chain. In particular, we claim the hypothesis of Lemma 9 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 9 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 10. For any constant $c_3 > 0$, let s be a value such that $|s| < c_3 \sqrt{n} \log n$. Then,

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

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

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

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

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

$$\mathbb{P}\left[|s_t| < h\sqrt{n} \mid s_{t-1} = s\right] \le \mathbb{P}\left[|s_t| < h\sqrt{n} \mid s_{t-1} = 0\right].$$
(5)

Thus, we can bound just the second probability, where the initial bias is zero. 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

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

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

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

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

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

By setting $c_1 = c/2$ and from Eq. (5) we get claim (i).

As for the second claim, assume s > 0 and $h\sqrt{n} \le s \le h\sqrt{n} \log n$. By Lemma 4 and the fact that $h\sqrt{n} \le s \le h\sqrt{n} \log n \le (1 - \sqrt{\varepsilon})s_{eq}$, we have (as in Lemma 5)

$$\mathbb{E}\left[s_t \mid s_{t-1} = s\right] = \frac{s(1-2p)}{2} \left(3 - \frac{s^2}{n^2}(1-2p)^2\right) \ge s\left(\frac{3}{2} - 3p - \frac{1-6p-2\varepsilon}{2}\right) = s(1+\varepsilon).$$

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

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

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

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

Proof of Theorem 2. Apply Lemmas 9 and 10 with $h = c_3 = \gamma$.

4.3 Victory of noise

In this subsection, we prove Theorem 3: so, in each statement, 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 make use of 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 constant probability that the majority opinion switches 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 $s \ge 2\sqrt{n}/(\varepsilon^2)$, the following holds

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

The proof can be found in Appendix B and makes use of some probabilistic inequalities such as the Jensen's one. 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 $s_0 \ge s_{min}$, with probability at least 1-1/n there exists a time $\tau = \mathcal{O}_{\varepsilon}(\log n)$ such that $|s_{\tau}| \le s_{min}$.

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

$$\mathbb{E} \left[X_t - X_{t-1} \mid X_{t-1} \ge s_{\min}, \mathcal{F}_{t-1} \right] = \mathbb{E} \left[X_t \mid X_{t-1} \ge s_{\min}, \mathcal{F}_{t-1} \right] - X_{t-1}$$

$$\stackrel{(a)}{\leq} \mathbb{E} \left[|s_t| \mid s_{t-1} \ge s_{\min}, \mathcal{F}_{t-1} \right] - s_{t-1} \stackrel{(b)}{\leq} \mathbb{E} \left[s_t \mid s_{t-1} \ge s_{\min}, \mathcal{F}_{t-1} \right] \cdot \left(1 + \frac{\varepsilon}{2} \right) - s_{t-1}$$

$$\stackrel{(c)}{\leq} s_{t-1} (1 - \varepsilon) \left(1 + \frac{\varepsilon}{2} \right) - s_{t-1} \le - \frac{\varepsilon \cdot s_{t-1}}{2},$$

where (a) holds because $X_t \leq |s_t|$, (b) holds for Lemma 11, and (c) holds for Lemma 4. Thus,

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

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

$$\begin{split} \mathbb{P}\left[\tau > t \mid s_{0}\right] &< \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \int_{s_{\min}}^{s_{0}} \frac{2}{\varepsilon \cdot y} \,\mathrm{d}y\right)\right] \leq \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \int_{s_{\min}}^{n} \frac{2}{\varepsilon \cdot y} \,\mathrm{d}y\right)\right] \\ &= \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \frac{2}{\varepsilon}(\log n - \log s_{\min})\right)\right] \\ &= \exp\left[-\frac{\varepsilon}{2} \cdot \left(t - \frac{2}{\varepsilon} - \frac{2}{\varepsilon}\left((\log n)/2 + 2\log\varepsilon\right)\right)\right] \\ &\leq \exp\left[-\frac{\varepsilon \cdot t}{2} + 1 + \frac{\log n}{2}\right]. \end{split}$$

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

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. Wlog we assume $s_{t-1} > 0$. Now, $s_{t-1} = b_{t-1} - a_{t-1}$, with $n/2 < b_{t-1} \le n/2 + \sqrt{n}/(2\varepsilon)$ and $n/2 - \sqrt{n}/(2\varepsilon) \le a_{t-1} < n/2$. Both b_{t-1} and a_{t-1} can be expressed as the sum of i.i.d. Bernoulli r.v.s. Since $\mathbb{E}[a_t \mid n/2 - \sqrt{n}/(2\varepsilon) \le a_{t-1} < n/2] \le n/2$, we have

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

where the latter inequality holds for 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.

Next lemma shows that the signed bias decreases each round. Its proof can be found in Appendix B.

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

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

We are ready to prove Theorem 3.

Proof of Theorem 3. Claim (i) follows directly from Lemmas 12 and 13. As for claim (ii), whenever the bias at some round $t = \tau + k$ becomes $|s_t| \ge (\gamma/2)\sqrt{n \log n}$, from Lemma 14.(ii) (and its symmetric statement), we have that $|s_t| \le \gamma \sqrt{n \log n}$ 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 $(\gamma/2)\sqrt{\log n}$. 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 $1 - 1/(2n^{\gamma^2 \varepsilon^2/2^8})$ for the chain rule.

References

- M. A. Abdullah and M. Draief. Global majority consensus by local majority polling on graphs of a given degree sequence. *Discrete Appl. Math.*, 180(C):1–10, Jan. 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. *CoRR*, abs/2103.10366, 2021.
- [4] B. L. Bassler. Small talk: Cell-to-cell communication in bacteria. Cell, 109(4):421–424, 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):58–104, 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 2017, Washington, DC, USA*, pages 335–344. ACM, 2017.
- [11] 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.
- [12] M. C. Carroll. The complement system in regulation of adaptive immunity. Nature Immunology, 5:981–986, 10 2004.
- [13] 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.

- [14] 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.
- [15] A. Condon, M. Hajiaghayi, D. Kirkpatrick, and J. Manuch. Approximate majority analyses using tri-molecular chemical reaction networks. *Natural Computing*, Aug. 2019.
- [16] E. Cruciani, H. A. Mimun, M. Quattropani, and S. Rizzo. Phase transitions of the kmajority dynamics in a biased communication model. In *International Conference on Distributed Computing and Networking 2021*, ICDCN '21, page 146–155, New York, NY, USA, 2021. Association for Computing Machinery.
- [17] 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, Paderborn, Germany, Proceedings, volume 12156 of Lecture Notes in Computer Science, pages 255–272. Springer, 2020.
- [18] M. Dietzfelbinger, A. Goerdt, M. Mitzenmacher, A. Montanari, R. Pagh, and M. Rink. Tight thresholds for cuckoo hashing via XORSAT. In S. Abramsky, C. Gavoille, C. Kirchner, F. M. auf der Heide, and P. G. Spirakis, editors, Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, Proceedings, Part I, volume 6198 of Lecture Notes in Computer Science, pages 213–225. Springer, 2010.
- [19] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. Stabilizing consensus with the power of two choices. SPAA '11, page 149–158, New York, NY, USA, 2011. Association for Computing Machinery.
- [20] J. Dong, D. Fernández-Baca, F. McMorris, and R. C. Powers. Majority-rule (+) consensus trees. *Mathematical Biosciences*, 228(1):10–15, 2010.
- [21] D. P. Dubhashi and A. Panconesi. Concentration of Measure for the Analysis of Randomized Algorithms. Cambridge University Press, 2009.
- [22] Emanuele Natale. On the Computational Power of Simple Dynamics. PhD Thesis, Sapienza University of Rome, 2017.
- [23] O. Feinerman, B. Haeupler, and A. Korman. Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication. *Distributed Comput.*, 30(5):339–355, 2017.
- [24] P. Fraigniaud and E. Natale. Noisy rumor spreading and plurality consensus. Distributed Comput., 32(4):257–276, 2019.
- [25] 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.
- [26] A. E. Gamal and Y. Kim. Network Information Theory. Cambridge University Press, 2011.
- [27] 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, Egham, United Kingdom, pages 305– 313. ACM, 2018.

- [28] Y. Hassin and D. Peleg. Distributed Probabilistic Polling and Applications to Proportionate Agreement. Information and Computation, 171(2):248–268, Dec. 2001.
- [29] N. Kang and N. Rivera. Best-of-three voting on dense graphs. SPAA '19, page 115–121, New York, NY, USA, 2019. Association for Computing Machinery.
- [30] 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.
- [31] R. Koetter and F. R. Kschischang. Coding for errors and erasures in random network coding. *IEEE Trans. Inf. Theory*, 54(8):3579–3591, 2008.
- [32] 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.
- [33] W. Lin, L. Zhixin, and G. Lei. Robust consensus of multi-agent systems with noise. In 2007 Chinese Control Conference, pages 737–741, 2007.
- [34] M. Mitzenmacher and E. Upfal. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. Cambridge University Press, 2005.
- [35] M. Mobilia. Does a Single Zealot Affect an Infinite Group of Voters? Physical Review Letters, 91(2):028701, July 2003.
- [36] 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.
- [37] T. K. Moon. Error Correction Coding: Mathematical Methods and Algorithms. Wiley-Interscience, USA, 2005.
- [38] 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.
- [39] A. Mukhopadhyay, R. Mazumdar, and R. Roy. Voter and majority dynamics with biased and stubborn agents. *Journal of Statistical Physics*, 181, 11 2020.
- [40] 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.
- [41] 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, Cancún, Mexico, pages 3613–3619. IEEE, 2008.
- [42] 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.
- [43] 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.
- [44] 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.
- [45] L. Wang and Z. Liu. Robust consensus of multi-agent systems with noise. Sci. China Ser. F Inf. Sci., 52(5):824–834, 2009.
- [46] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. Binary Opinion Dynamics with Stubborn Agents. ACM Trans. Econ. Comput., 1(4):19:1–19:30, Dec. 2013.

A Tools

We make use of the following general result on (super/sub)-martingales, which can be found in [32].

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 $x_{\min} \ge 0$. Let $h : [x_{\min}, x_{\max}] \to \mathbb{R}^+$ 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 $\mathbb{E}[X_t - X_{t+1} \mid X_t \ge x_{min}, \mathcal{F}_t] \ge h(X_t)$ and $\frac{\mathrm{d}}{\mathrm{d}x}h(x) \ge 0$, then $\mathbb{E}[T \mid X_0] \le \frac{x_{min}}{h(x_{min})} + \int_{x_{min}}^{X_0} \frac{1}{h(y)} \mathrm{d}y.$

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

$$\mathbb{E}\left[T \mid X_0\right] \ge \frac{x_{min}}{h(x_{min})} + \int_{x_{min}}^{X_0} \frac{1}{h(y)} \, \mathrm{d}y.$$

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

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

(iv) If $\mathbb{E}[X_t - X_{t+1} \mid X_t \ge x_{min}, \mathcal{F}_t] \le h(X_t)$ and $\frac{\mathrm{d}}{\mathrm{d}x}h(x) \le -\lambda$ for some $\lambda > 0$, then $\mathbb{P}[T < t \mid X_0] < \frac{e^{\lambda t} - e^{\lambda}}{e^{\lambda} - 1} \exp\left(-\lambda \left(\frac{x_{min}}{h(x_{min})} + \int_{x_{min}}^{X_0} \frac{1}{h(y)} \mathrm{d}y\right)\right).$

For an overview on the forms of Chernoff bounds see [21].

Lemma 16 (Multiplicative forms of Chernoff bounds). Let X_1, X_2, \ldots, 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

$$\mathbb{P}\left[X \ge (1+\delta)\mu_+\right] \le e^{-\frac{1}{3}\delta^2\mu_+};\tag{6}$$

(ii) for any $\delta \in (0,1)$ and $0 \le \mu_{-} \le \mu$, it holds that

$$\mathbb{P}[X \le (1-\delta)\mu_{-}] \le e^{-\frac{1}{2}\delta^{2}\mu_{-}}.$$
(7)

We also make use of the Hoeffding bound [34].

Lemma 17 (Hoeffding bounds). Let 0 < a < b be two constants. Let X_1, X_2, \ldots, 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 $\mu \le \mu_+$, it holds that

$$\mathbb{P}\left[X \ge \mu_+ + t\right] \le \exp\left(-\frac{2t^2}{n(b-a)^2}\right);\tag{8}$$

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

$$\mathbb{P}\left[X \le \mu_{-} - t\right] \le \exp\left(-\frac{2t^2}{n(b-a)^2}\right).$$
(9)

We make use of the following result which explicit the convergence "speed" in the centrallimit theorem.

Lemma 18 (Berry-Esseen). Let X_1, \ldots, 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

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

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 [30].

Lemma 19 (Reverse Chernoff bounds). Let X_1, X_2, \ldots, 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

$$\mathbb{P}\left[X \ge (1+\delta)\mu_+\right] \ge e^{-9\delta^2\mu_+};\tag{10}$$

(ii) for any $0 \le \mu_{-} \le \mu$, it holds that

$$\mathbb{P}\left[X \le (1-\delta)\mu_{-}\right] \ge e^{-9\delta^{2}\mu_{-}}.$$
(11)

B Missing proofs

Proof of Lemma 4. 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 $(b')^3 + 3a'(b')^2$. So, for Eq. (1), we have that

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

where the last equation follows from simple calculations.

B.1 Proofs: victory of the majority

Proof of Lemma 7. Let k be any arbitrarily large constant. As for (i), Lemma 4 gives that

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

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

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

where the latter inequality holds since $s_{eq} = \Theta(n)$ and $s_{eq} > (2k/\delta) \log n$ for a sufficiently large n. As for (ii), Lemma 4 implies that

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

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

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

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

Proof of Corollary 8. First, we prove (i). From Lemma 6 and the union bound, we have that

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

The claim follows by osberving that

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

As for claim (ii), we divide the proof in two different cases. Suppose, first, that $2s_{\rm eq}/3 \le s \le (1 - \varepsilon/4)s_{\rm eq}$. A constant $\varepsilon/4 \le \delta \le 1/3$ exists such that $s = (1 - \delta)s_{\rm eq}$. Then, from Lemmas 5 and 6

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

Notice that

$$\begin{aligned} \left| s_{\rm eq} - (1-\delta) \left(1 + \frac{3\varepsilon^2}{2^6} \right) s_{\rm eq} \right| &= s_{\rm eq} - (1-\delta) \left(1 + \frac{3\varepsilon^2}{2^6} \right) s_{\rm eq} \\ &= \left(s_{\rm eq} - (1-\delta) s_{\rm eq} \right) \cdot \left[1 - \frac{(1-\delta) \cdot \frac{3\varepsilon^2}{2^6} \cdot s_{\rm eq}}{\delta \cdot s_{\rm eq}} \right] \le \left(s_{\rm eq} - s \right) \cdot \left[1 - \frac{3\varepsilon^2}{2^5} \right] \,, \end{aligned}$$

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

$$\mathbb{P}\left[|s_{\rm eq} - s_t| \le |s_{\rm eq} - s| \cdot \left[1 - \frac{3\varepsilon^2}{2^5}\right] \ \middle| \ s_{t-1} = s\right] \ge 1 - \frac{1}{n^{\gamma^2 p^2 \varepsilon^4 / 2^{14}}}.$$
(12)

Second, suppose $(1 + \varepsilon/4)s_{eq} \le s \le 3s_{eq}/2$. Again, a constant $\varepsilon/4 \le \delta \le 1/3$ exists such that $s = (1 + \delta)s_{eq}$. From Lemmas 5 and 6, it holds that

$$\mathbb{P}\left[(1+\delta) \left(1 - \delta - \delta^2 \right) s_{\text{eq}} \le s_t \le (1+\delta) \left(1 - \frac{3\varepsilon^2}{2^6} \right) s_{\text{eq}} \ \left| \ s_{t-1} = s \right] \ge 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{18} 3^2)}} \,,$$

for the union bound, since $\delta \leq 1/3$. Notice that

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

where the inequality holds since $\delta \leq 1/3$. By simple calculations, it can be seen that $(1 + \delta)\left(1 - \frac{3\varepsilon^2}{2^6}\right) \geq 1$. Then, we have also that

$$\begin{aligned} \left| s_{\rm eq} - (1+\delta) \left(1 - \frac{3\varepsilon^2}{2^6} \right) s_{\rm eq} \right| &= (1+\delta) \left(1 - \frac{3\varepsilon^2}{2^6} \right) s_{\rm eq} - s_{\rm eq} \\ &= \left((1+\delta) s_{\rm eq} - s_{\rm eq} \right) \cdot \left[1 - \frac{(1+\delta) \cdot \frac{3\varepsilon^2}{2^6} s_{\rm eq}}{\delta s_{\rm eq}} \right] \stackrel{(a)}{\leq} \left((1+\delta) s_{\rm eq} - s_{\rm eq} \right) \cdot \left[1 - 3 \left(1 + \frac{\varepsilon}{4} \right) \cdot \frac{3\varepsilon^2}{2^6} \right] \\ \stackrel{(b)}{\leq} \left((1+\delta) s_{\rm eq} - s_{\rm eq} \right) \cdot \left[1 - \frac{9\varepsilon^2}{2^6} \right] , \end{aligned}$$

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

$$\mathbb{P}\left[|s_{\rm eq} - s_t| \le |s_{\rm eq} - s| \cdot \left[1 - \frac{9\varepsilon^2}{2^6}\right] \ \middle| \ s_{t-1} = s\right] \ge 1 - \frac{1}{n^{\gamma^2 \varepsilon^4 p^2 / (2^{18} 3^2)}}.$$
(13)

Combining Eqs. (12) and (13), we get that, whenever $(\varepsilon/4)s_{eq} \leq |s_{eq} - s| \leq s_{eq}/3$, then

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

B.2 Proofs: symmetry-breaking

Proof of Lemma 9. Define a set of hitting times $T := \{\tau(i)\}_{i \in \mathbb{N}}$, where

$$\tau(i) = \inf_{i \in \mathbb{N}} \left\{ t : t > \tau(i-1), f(X_t) \ge h\sqrt{n} \right\},\$$

setting $\tau(0) = 0$. By the first hypothesis, for every $i \in \mathbb{N}$, the expectation of $\tau(i)$ is finite. Now, define the following stochastic process which is a subsequence of $\{X_t\}_{t\in\mathbb{N}}$:

$$\{R_i\}_{i\in\mathbb{N}} = \{X_{\tau(i)}\}_{i\in\mathbb{N}}.$$

Observe that $\{R_i\}_{i\in\mathbb{N}}$ is still a Markov chain. Indeed, if $\{x_1,\ldots,X_{i-1}\}$ be a set of states in Ω , then

$$\mathbb{P} \left[R_i = x \mid R_{i-1} = x_{i-1}, \dots, R_1 = x_1 \right]$$

$$= \mathbb{P} \left[X_{\tau(i)} = x \mid X_{\tau(i-1)} = x_{i-1}, \dots, X_{\tau(1)} = x_1 \right]$$

$$= \sum_{t(i) > \dots > t(1) \in \mathbb{N}} \mathbb{P} \left[X_{t(i)} = x \mid X_{t(i-1)} = x_{i-1}, \dots, X_{t(1)} = x_1 \right]$$

$$\cdot \mathbb{P} \left[\tau(i) = t(i), \dots, \tau(1) = t(1) \right]$$

$$= \mathbb{P} \left[X_{\tau(i)} = x \mid X_{\tau(i-1) = x_{i-1}} \right]$$

$$= \mathbb{P} \left[R_i = x \mid R_{i-1} = x_{i-1} \right] .$$

By definition, the state space of R is $\{x \in \Omega : f(x) \ge h\sqrt{n}\}$. Moreover, the second hypothesis still holds for this new Markov chain. Indeed:

$$\mathbb{P}\left[f(R_{i+1}) < (1+\varepsilon)f(R_i) \mid R_i = x\right]$$

= $1 - \mathbb{P}\left[f(R_{i+1}) \ge (1+\varepsilon)f(R_i) \mid R_i = x\right]$
= $1 - \mathbb{P}\left[f(X_{\tau(i+1)} \ge (1+\varepsilon)f(X_{\tau(i)}) \mid X_{\tau(i)} = x\right]$
 $\le 1 - \mathbb{P}\left[f(X_{\tau(i+1)} \ge (1+\varepsilon)f(X_{\tau(i)}), \tau(i+1) = \tau(i) + 1 \mid X_{\tau(i)} = x\right]$

$$= 1 - \mathbb{P}\left[f(X_{\tau(i)+1} \ge (1+\varepsilon)f(X_{\tau(i)}) \mid X_{\tau(i)} = x\right]$$

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

These two properties are sufficient to study the number of rounds required by the new Markov chain $\{R_i\}_{i \in \mathbb{N}}$ to reach the target value m. Indeed, by defining the random variable $Z_i = \frac{f(R_i)}{\sqrt{n}}$, and considering the following "potential" function, $Y_i = exp\left(\frac{m}{\sqrt{n}} - Z_i\right)$, we can compute its expectation at the next round as follows. Let us fix any state $x \in \Omega$ such that $h\sqrt{n} \leq f(x) < m$, and define $z = \frac{f(x)}{\sqrt{n}}$, $y = exp\left(\frac{m}{\sqrt{n}} - z\right)$. We have

$$\mathbb{E}\left[Y_{i+1}|R_{i}=x\right] \leq \mathbb{P}\left[f(R_{i+1}) < (1+\varepsilon)f(x)\right]e^{m/\sqrt{n}} \\ + \mathbb{P}\left[f(R_{i+1}) \geq (1+\varepsilon)f(x)\right]e^{m/\sqrt{n}-(1+\varepsilon)z} \\ \text{(from hypothesis (ii))} \leq e^{-c_{2}z^{2}} \cdot e^{m/\sqrt{n}} + 1 \cdot e^{m/\sqrt{n}-(1+\varepsilon)z} \\ = e^{m/\sqrt{n}-c_{2}z^{2}} + e^{m/\sqrt{n}-z-\varepsilon z} \\ = e^{m/\sqrt{n}-z}(e^{z-c_{2}z^{2}} + e^{-\varepsilon z}) \\ \leq e^{m/\sqrt{n}-z}(e^{-2} + e^{-2}) \\ \leq e^{m/\sqrt{n}-z} \\ = \frac{y}{e}, \tag{14}$$

where in (14) we used that z is always at least h and thanks to hypothesis (i) we can choose a sufficiently large h, which depends on c_2 and ε .

By applying the Markov inequality and iterating the above bound, we get

$$\mathbb{P}\left[Y_i > 1\right] \le \frac{\mathbb{E}\left[Y_i\right]}{1} \le \frac{\mathbb{E}\left[Y_{i-1}\right]}{e} \le \dots \le \frac{\mathbb{E}\left[Y_0\right]}{e^{\tau_R}} \le \frac{e^{m/\sqrt{n}}}{e^i}.$$

We observe that if $Y_i \leq 1$ then $R_i \geq m$, thus by setting $i = m/\sqrt{n} + \log n = (c_3 + 1) \log n$, we get:

$$\mathbb{P}\left[R_{(c_3+1)\log n} < m\right] = \mathbb{P}\left[Y_{(c_3+1)\log n} > 1\right] < \frac{1}{n}.$$
(15)

Our next goal is to give an upper bound on the hitting time $\tau_{(c_3+1)\log n}$. Note that the event " $\tau((c_3+1)\log n) > c_4\log n$ " holds if and only if the number of rounds such that $f(X_t) \ge h\sqrt{n}$ (before round $c_4\log n$) is less than $(c_3+1)\log n$. Thanks to Hypothesis (1), at each round t there is at least probability $1 - c_1$ that $f(X_t) \ge h\sqrt{n}$. This implies that, for any positive constant c_4 , the probability $\mathbb{P}[\tau((c_3+1)\log n) > c_4\log n]$ is bounded by the probability that, within $c_4\log n$ independent Bernoulli trials, we get less then $(c_3+1)\log n$ successes, where the success probability is at least $1 - c_1$. We can thus choose a sufficiently large $c_4 = c_4(c_1, c_3)$ and apply the multiplicative form of the Chernoff bound (Lemma 16), obtaining

$$\mathbb{P}\left[\tau\left((c_3+1)\log n\right) > c_4\log n\right] < \frac{1}{n}.$$
(16)

We are now ready to prove the Lemma using (15) and (16), indeed

$$\mathbb{P}[X_{c_4 \log n} \ge m] > \mathbb{P}[R_{(c_3+1)\log n} \ge m \land \tau ((c_3+1)\log n) \le c_4 \log n] \\= 1 - \mathbb{P}[R_{(c_3+1)\log n} < m \lor \tau ((c_3+1)\log n) > c_4\log n]$$

$$\geq 1 - \mathbb{P}\left[R_{(c_3+1)\log n} < m\right] + \mathbb{P}\left[\tau\left((c_3+1)\log n\right) > c_4\log n\right]$$
$$> 1 - \frac{2}{n}.$$

Hence, choosing a suitable large c_4 , we have shown that in $c_4 \log n$ rounds the process reaches the target value m, w.h.p.

B.3 Proofs: victory of noise

Proof of Lemma 11. Trivially, it holds that

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

Furthermore, from Lemma 4, we have that $\mathbb{E}[s_t \mid s_{t-1} = s] \ge 0$ as long as $s \ge 0$. By writing

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

and by using the Jensen's inequality for a concave function (i.e. the square root), it follows that

$$\mathbb{E}\left[|s_{t}| \mid s_{t-1} = s\right] \leq \sqrt{\mathbb{E}\left[\left(s_{t} - \mathbb{E}\left[s_{t} \mid s_{t-1} = s\right]\right)^{2} \mid s_{t-1} = s\right]} + \mathbb{E}\left[s_{t} \mid s_{t-1} = s\right]$$

= $\sigma\left(|s_{t}| \mid s_{t-1} = s\right) + \mathbb{E}\left[s_{t} \mid s_{t-1} = s\right],$ (17)

where $\sigma(x)$ represents the standard deviation of a r.v. x. As pointed out in the preliminaries (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:

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

where the latter inequality holds since $\sigma \left(Y_i^{(t)} \mid s_{t-1} = s\right)^2 \leq 1$ for every *i*. Furthermore, from Lemma 4, we deduce that

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

Since $s \ge \frac{2\sqrt{n}}{\varepsilon^2} \ge \frac{2\sqrt{n}}{\varepsilon(1-p)}$, we get that $\mathbb{E}[s_t \mid s_{t-1} = s] \ge \frac{2\sqrt{n}}{\varepsilon}$. By using the latter facts in Eq. (17), we obtain

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

Proof of Lemma 14. From Lemma 4, for each $s \ge 0$ it holds that

$$\mathbb{E}\left[s_t \mid s_{t-1} = s\right] \le \frac{3s(1-p)}{2} \le (1-\varepsilon)s,\tag{18}$$

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

$$\mathbb{P}\left[s_t \ge (1-\varepsilon)s + \varepsilon \cdot s/4\right] \le e^{-s^2\varepsilon^2/(32n)} \le e^{-\gamma^2\varepsilon^2\log n/2^7} \le \frac{1}{n^{\frac{\gamma^2\varepsilon^2}{27}}}.$$

As for the second claim, we notice that, from Eq. (18), $\mathbb{E}[s_t | s_{t-1} = s] \leq s$. The Hoeffding bound (Lemma 17) now implies that

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

Moreover, from Lemma 4, for any $0 \le s \le n$, $\mathbb{E}[s_t \mid s_{t-1} = s] \ge 0$. Applying again the Hoeffding bound, we get that

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

For the union bound, we get the second claim.