Consensus Problems in Population Protocol Model

Dissertation zur Erlangung des akademischen Grades
Dr. rer. nat.

an der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Informatik
Universität Hamburg

Betreut von: Prof. Petra Berenbrink

eingereicht von

Felix Biermeier

Hamburg 2024

Date of Disputation: 25.04.2025

Committee:

1st Examiner: Prof. Dr. Petra Berenbrink, University of Hamburg

2nd Examiner: Prof. Dr. Sören Laue, University of Hamburg

Chair: Prof. Dr. Peter Kling, Darmstadt University of Applied Sciences

Abstract

This thesis considers plurality consensus problems in the population protocol model. These are fundamental problems in distributed computing. We design and analyze protocols to solve these consensus problems time and space efficiently. In the population protocol model, we consider n agents that interact in randomly chosen pairs, one pair per time step. In the *plurality consensus* problem, each of the n agents initially has one of the k opinions. The goal is for all agents to agree on the initial most frequent opinion among the population. We consider both exact and approximate plurality consensus. This refers to the initial bias defined as the difference between the largest and second largest opinion. In the exact case, the initial largest opinion must win eventually. In the approximate case, only with a sufficiently large initial bias. We consider a synchronized variant of the Undecided State Dynamics in the population protocol model. Our main result solves approximate plurality consensus for $k = \mathrm{O}\big(\sqrt{n}/\log n\big)$ in $\mathrm{O}\big(n \log^2 n\big)$ interactions using $k \cdot \mathrm{O}\big(\log n\big)$ states w.h.p. In the presence of an initial additive (multiplicative) bias, all agents agree on the initial largest opinion, and we lose the constraint on k and improve the convergence time. We solve the approximate plurality consensus for $k > 2$ via the (unsynchronized) Undecided State Dynamics in the population protocols model. In fact, we study this protocol for the first time in this regime and model. The main result solves approximate plurality consensus in $\mathrm{O}\big(n^2/x_{\max} \log n\big)$ interactions w.h.p. under a mild constraint on the initial largest opinion. Similar holds under the assumption of an initial additive (multiplicative) bias. We design and analyze tournament-based protocols for exact plurality consensus that beat the quadratic lower bound on the states by allowing a negligible failure probability. Our main result solves exact plurality w.h.p. in $\mathrm{O}\big(n^2/x_{\max} \cdot \log n + n \cdot \log^2 n\big)$ interactions and $\mathrm{O}\big(k \cdot \log\log n + \log n\big)$ states under a mild constraint on the initial largest opinion. We solve a parameterized variant of the approximate majority consensus problem with a preferred opinion 1 and unpreferred opinion 2. Our main result shows a phase transition concerning the stubbornness parameter p. If the stubbornness is sufficiently larger (smaller) than the threshold $1 - x_1/x_2$, all agents agree on the preferred (unpreferred) opinion in $\mathrm{O}\big(n \log n\big)$ interactions w.h.p. Otherwise, all agents agree on the same opinion in $\mathrm{O}\big(n \log^2 n\big)$ interactions w.h.p.

Zusammenfassung

In dieser Thesis geht es um Pluralitäts-Konsensus Probleme im Populationsprotokoll-Modell. Dies sind fundamentale Probleme im verteilten Rechnen. Wir konstruieren und analysieren Protokolle, welche diese Probleme zeit- und platzeffizient lösen. In diesem Modell betrachten wir n Agenten, welche zufällig paarweise miteinander interagieren, jeweils ein Paar pro Zeitschritt. Beim Pluralitäts-Konsensus Problem hat jeder der n Agenten initial eine von k Meinungen. Das Ziel ist es, dass sich alle Agenten auf die initial am häufigsten auftretende Meinung einigen. Wir betrachten jeweils die exakte und approximierte Variante des Pluralitäts-Konsensus Problems. Diese Varianten beziehen sich auf den initialen Bias, welcher als Differenz zwischen der größten und zweitgrößten Meinung definiert ist. In der exakten Variante muss die initial größte Meinung letztendlich gewinnen. In der approximierten Variante muss dies nur erfüllt werden, falls der initiale Bias ausreichend groß ist.

Wir betrachten die synchronisierte Variante der Undecided State Dynamics im Populationsprotokoll-Modell. Unser Hauptresultat löst das approximierte Pluralitäts-Konsensus Problem für $k = \mathrm{O}\big(\sqrt{n}/\log n\big)$ in $\mathrm{O}\big(n \log^2 n\big)$ Interaktionen und $k \cdot \mathrm{O}\big(\log n\big)$ Zuständen mit hoher Wahrscheinlichkeit. Falls ein initialer additiver (multiplikative) Bias vorliegt, dann nehmen alle Agenten die initial größte Meinung an. Zudem ist k unbeschränkt und die Konvergenzzeit verbessert sich. Wir lösen das approximierte Pluralitäts-Konsensus Problem für $k > 2$ mit Hilfe der (unsynchronisierten) Undecided State Dynamics im Populationsprotokoll-Model. In der Tat untersuchen wir dieses Protokoll als Erstes in dieser Konstellation. Unser Hauptresultat löst das approximierte Pluralitäts-Konsensus Problem für $k > 2$ in $\mathrm{O}\big(n^2/x_{\max} \log n\big)$ Interaktionen mit hoher Wahrscheinlichkeit unter einer schwachen Annahme über die initial größte Meinung. Ähnliches gilt auch bei der Annahme eines initialen additiven (multiplikativen) Bias. Wir konstruieren und analysieren turnierbasierte Protokolle für das exakte Pluralität-Konsensus Problem, welche die quadratische untere Schranke der Zustände mit dem Zulassen von vernachlässigbarer Fehlerwahrscheinlichkeit unterbietet. Unser Hauptresultatt löst das exakte Pluralität-Konsensus Problem mit hoher Wahrscheinlichkeit in $\mathrm{O}\big(n^2/x_{\max} \cdot \log n + n \cdot \log^2 n\big)$ Interaktionen und $\mathrm{O}\big(k \cdot \log\log n + \log n\big)$ Zuständen unter einer schwachen Annahme über die initial größte Meinung. Wir lösen eine parametrisierte Variante des approximierten Majorität-Konsensus Problems mit einer bevorzugten Meinung 1 und einer unbevorzugten Meinung 2. Unser Hauptresultät weist einen Phasenübergang bezüglich des Hartnäckigkeitsparameters p. Falls dieser Parameter wesentlich größer (kleiner) als der Schwellenwert $1-x_1/x_2$ ausfällt, dann werden alle Agenten die bevorzugte (unbevorzugte) Meinung in $\mathrm{O}\big(n \log n\big)$ Interaktionen mit hoher Wahrscheinlichkeit annehmen. Andernfalls nehmen alle Agenten dieselbe Meinung in $\mathrm{O}\big(n \log^2 n\big)$ Interaktionen mit hoher Wahrscheinlichkeit an.

Acknowledgements

I sincerely thank my advisor, Petra Berenbrink, for the opportunity to work in this wonderful research area and pursue my PhD. I am really grateful for her guidance and support. Furthermore, I thank all my colleagues at the university and colleagues in the community for a great time and for having nice discussions about research and other topics. Especially the weekly cake seminars were very enjoyable. Finally, I thank my family and friends for their support.

Contents

Chapter 1.

Introduction

A distributed system consists of a set of devices that are connected by a communication network. The communication network provides the facility for information exchange among devices. The devices usually cannot access a shared (global) memory and can only communicate by passing messages over the communication network. This includes the observation of another device's state. Note that a shared memory model is a completely different approach towards distributed systems. In contrast to the message-passing approach, where the devices communicate directly with each other, communication in this model occurs indirectly via a shared memory.

The devices have no access to a physical global clock. Inherent to that is usually asynchronous communication among the devices. Each device works on its own but serves a common goal with the other devices. Roughly speaking, every device has its own state and a local view of the global problem by gathering information through communication. All device states as a whole unit form a configuration. Note that synchronous communication works in a similar way, except all devices simultaneously perform their actions.

Given a distributed system, the characteristic of such a system is to perform distributed computing. That is, to solve a problem based on the action and information of the whole collection of devices. There are various complexity measures of common interest. Naturally, time and space complexity have to be considered. Additionally, the number and size of messages and the number of faulty devices are relevant. Further information about distributed systems and distributed computing can be found in the books by Attiya and Welch [14] and Kshemkalyani and Singhal [62]. In this thesis, we focus strongly on consensus problems as one of the field's fundamental problems (see Section 1.2 for more details). To do that, we first formally introduce the population protocol model proposed by Angluin et al. [8] as a realization of distributed computing in a distributed system.

1.1. Population protocol model

A population protocol consists of a collection of agents, the so-called population, that interact with each other to compute a function over the whole population.

A way to think of an agent is a finite state machine. The agents work on a common state space $\mathbb{Q}$, and agents may change their state by interacting with another agent. Agents are rather limited in their storage and computational power for simplicity and robustness. Additionally, agents are anonymous and indistinguishable. In the real world, you can think of mobile ad hoc networks or sensor networks.

In general, given a certain problem, designing a population protocol that solves the problem eventually among the population. An output function provides an explicit solution while reaching the desired configuration, a more implicit solution. Note that a population protocol does not necessarily terminate but rather converges since the agents do not necessarily detect the end of the computation. Formally, such protocol consists of

- the state space $\mathbb{Q}$
- initial states of each agent
- output function, mapping from state to output value
- transition function $\delta : \mathbb{Q} \times \mathbb{Q} \mapsto \mathbb{Q} \times \mathbb{Q}$

An interaction involves an initiator agent u and responder agent v. Such an interaction allows both agents to view each other's state and update their own state according to the transition function. The transition function is usually deterministic and represented explicitly by state transitions or implicitly by an algorithmic protocol. A configuration C represents the current state of the agents in the population. It represents essentially the global state of the population that is not known to any of the agents. An execution of a population protocol is a sequence of configurations $C_0, C_1, \ldots$ where the next configuration is obtained by its previous one updated according to a single interaction,i.e., a sequence of interactions provides the changes in the sequence of configurations.

A scheduler provides the sequence of interactions. Originally, the schedule is given by an adversary under a certain fairness condition. The fairness condition forces the adversary to schedule interactions between two agents so that, eventually, every reachable configuration may appear.

Originally, the focus was on the spectrum of computable functions by population protocols. In fact, it suffices to consider a function's so-called predicate, i.e., a boolean-valued verifying function. It turns out that only semilinear predicates are computable in the original basic model as shown by Angluin et al. [11]. Essentially, functions that are not representable by the so-called Presburger arithmetic are not computable. This caused the introduction of the random scheduler as an alternative. Now, pairwise interactions between two agents are chosen uniformly at random. This leads to a natural notion of time by examining the number of interactions. Before that, the adversarial scheduler may cause long sequences of nonproductive interactions that do not lead to any progress. In fact, the random scheduler makes the model more powerful. For example, it allows the simulation of register machines. From here on, more complex problems like leader election and consensus are considered. The latter will be investigated throughout the remainder of the thesis. A formal definition definition follows in the next section.

The natural measurements of the quality of a population protocol are based on its space and time complexity: the number of states per agent and the number of interactions to reach the desired configuration. This leads to extensive trade-off discussions between space and time. Space efficiency is relevant in its own right in the context of population protocols using simple and robust agents.

1.2. Overview of the Problems

We study variants of the plurality consensus problem in the population protocols model. It is one of the fundamental problems in distributed computing. We are given a population of n agents connected via a complete graph and a set of k opinions. Initially, every agent has one of the k opinions. The goal is that all agents eventually agree on the same opinion, preferably on the initial largest opinion.

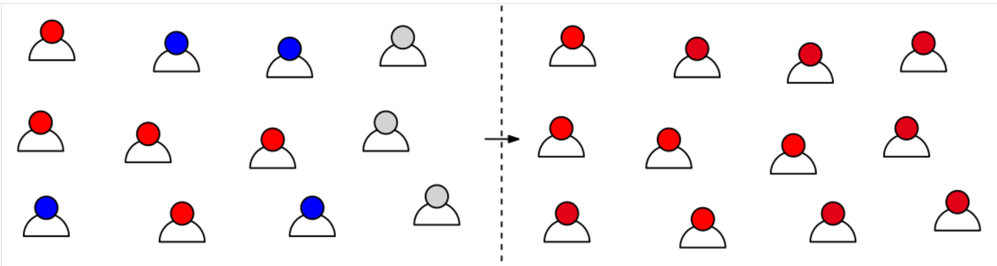

Figure 1.1.: Approximate Plurality Problem

For the record, plurality refers to the number of opinions larger than two. Otherwise, it is called the majority consensus problem for $k = 2$.

In Practical applications, it plays an important role in fault-tolerant sensor networks (where most sensors must confirm a trustworthy result) or majority-based conflict resolution (e.g., for CRCW PRAMs). They are also used in physics and biology to model massive dynamic systems of particles or bacteria or in social sciences to study how opinions form and spread through social interactions. See [19] for references and further applications. On an intuitive level, the initial larger opinions are supposed to be more likely to win than the smaller ones. This leads to the common notion of bias. In the literature, there is usually a distinction between additive and multiplicative bias. An additive bias is the difference between the support of the largest and second-largest opinion. The multiplicative bias refers to the ratio between the support of the largest and second-largest opinion. Regarding that notion, there are two popular forms of consensus problems: exact plurality and approximate plurality. In the first case, the initial largest opinion must win eventually, regardless of the initial (additive) bias. In the other case, the initial largest opinion is only required to win under the assumption that the initial bias is sufficiently large, like $\Omega\left(\sqrt{n \log n}\right)$. Otherwise, it is sufficient to reach a consensus, i.e., eventually, all agents agree on the same opinion.

Regardless of the consensus variant, the core aspect of a protocol solving such a problem is its convergence time, i.e., how long does it take until all agents agree on the same opinion? In particular, it is important to define and identify proper measurements for the convergence time. The literature has two rather obvious and established choices: the number of initial opinions k and the initial bias. The latter is more subtle and hidden in the convergence time and rather noticeable in a direct comparison with an unbiased case; i.e., a bias speeds up the convergence time of most protocols. Dealing with the number of opinions k is usually one of the main challenges in the analysis. There exist lots of configurations that are often not properly captured by these measurements alone. Consider the following example in Fig. 1.2. Both configurations have the same number of opinions, and the additive bias is roughly the same. Intuitively, the first configuration converges faster than the second one. The reason is that in the skewed configuration, almost all opinions are relatively small compared to the largest opinion. On the other hand, it usually requires non-trivial approaches and techniques to deal with larger subsets of opinions at once instead of one after the other.

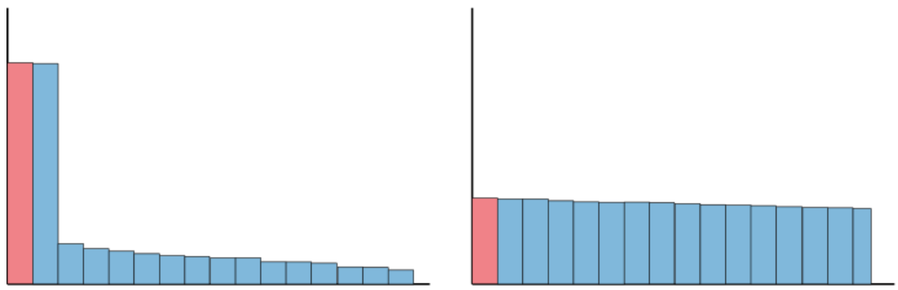

Figure 1.2.: Skewed and almost uniform configurations

Another quantity to consider is the space complexity of protocols. In the population protocol model, we are looking at the state space complexity. Actually, the state overhead of an agent is more interesting in this setting, i.e., the information in addition to storing one of the mandatory opinions. A typical example is an additional synchronization tool, especially in an asynchronous model like the population protocol model.

At last, we also consider a variant of the approximate majority consensus problem with preference. It is a form of biased opinion dynamics in that one opinion is preferable. Initially, each agent holds one of two possible opinions (1 and 2 in the following), but Opinion 1 is the preferred one. The notion of preference is rather abstract. There are various ways to model preference. Usually, the preferred opinion is more powerful than the unpreferred opinion, e.g., agents with the preferred opinion are more resilient in maintaining their opinion, or agents adopt the preferred opinion more likely than another opinion. The primary goal remains that all agents eventually agree on the same opinion. In contrast to the classical (approximate) majority consensus problem, the focus is not necessarily on the initial largest opinion but rather the preferred opinion. In this sense, how the disadvantage of initial support being the minority can be balanced out by preference and eventually even lead to being the winning opinion.

1.3. Undecided State Dynamics

The Undecided State Dynamics (USD) is a popular protocol for solving consensus problems in various distributed models. It plays a crucial role throughout the remainder of this thesis. Given its suitability as a primitive for other distributed tasks, a substantial amount of recent work has analyzed this process as a protocol for *consensus* under varying settings. The USD was originally introduced by Angluin et al. [9] for $k=2$ opinions in the *population protocol* model. Independently, Perron et al. [75] analyzed the two opinions USD in the *asynchronous gossip* model of Boyd et al. [33], which can be viewed as the continuous-time variant of the population protocol model. It belongs to the classes of protocols using additional information in the form of an extra state, the so-called undecided state. The undecided state can be viewed as an additional (special) opinion. The protocol in the population protocol model is rather simple. There are two types of state changes (see Fig. 1.3):

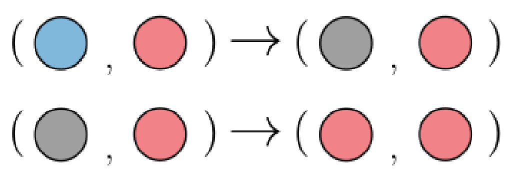

Figure 1.3.: Undecided State Dynamics: Pairwise interactions between agents with opposite opinions (blue and red) and undecided agent (grey)

First, whenever two agents with opposite opinions interact, the initiator loses its opinion and becomes undecided, i.e., it temporally does not have any opinion. Second, whenever the initiator is an undecided agent, it adopts the responder's opinion. In a certain sense, this resembles a delayed Voter process. Recall that the Voter rule states that the initiator adopts the responder's opinion regardless of its own opinion. In this case, adopting another opinion is delayed by becoming undecided first.

A couple of variants of the original USD are introduced in the literature. One of them can be viewed as a synchronized version introduced by Berenbrink et al. [29] and Ghaffari and Parter [53] in the gossip model. The gossip model is a synchronous time model, where in each (parallel) round, every agent selects uniformly at random interaction partners. In the synchronized USD, a sequence of (parallel) rounds forms a phase of two parts. In the first part, an agent becomes undecided if its interaction partner has a different opinion. In the second part, an undecided agent adopts the opinion of its interaction partner. Note that the first part only lasts a single round due to the gossip model. Essentially, if an agent loses its opinion in the first part, it tries to recover an opinion during the remainder of a phase.

1.4. Results

We investigate plurality consensus problems in the population protocol model. In Chapter 2, we consider a synchronized variant of the Undecided State Dynamics in the population protocol model. This variant has been introduced by Berenbrink et al. [29] and Ghaffari and Parter [53] in the gossip model. Both works provide a polylogarithmic time bound on the number of parallel rounds solving plurality majority for $k > 2$ and initial additive bias of order $\Omega\big(\sqrt{n \log n}\big)$. In contrast to the gossip model, we need an additional tool, so-called phase clocks, to properly synchronize the population. We solve the approximate plurality consensus problem for $k = \mathrm{O}\big(\sqrt{n}/\log n\big)$ with no assumption on the bias in $\mathrm{O}\big(n \log^2 n\big)$ interactions using $k \cdot \mathrm{O}\big(\log n\big)$ states w.h.p. Furthermore, despite the absence of any bias, we can restrict the set of contenders from being the winning opinion. We call such opinions significant. Roughly speaking, a significant opinion is almost as large as the largest opinion except for an additive difference of order $\mathrm{o}\big(\sqrt{n \log n}\big)$. In the spirit of approximate plurality, all agents agree only on one of the significant opinions. In addition, we provide improved runtimes in case the initial configuration has a sufficiently large additive or multiplicative bias (replacing a $\log n$ by $\log\log n$) for the whole range of k.

In Chapter 3, we solve the plurality consensus problem for $k > 2$ opinions via the (unsynchronized) Undecided State Dynamics in the population protocols model. Surprisingly, in contrast to the majority problem, the $k > 2$ regime has not been studied under these circumstances. Under mild constraints on the initial largest opinion and number of undecided agents, we show that the Undecided State Dynamics solves plurality consensus for one of the initial significant opinions in $\mathrm{O}\big(n^2/x_{\max} \log n\big)$ interactions w.h.p. Furthermore, all agents agree on the initial largest opinion in $\mathrm{O}\big(n^2/x_{\max} \log n\big)$ interactions w.h.p. if there exists an additive bias of order $\Omega\big(\sqrt{n} \log n\big)$. In case of a constant $(1+\varepsilon)$ multiplicative bias, the runtime improves to $\mathrm{O}\big(n^2/x_{\max} + n \log n\big)$. The runtime dependency on the support of the initial largest opinion provides a more detailed perspective on the convergence time than solely relying on the initial number of opinions k.

In Chapter 4, we consider the exact plurality consensus problem in the population protocol model. In contrast to the approximate plurality problem, the initial largest opinion must win regardless of bias. There is a well-known lower bound $\Omega\big(k^2\big)$ on the number of states stable protocols by Natale and Ramezani [72]. We have to deal with this negative obstacle to explore the plurality variant with $k > 2$ opinions. To overcome this issue, we relax the statement by allowing a negligible failure probability to significantly beat the state space lower bound. We provide protocols that solve the exact plurality problem w.h.p. The first protocol is tournament-based, assuming an ordering on the opinions and solving the problem in $\mathrm{O}\big(kn \log n\big)$ interactions using $\mathrm{O}\big(k + \log n\big)$ states. In the second protocol, we get rid of the ordering assumption at the cost of a slightly increased number of interactions $\mathrm{O}\big(kn \log n + n \log^2 n\big)$. At last, we reduce the number of tournaments by pruning the number of participating opinions at the cost of a slightly increased state space $\mathrm{O}\big(k \log\log n + \log n\big)$. Essentially, we filter out smaller opinions during the preprocessing. This leads to a trade-off between time and space: the number of interactions is $\mathrm{O}\big(n^2/x_{\max}\log n + n\log^2 n\big)$ where $x_{\max}$ denote the support of the initial largest opinion. On the other hand, it improves the runtime of a broad range of configurations (e.g., large initial bias in a skewed configuration with many small opinions).

In Chapter 5, we solve a variant of the approximate majority consensus problem with a preferred opinion 1 and unpreferred opinion 2. We characterize agents with the preferred opinion as stubborn. The stubbornness is described by a parameter p and used in the so-called stubborn USD. This variant works the same way as the known USD with one exception (see Fig. 1.4): Whenever an agent with the preferred opinion interacts with an agent with the opposite opinion, it keeps its opinion with probability p. Otherwise, it becomes undecided.

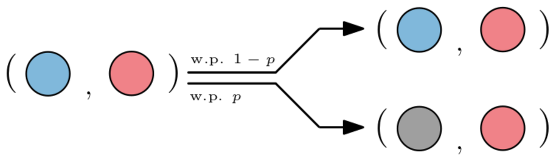

Figure 1.4.: Stubborn-USD: preferred opinion (blue), unpreferred opinion (red), undecided (grey)

Due to this advantage of the preferred opinion, we are especially interested in initial configurations where the unpreferred opinion is the initial majority.

We show that for any configuration with linear support for both opinions, there exists a threshold p_s such that Opinion 1 wins after $\mathrm{O}\big(n\log n\big)$ interactions if the stubbornness parameter p is slightly larger than p_s. If p is slightly smaller than p_s, Opinion 2 wins after $\mathrm{O}\big(n\log n\big)$ interactions. If $p \approx p_s$, we show that one of the two opinions wins in $\mathrm{O}\big(n\log^2 n\big)$ interactions, but which of the two opinions wins is unclear. Our results show that even if the initial support for Opinion 2 is larger, for sufficiently large p, the agents will still agree on Opinion 1.

1.5. Related Work

In this section, we provide an overview of the related research.

Undecided State Dynamics and its variants The two-opinion USD was introduced independently by Angluin et al. [9] for the population protocol model and by Perron et al. [75] for the closely related (continuous time) asynchronous gossip model. Both works show that the process converges w.h.p. in $\mathrm{O}\big(n\log n\big)$ steps (respectively, $\mathrm{O}\big(\log n\big)$ continuous time). Condon et al. [35] give an improved analysis for the two-opinion case in the population model and show the process solves the approximate majority of the problem assuming an initial additive bias of $\Omega\big(\sqrt{n\log n}\big)$, which improves over the additive bias of $\omega\big(\sqrt{n}\log n\big)$ needed in the analysis of Angluin et al. It is worth mentioning that they also consider a variant of plurality consensus with k opinions in a communication model in which three randomly chosen agents interact in a step. In this variant, if an agent interacts with two other agents of the same opinion, it adopts this opinion. They show that the system converges to the initial majority within $\mathrm{O}(kn \log n)$ interactions w.h.p., provided the initial bias is large enough.

Clementi et al. [34] study the USD in the gossip model. Note that the gossip model is a synchronous time model, where in each (parallel) round, every agent selects uniformly at random interaction partners. They consider the unbiased case with $k = 2$ opinions and show that, w.h.p., the protocol reaches consensus in $\mathrm{O}(\log n)$ rounds. Moreover, they show that the plurality opinion prevails if the initial bias is $\Omega(\sqrt{n \log n})$. Their analysis partitions the configuration space into seven cases, depending on the magnitude of a possible bias and the number of undecided agents, making it hard to apply the approach to arbitrary values of k. Another recent work by D'Amore et al. [42] also considers the USD with $k = 2$. While they require the usual initial bias of $\Omega(\sqrt{n \log n})$, they introduce *noise* in their model, which may modify sent messages with a certain probability p. Within $\mathrm{O}(\log n)$ time, their protocol reaches a $\Theta(n)$ bias towards the initial majority. Becchetti et al. [20] adopt the USD to the gossip model and generalize it to $k = \mathrm{O}((n/\log n)^{1/3})$ opinions. Central to their analysis is the introduction of the *monochromatic distance*, which measures the uniformity (i.e., lack of bias) of a configuration. Roughly speaking, this distance is the sum of squares of the support of each opinion, normalized by the square of the most popular opinion. They show convergence within $\mathrm{O}(\mathrm{md}(\mathbf{x}) \cdot \log n)$ parallel rounds, where $\mathrm{md}(\mathbf{x})$ is the monochromatic distance of the initial configuration, which is always bounded above by k. This analysis only holds when the initial configuration has a multiplicative bias.

Deviating slightly from the original USD definition, Berenbrink et al. [29] and Ghaffari and Parter [53] consider a *synchronized* version of the USD. For the synchronization, both suggest basically the same protocol, which uses counters to partition time into phases of length $\Theta(\log k)$. Agents can become undecided only at the start of such a phase and use the rest of the phase to obtain a new opinion. Both protocols achieve consensus in $\mathrm{O}(\log k \log n)$ rounds w.h.p., using $\log k + \mathrm{O}(\log \log k)$ bits per agent. The runtime can be slightly reduced to $\mathrm{O}(\log k \cdot \log \log_\alpha n + \log \log n)$, where α denotes the initial multiplicative bias [29]. Both [29, 53] proceed to use (different) more sophisticated synchronization mechanisms to design protocols that require only $\log k + \mathrm{O}(1)$ bits and maintain (essentially) the same runtime bounds (the refined runtime of [29] becomes $\mathrm{O}(\log(n) \cdot \log \log_\alpha n)$). Note that neither [53] nor [29] extend to the case without bias or very large k: their techniques are based on chains of concentration bounds, which are no longer applicable in the general setting.

Other Consensus Dynamics A large number of works [1, 3, 22, 27, 28, 47] aim to identify the majority opinion with $k = 2$ even if the initial winning margin is as small as 1. The best-known result [47] solves this *exact* majority problem in $\mathrm{O}(n \log n)$ interactions using $\mathrm{O}(\log n)$ states, both in expectation.

The *majority problem* is a special case of plurality consensus.As a fundamental problem in distributed computing, a lot of work has been invested to find an (asymptotically) optimal, stable population protocol for exact majority [1, 3, 22, 27, 28], culminating in [47], which solves majority using both $\mathrm{O}(\log n)$ states and expected parallel time. This is optimal in that no stable protocol can solve majority faster ($\Omega(\log n)$ is the time until each agent interacted at least once), and any polylogarithmic-time stable majority protocol requires $\Omega(\log n)$ states (under two natural conditions, see [1]). Note that the difficulty here stems from requiring *exactness.* Focusing on constant-state protocols that might fail, [61] mentions a protocol with constant state space and which w.h.p. determines the exact majority in time $\mathrm{O}(\log^3 n)$.

Population protocols for general plurality consensus are scarce. One line of research studies the state complexity (ignoring time) required to *always* identify the plurality opinion. While one needs at least k states to represent k opinions, Natale and Ramezani [72] show that always correct plurality consensus requires even $\Omega(k^2)$ states. The currently best always correct protocol needs $\mathrm{O}(k^{11})$ states, which can be reduced to $\mathrm{O}(k^6)$ if there is a total ordering on the opinions [50]. The quadratic lower bound makes it apparent that *always* guaranteeing a correctly identified plurality opinion comes at the cost of high space complexity.

A related family of protocols is the j-Majority processes. The idea is that every agent adopts the majority opinion among a random sample of j other agents (breaking ties randomly). The most simple variant (for $j{=}1$) is also known as the so-called Voter process [31, 36, 55, 59, 71]. Here, every agent adopts the opinion of a single, randomly chosen agent. The protocols for $j = 2$ and $j = 3$ have been analyzed under the names of TwoChoices process [37, 38, 39] and the 3-Majority dynamics [21, 25, 52]. In the TwoChoices process, lazy tie-breaking towards an agent's original opinion is assumed. Ghaffari and Lengler [52] show for the TwoChoices process with $k = \mathrm{O}\big(\sqrt{n/\log n}\big)$ and for 3-Majority with $k = \mathrm{O}\big(n^{1/3}/\log n\big)$ that consensus is reached in $\mathrm{O}(k \cdot \log n)$ rounds w.h.p. For arbitrary k they show that 3-Majority reaches consensus in $\mathrm{O}\big(n^{2/3} \log^{3/2} n\big)$ rounds w.h.p. Schoenebeck and Yu [76] analyze the convergence time of a generalization of multi-sample consensus protocols for two opinions on complete graphs and Erdős-Rényi graphs. In the MedianRule process [46], the authors assume that opinions are ordered. In every step, every agent adopts the median of its own opinion and two randomly sampled opinions. This protocol reaches consensus in $O(\log k \log\log n + \log n)$ rounds w.h.p. In contrast to the Median Rule, we remark that the USD does not require a total order among the opinions.

The authors in [43] study the 2-USD with uniform noise in a synchronous time model (gossip). Whenever an agent communicates with another agent, it observes its actual state only with probability $1-p$. Otherwise, it observes any state uniformly at random. The main result is a phase transition regarding the probability p. When the probability p is less than $1/6$, the configuration quickly reaches a meta-stable almost consensus. On the other hand, when the probability is greater than $1/6$, the initial majority is lost in $O(\log n)$ rounds.

D'Amore and Ziccardi [44] consider uniform communication noise for the 3-Majority dynamics. They observe a similar phase transition as in [43].

In [67], Mobilia examined the role of a single so-called *zealot* – an agent that never changes its opinion – for the Voter dynamics. Mobilia et al. [68] pursued this further for several zealots. Yildiz et al. [77] examined the role of two sets of zealots with opposing opinions – which they named *stubborn agents*.

In [18], Becchetti et al. consider a constant number of opinions and a single stubborn agent. In each time step, an agent is activated uniformly at random and samples ℓ opinions of other agents uniformly at random. They show for a constant number of initial opinions that every memoryless dynamics requires $\Omega(n^2)$ time steps in expectation to converge.

Alistarh et al. [2] introduce the *catalytic input model* (CI model) as a special variant of the population protocols model. In this model, the agents are either catalytic or non-catalytic agents. Catalytic agents never change their state. On the other hand, non-catalytic agents are allowed to change their state. Additionally, non-catalytic agents can perform spontaneous state changes; the so-called leak rate specifies the frequency of the spontaneous interactions. The goal of the non-catalytic agents is to compute a function over the states of the catalytic agents. The authors develop a protocol to detect whether there is a catalytic agent in a given state D. Note that non-catalysts can compute false positives due to the leaky interactions.

Amir et al. [6] considerthe catalytic input model with n catalytic agents and m non-catalytic agents which they call worker agents ($N = n+m$). They solve the approximate majority problem for two opinions w.h.p. in $\mathrm{O}\big(N \log N\big)$ interactions when the initial bias among the catalytic agents is $\Omega\big(\sqrt{N \log N}\big)$ and $m = \Theta\big(n\big)$. They show that the size of the initial bias is tight up to a $\mathrm{O}\big(\sqrt{\log N}\big)$ factor. Additionally, they consider the approximate majority problem in the CI and population models with leaks. Their protocols tolerate a leak rate of at most $\beta = \mathrm{O}\big(\sqrt{N \log N}/N\big)$ in the CI model and a leak rate of at most $\beta = \mathrm{O}\big(\sqrt{n \log n}/n\big)$ in the population model. They also show a separation between the CI and population models' computational power.

Alistarh et al. [4] consider the CI model and introduce the *robust comparison problem*. The catalytic agents are either in state A or B. The goal of the worker agents is to decide the majority state, i.e., whether A or B has the larger support. In the dynamic version, the number of agents in state A or B can change during the execution as long as the counts for A and B remain stable sufficiently long, allowing the algorithm to stabilize on an output. If at time t at least $\Omega\big(\log n\big)$ catalytic agents are in either A or B and the ratio between the numbers of agents supporting agents A and B is at least a constant, then most non-catalytic agents (up to $\mathrm{O}\big(n \log n\big)$ agents) outputs w.h.p. the correct majority. The protocol needs $\mathrm{O}\big(\log n \cdot \log\log n\big)$ states per agent, assuming that the number of catalytic agents in A and B does not change in the meantime. Additionally, the authors show that their protocol is robust to leaky transitions at a rate of $\mathrm{O}\big(1/n\big)$. If the initial support of A and B states is $\Omega\big(\log^2 n\big)$ the authors can strengthen their results such that a ratio between the two base states of $1 + \mathrm{o}\big(1\big)$ is sufficient.

Dudek and Kosowski [48] develop population protocols for broadcasting and source detection (the agents must decide if at least one agent in a dedicated source state is present in the population). Both protocols are based on oscillatory dynamics.

Biased Opinions All results for biased opinion dynamics consider the case with two opinions. One of the opinions is the preferred one. Anagnostopoulos et al. [7] consider biased opinion dynamics consisting of two steps. In each time step, an agent is selected uniformly at random and adopts the preferred opinion with probability α. Otherwise, the agent adopts the majority of its neighbors' opinions. Note that an agent can adopt the preferred opinion in this setting even if none of its neighbors share it. Hence, in contrast to our model, the system has only one absorbing state where all agents agree on the preferred opinion. The authors show a phase transition for $\alpha = 1/2$ in dense graphs for the majority rule. For the process with voter rule (the node adopts a random neighbor's opinion), they do not observe any phase transition; the absorption time is $O(n \log n)$.

Mukhopadhyay et al. [70] consider the same biased opinion variant as [7] but with the 2-choices rule. They bound the expected absorption time for the complete graph and observed a phase transition around $\alpha = 1/9$. For large $\alpha > 1/9$, the process converges to the preferred opinion in time $O(n \log n)$. On the other hand, for small $\alpha < 1/9$, the convergence time depends on the initial fraction of the preferred opinion among the population. Suppose the initial support of the preferred opinion is sufficiently large. In that case, the process converges fast ($O(n \log n)$), and for a small initial support the process needs at least $\Omega(\exp(n))$ steps.

Cruciani et al. [41] consider a variant of the biased opinion dynamics on core-periphery networks. The network consists of a core that is a densely-connected subset of agents with the same opinion. The remaining agents form the periphery with another opinion. They observe a phase transition that depends on the cut between the core and periphery. Either all agents agree relatively fast on the initial opinion of the core agents. Otherwise, the process remains in a meta-stable where both opinions remain in the network for at least polynomial many rounds.

Cruciani et al. [40] study the j-majority dynamics. In each time step, each agent simultaneously samples j neighbors' uniform at random, and it adopts the majority opinion. They consider two different noise models that alter the communication between agents with probability p. In the first variant, an agent may observe the preferred opinion instead of the sampled opinion. In the second variant, an agent's opinion may change directly to the preferred one. Both variants show phase transitions for some p^*, with the preferred opinion being the initial minority. Eventually, all agents agree on the preferred opinion w.h.p. For small $p < p^*$ and $j \geq 3$, it requires $n^{\omega(1)}$ parallel rounds. Large $p > p^*$ and $j \geq 3$ only require $O(1)$ parallel rounds. At last, for $j < 3$, it requires $O(1)$ parallel rounds for every $p > 0$.

1.6. Own Publications

Parts of the thesis are based on the following publications [5, 15, 16, 23]:

[5] T. Amir, J. Aspnes, P. Berenbrink, F. Biermeier, C. Hahn, D. Kaaser, and J. Lazarsfeld. "Fast Convergence of k-Opinion Undecided State Dynamics in the Population Protocol Model". In: *Proceedings of the 2023 ACM Symposium on Principles of*

Distributed Computing, PODC 2023, Orlando, FL, USA, June 19-23, 2023. Ed. by R. Oshman, A. Nolin, M. M. Halldórsson, and A. Balliu. ACM, 2023, pp. 13–23. DOI: 10.1145/3583668.3594589.

[15] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling. "Fast Consensus via the Unconstrained Undecided State Dynamics". In: *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. SIAM, 2022, pp. 3417–3429. DOI: 10.1137/1.9781611977073.135.

[16] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling. "Population Protocols for Exact Plurality Consensus: How a small chance of failure helps to eliminate insignificant opinions". In: *PODC '22: ACM Symposium on Principles of Distributed Computing*. ACM, 2022, pp. 224–234. DOI: 10.1145/3519270.3538447.

[23] P. Berenbrink, F. Biermeier, and C. Hahn. *Undecided State Dynamics with Stubborn Agents*. 2024. arXiv: 2406.07335 [cs.DC].

Further publications [12, 24] that are not part of the thesis :

[12] A. Antoniadis, F. Biermeier, A. Cristi, C. Damerius, R. Hoeksma, D. Kaaser, P. Kling, and L. Nölke. "On the Complexity of Anchored Rectangle Packing". In: *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*. Ed. by M. A. Bender, O. Svensson, and G. Herman. Vol. 144. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 8:1–8:14. DOI: 10.4230/LIPICS.ESA.2019.8.

[24] P. Berenbrink, F. Biermeier, C. Hahn, and D. Kaaser. "Loosely-Stabilizing Phase Clocks and The Adaptive Majority Problem". In: *1st Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2022, March 28-30, 2022, Virtual Conference*. Ed. by J. Aspnes and O. Michail. Vol. 221. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 7:1–7:17. DOI: 10.4230/LIPICS.SAND.2022.7.

Almost all works are accepted publications of peer-reviewed conferences during my time as a PhD student.

Declaration on the Own Contribution The papers are a joint work of all authors. The thesis mostly contains parts of the papers in which I was primarily involved. There are certain exceptions, like necessary technical results to remain self-contained.

In the publication [15], the concentration bounds of a single phase are a necessity with no major contribution by myself. My focus was the unbiased case with a small number of opinions k and the proof draft of the biased case.

In the publication [16], we developed the protocols jointly. The proof of the initialization phase in the protocols is necessary, with no major contribution by myself. I mainly focused on the tournament phase's proofs and removing the ordering assumption. Furthermore, I was involved in drafting the last protocol's proof of the phase clock subpopulation approach.

In the publication [5], the bounds on the undecided agents are a necessity with no major contribution by myself. I focused on creating a sufficiently large additive and multiplicative bias. In contrast to the original paper, I extended the former result by removing the constraint on the initial number of opinions k. I replaced it with a mild constraint on the initial largest opinion. Furthermore, I also relaxed the constraint on the initial number of undecided. In this course, I especially revised the first phase, which mainly deals with the evolution of the undecided agents.

In the unpublished work [23], I proposed the protocol and focused on the proof of the biased case. Furthermore, I was involved in the discussion and draft of creating an initial bias.

1.7. Further Notation

A configuration $\mathbf{x}(t)$ at time t is a vector $(x_1(t), x_2(t), \ldots x_k(t), u(t))$ of length $k+1$. For $1 \leq i \leq k$, $x_i(t)$ is the number of agents of Opinion i and $u(t) = n - \sum_{i=1}^{k} x_i(t)$ is the number of undecided agents. In the beginning we assume $x_1(0) \geq x_2(0) \geq \cdots \geq x_k(0)$. For $t > 0$, we define $\max(t)$ as the index of the opinion with the largest support at step t (if there are several opinions with the same maximum support, we pick an arbitrary one). Furthermore we introduce the notation $x_{\max}(t) = x_{\max(t)}(t) = \max_{i \in [k]} \{ x_i(t) \}$ for the support of the largest opinion at time t. Note that $x_{\max}(t)$ can refer to the support of different opinions over time.

A configuration $\mathbf{x}$ has an *additive bias* β if there exists an Opinion m such that for all other opinions $i \neq m$, we have $x_m \geq x_i + \beta$. We say that a configuration $\mathbf{x}$ has a *multiplicative bias* α if there exists an Opinion m such that for all other opinions $i \neq m$ we have $x_m \geq \alpha \cdot x_i$. In the following, we use upper case letters for random variables (for example, $\mathbf{X}(t)$ and $U(t)$) and lower case letters ($\mathbf{x}(t)$ and $u(t)$) for fixed configurations or values.

We use $\mathcal{F}_t = (F_i)_{i=0}^{t}$ to denote the natural filtration consisting of the initial configuration at time 0 and all random choices up to time t: the interacting agents, and the outcome of the interaction based on the stubbornness p whenever it is part of the model (see Chapter 5). I.e., We write $\mathcal{F}_t$ for $\mathbf{X}(0) = \mathbf{x}(0), \mathbf{X}(1) = \mathbf{x}(1), \ldots, \mathbf{X}(t) = \mathbf{x}(t)$ and for the sake of readability we may use $\mathbf{x}$ instead of $\mathbf{x}(t)$.

Chapter 2.

Approximate Plurality Consensus via Synchronized Undecided State Dynamics

In this part, we consider approximate plurality consensus in the population protocol model. Initially, all n agents have one of the k opinions. Recall that agents interact in pairs (randomly chosen) and update their state according to the underlying protocol. Eventually, all agents must agree on the same opinion. Moreover, if there is a sufficiently large *bias* – i.e., the difference between the number of agents initially assigned to the largest and second largest opinion – the initial largest opinion should prevail.

Berenbrink et al. [29] and Ghaffari and Parter [53] introduced a *synchronized* version of the Undecided State Dynamics (USD) in the gossip model. Recall that the gossip model is a synchronous time model, where in each (parallel) round, every agent selects uniformly at random interaction partners. In the synchronized USD, a sequence of (parallel) rounds forms a phase of two parts. In the first part, an agent becomes undecided if its interaction partner has a different opinion. In the second part, an undecided agent adopts the opinion of its interaction partner. Note that the first part only lasts a single round due to the gossip model. Essentially, if an agent loses its opinion in the first part, it tries to recover an opinion during the remainder of a phase.

We consider the synchronized USD protocol in the population protocol model. Here, a synchronization mechanism known as *phase clocks* is used to make the agents jointly progress through *phases* of length $\Theta(n \log n)$, alternating between *decision parts* (where agents become undecided if they encounter a different opinion) and *boosting parts* (where undecided agents adopt one of the remaining opinions).

Results and Methodology Our protocol reaches, w.h.p., plurality consensus in $O(n \log^2 n)$ interactions under a mild assumption on the initial number of opinions k (see Theorem 2.1). To be more precise, all agents agree on what we call a *significant* opinion (an opinion whose support is at most $O(\sqrt{n \log n})$ smaller than the largest opinion).

Theorem 2.1. *Consider Algorithm 2.1 on an initial configuration* $\mathbf{x}$ *with* $k = O(\sqrt{n}/\log n)$ *opinions. The algorithm uses* $k \cdot \Theta(\log n)$ *states per agent and has the following properties: All agents agree on a significant opinion in* $O(n \cdot \log^2 n)$ *interactions, w.h.p.*

If the initial configuration has an additive bias of order $\Omega(\sqrt{n \log n})$ (see Theorem 2.2), all agents agree on the initial largest opinion. A similar result holds for a multiplicative bias and essentially matches the known results from [29, 53] in the gossip model.

Theorem 2.2. *Consider Algorithm 2.1 on an initial configuration* $\mathbf{x}$ *with* $k \leq n$ *opinions. The algorithm uses* $k \cdot \Theta(\log n)$ *states per agent and has the following properties:*

1. *Assume* x *has additive bias of* $\xi \cdot \sqrt{n \cdot \log n}$ *and multiplicative bias of* α*. W.h.p. the algorithm reaches a configuration in which all agents agree on the initial plurality opinion in time*
   - $\mathrm{O}(n \cdot \log n \cdot \log\log_\alpha n)$ *if* $k \leq \sqrt{n}/\log n$ *and*
   - $\mathrm{O}(n \cdot \log n \cdot (\log\log_\alpha n + \log\log n))$ *if* $k > \sqrt{n}/\log n$.

Considering the parameter α, an initial additive bias of order $\Omega(\sqrt{n \log n})$ implies $\alpha = \Omega(\sqrt{n/\log n})$ and hence, $\log\log_\alpha n = \mathrm{O}(\log n)$. On the other hand, a constant multiplicative bias α improves the runtime in Theorem 2.1 by replacing a $\log n$ term with a $\mathrm{O}(\log\log n)$ term.

The main challenge is to handle the case without a clear bias, i.e., there is almost no difference between the number of agents initially assigned to the most common and second most common opinion. The analysis of our protocol relies on specialized tail bounds based on the Pólya-Eggenberger distribution.

For the unbiased case with $k = \mathrm{O}(\sqrt{n}/\log n)$, we adopt an idea suggested (but not followed completely through) by Ghaffari and Lengler [52] for the TwoChoices and 3Majority process: Depending on their support, opinions are categorized as either *strong*, *weak*, or *super-weak*. Roughly speaking, the support of strong opinions is at least a constant fraction of the largest opinion. We first show that w.h.p., weak and super-weak opinions never become strong. Then, we consider an arbitrary pair of opinions and show that, w.h.p., one of them becomes super-weak after $\mathrm{O}(\log n)$ phases. Thus, a union bounds yields that, w.h.p., a single strong opinion (note that the opinion of maximum support is always strong) prevails and eventually becomes the consensus opinion. A crucial part is dealing with two strong opinions of roughly the same support. We combine anti-concentration and concentration bounds to create a sufficient gap between both opinions. This allows us to apply a known drift result such that their support drifts further apart. Although we cannot guarantee that the initial largest opinion will win, as mentioned before, the contenders must be significant.

Note that we cannot necessarily apply this categorization of the opinions to larger values of k, e.g., $k = \Omega(\sqrt{n})$. In this case, the variance becomes too high. It makes it likely that weak or even super-weak opinions become strong again, violating the key invariant used to analyze that case. In fact, there are configurations in which even the initial largest opinion is, by definition, both strong and super-weak.

In contrast to the unbiased case, assuming an initial (large enough) bias, we can track the ratio between the largest and second-largest opinions over the phases. This allows us to cover a large range of values for k.

2.1. Synchronized Undecided State Dynamics

This section formally presents the synchronized variant of the undecided state dynamics.

Algorithm 2.1: Synchronized USD

```
Actions performed when agents (u, v) interact:
 ▷ Decision Part: clock[u] < 2τ log n
if clock[u] < 2τ log n and not decision[u] then
if opinion[u] ≠ opinion[v] then
  undecided[u] ← True
else
    undecided[u] ← False
   decision[u] ← True

 ▷ Boosting Part: clock[u] ≥ 2τ log n
 if clock[u] ≥ 2τ log n and undecided[u] then
  if not undecided[v] then
    undecided[u] ← False
    opinion[u] ← opinion[v]
  decision[u] ← False

 ▷ Leaderless Phase Clock [1]
 if clock[u] ≤_(6τ log n) clock[v] then
     clock[u] ← (clock[u] + 1)  mod 6τ log n
 else
     clock[v] ← (clock[v] + 1)  mod 6τ log n
```

The *state* of an agent u is a tuple $(\mathsf{clock}[u], \mathsf{opinion}[u], \mathsf{decision}[u], \mathsf{undecided}[u])$ (see Algorithm 2.1). $\mathsf{opinion}[u] \in \{1, 2, \ldots, k\}$ stores the current opinion of agent u. The Boolean variable $\mathsf{undecided}[u]$ indicates whether agent u is currently undecided, and $\mathsf{decision}[u]$ indicates whether agent u has already performed an interaction in the decision part. Both flags are initialized to False.

Our protocol uses the leaderless phase clock from [1] (which runs on every agent) to divide the interactions into *phases*, each consisting of $\mathrm{O}(n \log n)$ interactions. The first part of a phase is called *decision part*, and the second part is called *boosting part*. In the decision part, every agent becomes undecided if and only if its first interaction partner has a different opinion. It sets its $\mathsf{undecided}$ bit in that case. In the boosting part, every undecided agent adopts the opinion of a randomly sampled agent that is not undecided. This agent propagates its opinion to other undecided agents in subsequent steps. The clock of agent u uses the variable $\mathsf{clock}[u]$ (initially 0) which can take values in $\{0, \ldots, 6\tau \log n - 1\}$ for a suitably chosen constant τ. In each interaction, the smaller[1] of the two values $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ is increased by one modulo $6\tau \log n$. For a polynomial number of interactions it guarantees [1, see Section 4] that for any pair of agents u and v the distance[2] between $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ is at most $\tau \log n$, and every agent participates

[1]Smaller w.r.t. the circular order modulo $m = 6\tau \log n$, defined as $a \leq_{(m)} b \equiv (a \leq b \text{ XOR } |a-b| > m/2)$.
[2]Distance w.r.t. the circular order modulo $m = 6\tau \log n$, defined as $|a - b|_{(m)} =$

in $\Omega(\log n)$ interactions in every phase. Hence, it cleanly separates the decision and boosting parts.

2.2. Analysis for the Population Protocol Model

This section introduces preliminary notions and observations to handle the synchronized USD.

We consider a system of n identical, anonymous agents. Initially, each agent has one of k possible opinions, which we represent as numbers from the set $\{1, 2, \ldots, k\}$. We do not assume an order among the opinions. A *configuration* describes the current state of the system and can be represented as an (unsorted) vector $\mathbf{x} = (x_i)_{i=1}^{k} \in \{0, 1, \ldots, n\}^k$, where x_i is the *support* of opinion i, defined as the number of agents with opinion i.

The (additive) *bias* of a configuration $\mathbf{x}$ is $x_{(1)} - x_{(2)}$, where $x_{(i)}$ denotes the *support* of the i-th largest opinion (ties broken arbitrarily but consistently). The *multiplicative bias* is defined as $x_{(1)}/x_{(2)}$. In the analysis, we also use $x_{max} = x_{(1)}$ to denote the support of the largest opinion. For any opinion i with $x_i = x_{max}$ we say opinion i *provides* x_{max}. We use $\mathcal{S}(\mathbf{x})$ to denote the set of significant opinions in configuration $\mathbf{x}$.

In our analysis, we assume that the phase clocks properly separate the boosting and decision parts of the considered phases. This follows from [1, 74], where it is shown that for a polynomial number of phases and any pair of agents u and v, the distance between $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ w.r.t. the circular order modulo $6\tau \log n$ is less than $\tau \log n$, w.h.p. The choice of τ also ensures that every undecided agent can adopt an opinion in the boosting part of a phase, w.h.p.

The strict phase synchronization allows us to define a series of random vectors $\mathcal{X} = (\mathbf{X}(t))_{t\in\mathbb{N}}$ that describe the configurations at the beginning of phase t where the i-th entry $X_i(t)$ is the number of agents with opinion i. For the analysis, we also define a series of random vectors $\mathcal{Y} = (\mathbf{Y}(t))_{t\in\mathbb{N}}$ where $Y_i(t)$ is the number of decided agents with opinion i at the beginning of the *boosting* part of phase t. Finally, the series $\mathcal{X}_{max} = (X_{max}(t))_{t\in\mathbb{N}}$ describes the size of the support of the largest opinion. We generally use bold font to denote vectors, non-bold font to denote vector components, and capital letters for random variables. When we fix the value of a random variable at the beginning of a phase t, we use lowercase letters, e.g., $\mathbf{X}(t) = \mathbf{x}(t)$. When it is clear from the context, we omit the parameter t in the proofs.

The following observation shows that a binomial distribution can describe the opinion distribution after the decision part. Note that $\|\mathbf{Y}(t)\|_1$ denotes the number of decided agents at the beginning of the t-th boosting part.

Observation 2.3 (Decision Part)**.** *Assume $\mathbf{X}(t) = \mathbf{x}(t)$ is fixed and let $\mathbf{Y}(t)$ be the configuration at the beginning of the boosting part of phase t. Then, for $1 \le i \le k$, the $Y_i(t)$ have an independent binomial distribution with $Y_i(t) \sim \mathrm{Bin}(x_i(t), x_i(t)/n)$. Additionally, for $\psi(t) := \mathbb{E}[\|\mathbf{Y}(t)\|_1] = \sum_{i=1}^{k} x_i(t)^2/n$ we have $n/k \le \psi(t) \le x_{max}(t)$.*

$\min\{|a-b|, m-|a-b|\}$.

The opinion distribution after the boosting part can be modeled by a so-called *Pólya-Eggenberger distribution.* The Pólya-Eggenberger process is a simple urn process that runs in multiple steps. Initially, the urn contains a red and b blue balls, where $a, b \in \mathbb{N}_0$. In each process step, one ball is drawn uniformly at random from the urn, its color is observed, and it is returned with one additional ball of the same color. The corresponding *Pólya-Eggenberger distribution* $\mathrm{PE}(a, b, m)$ describes the number of *total* red balls contained in the urn after m steps. To bound $X_i(t+1)$, we use tail inequalities (Theorem 1 and Theorem 47) shown in the full version of [17]. We state these bounds in Appendix A.3 for convenience.

Observation 2.4 (Boosting Part). *Assume $\mathbf{Y}(t) = \mathbf{y}(t)$ is fixed and $\|\mathbf{y}(t)\|_1 \geq 1$. Let $\mathbf{X}(t+1)$ be the configuration at the beginning of the decision part of phase $t+1$. Then, for $1 \leq i \leq k$, $X_i(t+1)$ has Pólya-Eggenberger distribution $X_i(t+1) \sim \mathrm{PE}(y_i(t), \|\mathbf{y}(t)\|_1 - y_i(t), n - \|\mathbf{y}(t)\|_1)$ with $\mathbb{E}[X_i(t+1)] = y_i(t) \cdot (n/\|\mathbf{y}(t)\|_1)$.*

Proof. This follows from an easy coupling of the boosting part with the Pólya-Eggenberger process, defined as follows. Let $\ell_0 \geq 1$ be the number of decided agents at the beginning of the boosting part. For $\ell_0 < i \leq n$, the process picks an arbitrary one of the undecided agents. This agent chooses one of the ℓ_{i-1} decided agents uniformly at random and adopts its opinion, resulting in $\ell_i := \ell_{i-1} + 1$. Our process's coupling with this process is now straightforward, as we discard all interactions that do not change the number of decided agents. $\square$

Finally, we introduce some important constants that we use throughout our analysis.

Definition 2.5. *We define $\xi := (160 \cdot c_w)^2 + (148 \cdot c_p)^2$, $c_w := 8\sqrt{1 + 2/\varepsilon^*}$, $\varepsilon^* := \varepsilon_p/192$ and $c_k := 4 \cdot (2625 + c_p)^2$. The constants $1 > \varepsilon_p > 0$ and $c_p > 1$ originate from the Pólya-Eggenberger concentration results of [17].*

Definition 2.6. *Opinion i in configuration $\mathbf{x}$ is called* super-weak *iff $x_i \leq c_w \cdot \sqrt{n \log n}$,* weak *iff $c_w \cdot \sqrt{n \log n} < x_i < 0.9 \cdot x_{max}$, and* strong *iff $x_i \geq 0.9 \cdot x_{max}$. Additionally, we call an opinion i* significant *if $x_{(i)} \geq x_{(1)} - \xi \cdot \sqrt{n \log n}$ (the constant ξ is specified in Definition 2.5 and originates from our analysis in Section 2.2). An opinion that is not significant is called* insignificant.

2.3. Analysis of a Single Phase

In this section, we analyze the evolution of opinions throughout some fixed phase t. The following lemma gives Chernoff-like guarantees for a large range of deviations and opinion sizes.

Lemma 2.7. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$ and an opinion i with support $x_i(t)$. Furthermore, let $\psi = \sum_{j=1}^{k} x_j(t)^2/n$. Then, for any $0 < \delta < x_i(t)/\sqrt{n}$ and a suitable small constant*

$\varepsilon^* > 0$

$$\Pr\left[X_i(t+1) < \frac{x_i(t)^2}{\psi} - \frac{x_i(t)}{\psi}\sqrt{n}\delta\right] \leq 7e^{-\varepsilon^*\delta^2} \ ,$$
$$\Pr\left[X_i(t+1) > \frac{x_i(t)^2}{\psi} + \frac{x_i(t)}{\psi}\sqrt{n}\delta\right] \leq 7e^{-\varepsilon^*\delta^2}.$$

Proof. We focus on the upper bound and consider some fixed opinion i throughout phase t. At first we analyze the evolution of i throughout decision part t and consider $Y_i(t)$ and $\|\mathcal{Y}(t)\|$. Recall (see Observation 2.3) that $Y_i(t) \sim \mathrm{Bin}(x_i, x_i/n)$ and $\|\mathcal{Y}(t)\|_1 = \sum_{j=1}^{k} Y_j(t)$. We define the event $\mathcal{E}_i$ as follows

$$\mathcal{E}_i = \left\{Y_i(t) < \frac{x_i^2}{n} \cdot \left(1 + \frac{\delta\sqrt{n}}{8 \cdot x_i}\right) \text{ and } \|\mathcal{Y}(t)\|_1 > \psi \cdot \left(1 - \frac{\delta\sqrt{n}}{8 \cdot x_i}\right)\right\}.$$

First, we bound the probability of the event $\mathcal{E}_i$ complement. To do so, we apply Chernoff Bounds (Theorem A.1) to $Y_i(t)$ and $\|\mathcal{Y}(t)\|_1$. Hence, for $\delta' = \delta \cdot \sqrt{n}/(8 \cdot x_i) < 1$ and using $\psi \geq x_i^2/n$

$$\Pr\left[Y_i(t) \geq \frac{x_i^2}{n} \cdot (1 + \delta')\right] \leq \exp\left(-\frac{x_i^2 \cdot \delta'^2}{3 \cdot n}\right) \leq \exp\left(-\frac{x_i^2 \cdot n \cdot \delta^2}{192 \cdot x_i^2 \cdot n}\right) \leq \exp\left(-\frac{\delta^2}{192}\right)$$

$$\Pr\left[\|\mathcal{Y}(t)\|_1 \leq \psi \cdot (1 - \delta')\right] \leq \exp\left(-\frac{\delta'^2 \cdot \psi}{2}\right) \leq \exp\left(-\frac{x_i^2 \cdot n \cdot \delta^2}{128 \cdot x_i^2 \cdot n}\right) \leq \exp\left(-\frac{\delta^2}{192}\right)$$

An application of the union bound yields

$$\Pr\left[\bar{\mathcal{E}}_i\right] \leq 2\exp\left(-\frac{\delta^2}{192}\right). \tag{2.1}$$

Now, we deal with the outcome of the boosting part conditioned on the event $\mathcal{E}_i$. We fix $Y_i(t) = y_i$ and define $d := \|\mathcal{Y}(t)\|_1 = \sum_{j=1}^{k} y_j$. As mentioned in Observation 2.4 we model $X_i(t+1) \sim \mathrm{PE}(y_i, d-y_i, n-d)$. Applying the tail bound for the Pólya Eggenberger distribution from Theorem A.11 we get for $0 < \frac{\delta}{8} < \sqrt{y_i}$ and some constant $1 > \varepsilon_p > 0$ that

$$\Pr\left[X_i(t+1) > \frac{y_i}{d} \cdot n + \frac{\sqrt{y_i}}{d} \cdot n \cdot \frac{\delta}{8} \mid \mathcal{E}_i\right] < 4\exp\left(-\varepsilon_p \cdot \frac{\delta^2}{64}\right). \tag{2.2}$$

Since $y_i < \frac{x_i^2}{n} \cdot \left(1 + \frac{\delta\sqrt{n}}{8x_i}\right)$ and $d > \psi \cdot \left(1 - \frac{\delta\sqrt{n}}{8x_i}\right)$ we get

$$\begin{aligned}
\frac{y_i}{d} \cdot n + \frac{\sqrt{y_i}}{d} \cdot n \cdot \frac{\delta}{8} &< \frac{x_i^2}{\psi} \cdot \frac{\left(1 + \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)}{\left(1 - \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)} + \frac{\sqrt{\frac{x_i^2}{n} \cdot \left(1 + \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)}}{\psi \cdot \left(1 - \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)} \cdot n \cdot \frac{\delta}{8} \\
&= \frac{x_i^2}{\psi} \cdot \frac{\left(1 + \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)}{\left(1 - \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)} + \frac{x_i}{\psi} \cdot \frac{\sqrt{\left(1 + \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)}}{\left(1 - \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right)} \cdot \sqrt{n} \cdot \frac{\delta}{8} \\
&\overset{(*)}{<} \frac{x_i^2}{\psi} \cdot \left(1 + 3 \cdot \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right) + \frac{x_i \cdot \sqrt{n} \cdot \delta}{8\psi} \cdot \left(1 + 3 \cdot \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right) \\
&< \frac{x_i^2}{\psi} \cdot \left(1 + 3 \cdot \frac{\delta \cdot \sqrt{n}}{8 \cdot x_i}\right) + 2 \cdot \frac{x_i \cdot \sqrt{n} \cdot \delta}{8 \cdot \psi} \\
&= \frac{x_i^2}{\psi} + 5 \cdot \frac{x_i \cdot \sqrt{n} \cdot \delta}{8 \cdot \psi} < \frac{x_i^2}{\psi} + \frac{x_i \cdot \sqrt{n} \cdot \delta}{\psi}.
\end{aligned}$$

In (*) we apply the inequality $(1 + a)/(1 - a) \leq 1 + 3a$ which holds for all $a \leq 1/3$. Hence, a combination of this and Inequality (Eq. (2.2)) results in

$$\begin{aligned}
\Pr\left[X_i(t+1) > \frac{x_i^2}{\psi} + \frac{x_i}{\psi} \cdot \sqrt{n}\delta \mid \mathcal{E}_i\right] &< \Pr\left[X_i(t+1) > \frac{y_i}{d}n + \frac{\sqrt{y_i}}{d} \cdot n \cdot \frac{\delta}{8} \mid \mathcal{E}_i\right] \\
&< 4\exp\left(-\varepsilon_p \cdot \frac{\delta^2}{64}\right).
\end{aligned}$$

At last, we combine this with Inequality (Eq. (2.1)) by applying the law of total probability. Then we get that

$$\begin{aligned}
\Pr\left[X_i(t+1) > \frac{x_i^2}{\psi} + \frac{x_i^2}{\psi} \cdot \sqrt{n} \cdot \delta\right] &= \Pr\left[X_i(t+1) > \frac{x_i^2}{\psi} + \frac{x_i}{\psi} \cdot \sqrt{n} \cdot \delta \mid \mathcal{E}_i\right] \cdot \Pr\left[\mathcal{E}_i\right] \\
&\quad + \Pr\left[X_i(t+1) > \frac{x_i^2}{\psi} + \frac{x_i}{\psi} \cdot \sqrt{n} \cdot \delta \mid \bar{\mathcal{E}}_i\right] \cdot \Pr\left[\bar{\mathcal{E}}_i\right] \\
&< 4\exp\left(-\varepsilon_p \cdot \frac{\delta^2}{64}\right) + 2\exp\left(-\frac{\delta^2}{192}\right) < 7\exp\left(-\varepsilon^* \cdot \delta^2\right)
\end{aligned}$$

for some suitably chosen constant $\varepsilon^* = \varepsilon_p/192 > 0$. As a symmetric approach can develop a matching lower bound, we omit the detailed proof for this case. □

Lemma 2.7 does not give a high probability for opinions with small support. These cases are handled in the following lemma, providing a coarse bound for this regime.

Lemma 2.8. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$ and an opinion i with $x_i(t) \leq c\sqrt{n \log n}$. For any constant $c > 0$ and $\psi = \sum_{j=1}^{k} x_j(t)^2/n$, it holds that* $\Pr[X_i(t+1) > \left(12c^2 + 74c_p\right) \cdot n \cdot \log n/\psi] < 4n^{-2}$.

Proof. We track opinion i with support $x_i(t) \leq c\sqrt{n \log n}$ throughout the decision and boosting part t. Similar to the proof of Lemma 2.7 first we analyze $Y_i(t)$ and $\|\mathbf{Y}(t)\|$. Again we model $Y_i(t) \sim \mathrm{Bin}(x_i, x_i/n)$ and $\|\mathbf{Y}(t)\|_1 \sim \sum_{j=1}^{k} \mathrm{Bin}(x_j, x_j/n)$. Let $c' = \max\{c, \sqrt{6}\}$. Then, our goal is to bound $\Pr[X_i(t+1) > n/\psi \cdot (12c'^2 + 2c_p) \cdot \log n]$.

First, note that if $\psi < (12c'^2 + 2c_p) \cdot \log n$ then $n/\psi \cdot (12c'^2 + 2c_p) \cdot \log n > n$ and

$$\Pr\left[X_i(t+1) > \frac{n}{\psi} \cdot (12c'^2 + 2c_p) \cdot \log n\right] = 0$$

and the statement of the lemma follows immediately. Hence, in the following, we can assume that $\psi \geq (12c'^2 + 2c_p) \cdot \log n$. We define the event $\mathcal{E}_i$ as follows

$$\mathcal{E}_i = \left\{Y_i(t) < 2c'^2 \cdot \log n \wedge \|\mathbf{Y}(t)\|_1 > \frac{\psi}{2}\right\}.$$

First, we bound the probability of $\bar{\mathcal{E}}_i$. We apply general Chernoff upper Bound (Theorem A.2) and get for $\delta' = 1$ that

$$\Pr\left[Y_i(t) \geq 2c'^2 \cdot \log n\right] \leq \exp\left(-\frac{c'^2 \cdot \log n}{3}\right) \leq \exp\left(-\frac{6 \cdot \log n}{3}\right) \leq n^{-2}.$$

Due to the definition of c' we have $\psi \geq (12c'^2 + 2c_p) \log n \geq 10 \log n$. Applying Chernoff bounds with (Theorem A.1) $\delta' = \sqrt{(4 \log n)/\psi} \leq 1/2$ we get

$$\Pr\left[\|\mathbf{Y}(t)\|_1 \leq \frac{\psi}{2}\right] \leq \Pr\left[\|\mathbf{Y}(t)\|_1 \leq \psi \cdot \left(1 - \sqrt{\frac{4 \cdot \log n}{\psi}}\right)\right] \leq \exp\left(-\frac{4 \cdot \psi \cdot \log n}{2 \cdot \psi}\right) \leq n^{-2}.$$

An application of the union bound yields

$$\Pr\left[\bar{\mathcal{E}}_i\right] \leq 2n^{-2}. \tag{2.3}$$

The outcome of boosting part t remains to be considered. We fix $Y_i(t) = y_i$ and define $d := \|\mathbf{Y}(t)\|_1$.

Similar to Lemma 2.7, we model $X_i(t+1) \sim \mathrm{PE}(y_i, d - y_i, n - d)$ and apply Theorem A.12, which states a tail bound for this Pólya Eggenberger distribution, to deduce that

$$\Pr\left[X_i(t+1) > \frac{n}{d} \cdot (3y_i + c_p \cdot \log n) \;\middle|\; \mathcal{E}_i\right] < 2n^{-2}. \tag{2.4}$$

Conditioned on $\mathcal{E}_i$ we have $y_i < 2c'^2 \log n$ and $d > \frac{\psi}{2}$. Therefore,

$$\frac{n}{d} \cdot (3y_i + c_p \cdot \log n) < \frac{n}{\psi} \cdot (6y_i + 2c_p \cdot \log n) < \frac{n}{\psi} \cdot \left(\left(12c'^2 + 2c_p\right) \cdot \log n\right)$$

Together with Inequality (Eq. (2.4)), this yields

$$\begin{aligned}\Pr\left[X_i(t+1) > \frac{n}{\psi} \cdot ((12c'^2 + 2c_p) \cdot \log n) \;\middle|\; \mathcal{E}_i\right] &< \Pr\left[X_i(t+1) > \frac{n}{d} \cdot (3y_i + 2c_p \log n) \;\middle|\; \mathcal{E}_i\right] \\ &< 2n^{-2}.\end{aligned}$$

Similar to Lemma 2.7, the statement follows by applying the law of total probability. Note that $c_p > 1$, therefore $12c'^2 + 2c_p < 12c^2 + 74c_p$. $\square$

The following lemma shows that the bias between two opinions roughly squares if their size and the difference between them are $\Omega(\sqrt{n \log n})$.

Lemma 2.9. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$ and consider two opinions i, j with support $x_i(t) - x_j(t) \geq \xi\sqrt{n \log n}$ and $x_j(t) \geq c_w\sqrt{n \log n}$, then $\Pr[X_i(t+1)/X_j(t+1) \geq (x_i(t)/x_j(t))^{1.5}] \geq 1 - 2n^{-2}$.*

Proof. First we lower bound the support of opinion i and upper bound the support of opinion j. To do so, we apply Lemma 2.7 to both opinions with $\delta = \sqrt{(\ln 7 + 2\log n)/\varepsilon^*}$.

Note that $c_w \geq \frac{8\delta}{\sqrt{\log n}}$. In this way, we get

$$\Pr\left[X_i(t+1) \geq \frac{x_i^2}{\psi} - \frac{x_i}{\psi}\cdot\sqrt{n}\cdot\delta\right] \geq 1 - n^{-2} \text{ and } \Pr\left[X_j(t+1) \leq \frac{x_j^2}{\psi} + \frac{x_j}{\psi}\cdot\sqrt{n}\cdot\delta\right] \geq 1 - n^{-2}.$$

An application of the union bound yields with probability at least $1 - 2n^{-2}$

$$\frac{X_i(t+1)}{X_j(t+1)} \geq \left(\frac{\frac{x_i^2}{\psi} - \frac{x_i}{\psi}\cdot\sqrt{n}\cdot\delta}{\frac{x_j^2}{\psi} + \frac{x_j}{\psi}\cdot\sqrt{n}\cdot\delta}\right) \geq \left(\frac{x_i}{x_j}\right)^2 \cdot \left(\frac{1 - \frac{\sqrt{n}\delta}{x_i}}{1 + \frac{\sqrt{n}\delta}{x_j}}\right) \geq \left(\frac{x_i}{x_j}\right)^2 \cdot \left(1 - \frac{2\sqrt{n}\delta}{x_j}\right)$$

where we applied the inequality $(1-a)/(1+a) \geq 1 - 2a$ for any a. For the moment, let us assume $1 - (2\sqrt{n}\cdot\delta/x_j) \geq (x_i/x_j)^{-1/2}$, then we have

$$\frac{X_i(t+1)}{X_j(t+1)} \geq \left(\frac{x_i}{x_j}\right)^2 \cdot \left(1 - \frac{2\sqrt{n}\delta}{x_j}\right) \geq \left(\frac{x_i}{x_j}\right)^{1.5}.$$

and the statement follows immediately. Hence, we prove the following claim in the remaining part of the proof.

Claim 2.10. *If $x_i - x_j \geq \xi\sqrt{n \log n}$ and $x_j \geq c_w\sqrt{n \log n}$ then,*

$$1 - \left(\frac{x_i}{x_j}\right)^{-1/2} \geq \frac{2\sqrt{n}\delta}{x_j}.$$

First assume that $x_i/x_j \geq 2$.

Then

$$1 - \left(\frac{x_i}{x_j}\right)^{-1/2} \geq 1/4 \geq \frac{2\delta}{c_w\sqrt{\log n}} \geq \frac{2\sqrt{n}\delta}{x_j}$$

Next, assume that $x_i/x_j < 2$. Let $1 > \varepsilon > 0$ be defined s.t. $x_i/x_j = 1 + \varepsilon$. Observe that

$$1 - \left(\frac{x_i}{x_j}\right)^{-1/2} = 1 - (1+\varepsilon)^{-1/2} \geq \varepsilon/10 \tag{2.5}$$

By using the assumptions, it follows that

$$\varepsilon/10 = \frac{x_i - x_j}{10x_j} \geq \frac{\xi\sqrt{n \log n}}{10x_j} \geq \frac{c_w\sqrt{n \log n}}{x_j} \geq \frac{2\sqrt{n}\delta}{x_j}. \tag{2.6}$$

Then, it follows by combining (Eq. (2.5)) and (Eq. (2.6)). □

Recall that opinion j is insignificant in configuration $\mathbf{x}$ if $x_j < x_{max} - \xi\sqrt{n \log n}$, and $\mathcal{S}(\mathbf{x})$ is the set of significant opinions in configuration $\mathbf{x}$. Note that any significant opinion can win in our setting, and the largest opinion (which provides $\mathcal{X}_{max}$) can change over time. The next lemma shows that if an opinion becomes insignificant, it cannot become significant again w.h.p.

Lemma 2.11. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$. Then, $\mathcal{S}(\mathbf{X}(t+1)) \subseteq \mathcal{S}(\mathbf{x}(t))$ w.h.p.*

Proof. Recall that an opinion j is insignificant if $x_{max} - x_j > \xi \cdot \sqrt{n \cdot \log n}$.

Let $A(t)$ be the set of insignificant opinions with support larger than $4 \cdot c_w \cdot \sqrt{n \cdot \log n}$ and let $B(t)$ be the set of insignificant opinions with support smaller than $4 \cdot c_w \cdot \sqrt{n \cdot \log n}$ and larger than zero. We show that every member of $A(t)$ and $B(t)$ remains insignificant at the start of phase $t+1$.

First, we deal with $A(t)$ by fixing an opinion $j \in A(t)$. Note that $A(t) = \emptyset$ directly implies this case's statement. We lower bound the support of opinion with maximum support, and we upper bound the support of opinion j. To do so, we apply Lemma 2.7 to both opinions with $\delta = \sqrt{(\ln 7 + 2\log n)/\varepsilon^*}$ and yield

$$\Pr\left[X_{max}(t+1) \geq \frac{x_{max}^2}{\psi} - \frac{x_{max}}{\psi} \cdot \sqrt{n} \cdot \delta\right] \geq 1 - n^{-2} \tag{2.7}$$

and

$$\Pr\left[X_j(t+1) \leq \frac{x_j^2}{\psi} + \frac{x_j}{\psi} \cdot \sqrt{n} \cdot \delta\right] \geq 1 - n^{-2}.$$

By a simple union bound, we have with probability at least $1 - 2 \cdot n^{-2}$ that

$$\begin{aligned} X_{max}(t+1) - X_j(t+1) &> \frac{x_{max}^2 - x_j^2}{\psi} - \left(\frac{(x_{max} + x_j) \cdot \sqrt{n} \cdot \delta}{\psi}\right) \\ &= \frac{x_{max} + x_j}{\psi} \cdot ((x_{max} - x_j) - \delta). \end{aligned} \tag{2.8}$$

Since $x_{max} > \psi$ (see Observation 2.3) and $x_{max} > x_j$ we have

$$\begin{aligned} \frac{x_{max} + x_j}{\psi} \cdot ((x_{max} - x_j) - \delta) &\geq \left(1 + \frac{x_j}{x_{max}}\right) \cdot (x_{max} - x_j) - \left(1 + \frac{x_j}{x_{max}}\right) \cdot \delta \\ &= (x_{max} - x_j) + \frac{x_j}{x_{max}} \cdot (x_{max} - x_j) - \left(1 + \frac{x_j}{x_{max}}\right) \cdot \delta \\ &\geq (x_{max} - x_j) + \frac{x_j}{x_{max}} \cdot (x_{max} - x_j) - 2 \cdot \delta. \end{aligned}$$

In the observation below we show $x_j/x_{max} \cdot (x_{max} - x_j) - 2\delta > 0$. From that follows $X_{max}(t+1) - X_j(t+1) \geq (x_{max} - x_j) \geq \xi \cdot \sqrt{n \cdot \log n}$ where the last Inequality holds since opinion j is insignificant.

Observation 2.12. *Assume* $x_j > 4 \cdot c_w \cdot \sqrt{n \cdot \log n}$ *and* $x_{max} - x_j \geq \xi\sqrt{n \log n}$ *then* $\frac{x_j}{x_{max}} \cdot (x_{max} - x_j) - 2\delta > 0$.

Proof. To ease the calculations, let

$$\Delta' := (x_{max} - x_j)/\sqrt{n \log n} \quad \text{ and } \quad c_1 = \sqrt{(\ln 7 + 2\log n)/(\varepsilon^* \cdot \log n)}.$$

Recall that $\delta = c_1 \cdot \sqrt{\log n}$, then

$$\begin{aligned}
\frac{x_j}{x_{max}} \cdot (x_{max} - x_j) - 2\delta &\geq \sqrt{n \log n} \cdot \left(\frac{x_j \cdot (x_{max} - x_j)}{x_{max} \cdot \sqrt{n \cdot \log n}} - 2c_1 \right) \\
&\geq \sqrt{n \log n} \cdot \left(\frac{x_j \cdot \Delta'}{x_j + x_{max} - x_j} - 2c_1 \right) \\
&= \sqrt{n \log n} \cdot \left(\frac{\Delta'}{1 + \frac{x_{max} - x_j}{x_j}} - 2c_1 \right) \\
&\overset{(a)}{\geq} \sqrt{n \log n} \cdot \left(\frac{\Delta'}{1 + \frac{\Delta'}{4c_w}} - 2c_1 \right) \\
&= \sqrt{n \log n} \cdot \left(\frac{4c_w \cdot \Delta' - 2c_1 \cdot (4c_w + \Delta')}{4c_w + \Delta'} \right) \\
&= \sqrt{n \log n} \cdot \left(\frac{4c_w(\Delta' - 2c_1) + \Delta'(4c_w - 2c_1)}{4c_w + \Delta'} \right) \overset{(b)}{>} 0
\end{aligned}$$

where we use (a) $(x_{max} - x_j)/x_j \leq \frac{x_{max} - x_j}{4 \cdot c_w \cdot \sqrt{n \cdot \log n}} = \frac{\Delta'}{4c_w}$ due to $x_j \geq 4c_w\sqrt{n \log n}$ and (b) $\Delta' = (x_{max} - x_j)/\sqrt{n \log n} \geq \xi > 2 \cdot c_1$ and $c_w = 8\sqrt{1 + 2/\varepsilon^*} > c_1/2$ (Definition 2.5). $\square$

Another union bound application on (Eq. (2.8)) over all opinions $j \in A(t)$ yields with probability at least $1 - 2 \cdot n^{-1}$ that all opinions $j \in A(t)$ remains insignificant at the start of phase $t + 1$.

Next, we deal with $B(t)$ by fixing an opinion $j \in B(t)$ (assuming $B(t) \neq \emptyset$). Again, We lower bound the support of the largest opinion and upper bound the support of opinion j. For the largest opinion we have (Eq. (2.7)) where for opinion j we get by Lemma 2.8 for $c = 4 \cdot c_w$ that

$$\Pr\left[X_j(t+1) < \frac{n}{\psi} \cdot (192 \cdot c_w^2 + 74c_p) \cdot \log n \right] \geq 1 - 4n^{-2}. \tag{2.9}$$

An application of the union bound yields with probability at least $1 - 5n^{-2}$ that

$$\begin{aligned} X_{max}(t+1) - X_j(t+1) &\geq \frac{x_{max}^2}{\psi} - \frac{x_{max}}{\psi} \cdot \sqrt{n} \cdot \delta - \frac{n}{\psi} \cdot (192 \cdot c_w^2 + 74c_p) \cdot \log n \qquad (2.10) \\ &= \frac{x_{max}^2}{\psi} \cdot \left(1 - \frac{\delta \cdot \sqrt{n}}{x_{max}} - \frac{(192 \cdot c_w^2 + 74c_p) \cdot n \cdot \log n}{x_{max}^2}\right) \\ &\geq x_{max} \cdot \left(1 - \frac{\delta \cdot \sqrt{n}}{\sqrt{n} \cdot \log n} - \frac{(192 \cdot c_w^2 + 74c_p) \cdot n \cdot \log n}{n \cdot \log^2 n}\right) \\ &\geq x_{max} \cdot \left(1 - \frac{\delta}{\log n} - \frac{(192 \cdot c_w^2 + 74c_p)}{\log n}\right) \\ &\geq \xi \cdot \sqrt{n \log n} \end{aligned}$$

where we use $x_{max} > \psi$ (see Observation 2.3) and $x_{max} \geq n/k \geq \sqrt{n} \cdot \log n$. Another union bound application on (Eq. (2.10)) over all opinions $j \in B(t)$ yields with probability at least $1 - 5 \cdot n^{-1}$ that all opinions $j \in B(t)$ remains insignificant at the start of phase $t + 1$. □

2.4. Analysis of the Unbiased Case

This section provides the necessary statements to prove one of the main results Theorem 2.1. The analysis is inspired by the general approach from [52], where the authors analyze the majority process for k opinions. Opinions are classified as *strong*, *weak*, or *super-weak*, depending on their support. The authors of [52] divide time into epochs of length $\mathrm{O}\big((\frac{5}{6})^i \cdot k \log n\big)$ and show that at the end of the i-th epoch the support of the largest opinion grows by a constant factor and the fraction of *non*-super-weak opinions decreases by a constant factor. As super-weak opinions remain super-weak, this implies that eventually, consensus is reached in time $\mathrm{O}\big(k \log n\big)$.

Our approach is different and exploits the properties of the *undecided state dynamics (USD)*, which, for example, allows us to avoid both epochs of different lengths and a total runtime that is linear in k. Throughout our analysis, we consider all pairs of opinions. We show that during $\mathrm{O}\big(\log n\big)$ phases, at least one opinion in each pair becomes weak and, eventually, super-weak. If both opinions in a pair are initially strong, we apply (similar to [52]) the drift result of [46] to show that their support drifts apart. Hence, only one of the strong opinions prevails, which will be adopted by every agent within a constant number of additional phases.

Proposition 2.13. *Assume* $\mathbf{X}(t)$ *is a configuration with* $k < \sqrt{n}/\log n$ *opinions. Then, after* $\mathrm{O}\big(\log n\big)$ *phases, all agents agree on some opinion* $i \in \mathcal{S}(X(t))$*, w.h.p.*

Proof. First, we show that all agents agree on one opinion. We fix $\mathbf{X}(t) = \mathbf{x}(t)$ and consider two arbitrary opinions i and j. If both opinions are strong Lemma 2.14 shows that one of them becomes weak or super-weak within $\mathrm{O}\big(\log n\big)$ phases with probability at least $1 - \mathrm{O}\big(n^{-1.9}\big)$. As soon as either i or j are weak Lemma 2.16 shows that the weak opinion becomes super-weak within the next $\mathrm{O}(\log\log n)$ phases and remains super-weak for the rest of the process (Lemma 2.16). This happens again with a probability of $1-\mathrm{O}(n^{-1.9})$. Hence, after $\mathrm{O}(\log n)$ phases either i or j are super-weak. Since we have at most $k^2 \leq (\sqrt{n}/\log n)^2 = o(n)$ pairs of distinct opinions, we can apply the union bound over all such pairs to show that all but a single opinion are super-weak within $t' = \mathrm{O}(\log n)$ phases w.h.p. In Lemma 2.17, we show that the single remaining non-super-weak opinion wins within two additional phases.

It remains to show that the winning opinion is one of the initially significant opinions. Recall that $\mathcal{S}(X(t))$ denotes the set of all significant opinions at the start of phase t. We show for $t' = \mathrm{O}(\log n)$ that $\mathcal{S}(X(t+t')) \subseteq \mathcal{S}(X(t))$. This means that no insignificant opinion in phase t can become significant in phase $t+t'$. As the subset relation is transitive, we have that

$$\begin{aligned}\Pr\left[\mathcal{S}\big(X(t+t')\big) \subseteq \mathcal{S}(X(t))\right] &\geq \Pr\left[\forall t < t_1 \leq t+t' : \mathcal{S}(X(t_1)) \subseteq \mathcal{S}(X(t_1-1))\right] \\ &= 1 - \Pr\left[\exists t < t_1 \leq t+t' : \mathcal{S}(X(t_1)) \not\subseteq \mathcal{S}(X(t_1-1))\right].\end{aligned}$$

Furthermore, from Lemma 2.11 we have that for any $t_1 \geq 0$ it holds $\Pr[\mathcal{S}(X(t_1)) \subseteq \mathcal{S}(X(t_1-1))] \geq 1 - n^{-\Omega(1)}$. Together with the Union bound application, this implies that

$$\begin{aligned}&1 - \Pr\left[\exists t < t_1 \leq t+t' : \mathcal{S}(X(t_1)) \not\subseteq \mathcal{S}(X(t_1-1))\right] \\ &\geq 1 - \sum_{t_1=t+1}^{t+t'} \Pr\left[\mathcal{S}(X(t_1)) \not\subseteq \mathcal{S}(X(t_1-1))\right] \\ &\geq 1 - t' \cdot n^{-\Omega(1)} \geq 1 - n^{-\Omega(1)}.\end{aligned}$$

Therefore, we have w.h.p. that $\mathcal{S}(X(t+t')) \subseteq \mathcal{S}(X(t))$. In the first part of the proof, we also established that w.h.p., only a single opinion i remains in phase $t+t'$. Clearly, the remaining opinion i is significant in $\mathbf{X}(t+t')$. In other words, Opinion i belongs to the set $\mathcal{S}(X(t+t'))$, which is a subset of $\mathcal{S}(X(t))$ w.h.p. Therefore, Opinion i is also in $\mathcal{S}(X(t))$, and the result follows. $\square$

Lemma 2.14. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$ and any two distinct strong opinions i and j. Then, at least one of them will become weak or super-weak within $\mathrm{O}(\log n)$ phases with probability at least $1-\mathrm{O}(n^{-1.9})$.*

Proof. First, we show that at least one of the opinions i and j will become weak or super-weak. We consider the difference between opinion i and opinion j via a case study. If the difference is $\mathrm{o}(\sqrt{n\log n})$, we apply the drift result from Theorem A.8 to increase the difference up to $\Omega(\sqrt{n\log n})$. To be more precise, we map the difference $|X_i(t) - X_j(t)|$ to the state space of $W(t) = \lfloor |X_i(t)-X_j(t)|/(c_a \cdot \sqrt{n}) \rfloor \in \{0, \dots, (\xi/c_a)\cdot\sqrt{\log n}\}$ where c_a originates from Lemma 2.15. Observe that, for some t', $W(t') = (\xi/c_a)\cdot\sqrt{\log n}$ implies that $|X_i(t') - X_j(t')| \geq \xi \cdot \sqrt{n\log n}$. Now, we deal with the two requirements within the drift result with the help of Lemma 2.15. The first requirement is fulfilled by the first result in Lemma 2.15. That is,

$$\Pr\left[W(t+1) \geq 1\right] \geq \Pr\left[\lfloor |X_i(t+1) - X_j(t+1)| \rfloor \geq c_a \cdot \sqrt{n}\right] = \Omega(1).$$

The second requirement is fulfilled by the second result in Lemma 2.15. Assuming $c_a \cdot \sqrt{n} \leq |X_i(t) - X_j(t)| \leq \xi \cdot \sqrt{n \log n}$, it holds for a suitable constant $c_2 > 0$ that

$$\begin{aligned}
&\Pr\left[W(t+1) \geq \min\{(5/4) \cdot W(t), (\xi/c_a) \cdot \sqrt{\log n}\}\right] \\
&\quad \geq \Pr\left[|X_i(t+1) - X_j(t+1)| \geq \min\{(1+\varepsilon) \cdot |x_i - x_j|, \xi \cdot \sqrt{n \log n}\}\right] \\
&\quad \geq 1 - 14 \cdot \exp\left(-\varepsilon^* \cdot (|x_i - x_j|)^2/16n\right) \\
&\quad \geq 1 - 14 \cdot \exp\left(-\varepsilon^*/16 \cdot c_a \cdot (|x_i - x_j|)/\sqrt{n} \cdot (|x_i - x_j|)/c_a\sqrt{n}\right) \\
&\quad \geq 1 - \exp\left(-c_2 \cdot W(t)\right)
\end{aligned}$$

Thus, due to the drift result, the difference between opinion i and j is at least $\xi \cdot \sqrt{n \log n}$ in $\mathrm{O}(\log n)$ phases.

Now, we assume that $|X_i(t) - X_j(t)| \geq \xi \cdot \sqrt{n \log n}$. Since both opinions are strong and their difference is sufficiently large, we apply Lemma 2.9, which yields with probability at least $1 - \mathrm{O}(n^{-2})$ that

$$\frac{X_i(t+1)}{X_j(t+1)} \geq \left(\frac{x_i}{x_j}\right)^{1.5}. \tag{2.11}$$

Our goal is to apply the above result over multiple phases repeatedly. We establish that Lemma 2.9 may also be applied in the following phase. To this end, we need to check the two conditions that fulfill the requirements of Lemma 2.9: (i) $X_i(t+1) - X_j(t+1) \geq \xi \cdot \sqrt{n \log n}$ holds, and (ii) opinions i and j remain strong opinions in $\mathbf{X}(t+1)$. Note that the lemma statement immediately follows in case (ii) is violated. In the following, we will establish that (i) indeed holds. We apply Lemma 2.7 to both i and j and set $\delta = \sqrt{(\ln 7 + 2 \log n)/\varepsilon^*}$. This way, we get, w.h.p., that

$$\begin{aligned}
X_i(t+1) - X_j(t+1) &\geq \frac{1}{\psi} \cdot \left(x_i^2 - x_j^2 - (x_i + x_j) \cdot \sqrt{(\ln 7 + 2\log n)/\varepsilon^*} \cdot \sqrt{n}\right) \\
&\geq \frac{x_i + x_j}{\psi} \cdot \left((x_i - x_j) - \sqrt{(\ln 7 + 2\log n)/\varepsilon^*} \cdot \sqrt{n}\right)
\end{aligned}$$

Since both opinions are strong and $\psi \leq x_{max}$ (see Observation 2.3), it follows that $(x_i + x_j)/\psi > 9/5$. This, together with the above inequality chain, implies that, indeed, $X_i(t+1) - X_j(t+1) > \xi \cdot \sqrt{n \log n}$. The above argument can easily be translated into an induction, which yields that $X_i(t+t') - X_j(t+t') \geq \xi \cdot \sqrt{n \log n}$ until a round $t + t'$ is reached where opinion j becomes weak (i.e., condition (ii) above is violated). Note that, w.h.p., $t' = \mathrm{O}(\log n)$ must hold as otherwise, it follows by a repeated application of (Eq. (2.11)) that

$$\frac{X_i(t+t')}{X_j(t+t')} \geq \left(1 + \frac{\xi \cdot \sqrt{n \log n}}{x_j}\right)^{1.5^{t'}} \gg \left(1 + \frac{1}{\sqrt{n}}\right)^{1.5^{t'}} > n.$$

Hence opinion j will become either weak or super-weak within $\mathrm{O}(\log n)$ phases with probability at least $1 - \mathrm{O}(n^{-1.9})$. $\square$

Lemma 2.15. *Fix* $\mathbf{X}(t) = \mathbf{x}(t)$ *and two distinct strong opinions* i *and* j. *Let* $c_a = \max\{1/\varepsilon_p\ ,\ 100\}$.

1. *If* $|x_i(t) - x_j(t)| < c_a \cdot \sqrt{n}$, *then* $\Pr[|X_i(t+1) - X_j(t+1)| \geq c_a \cdot \sqrt{n}] = \Omega(1)$.
2. *If* $c_a \cdot \sqrt{n} \leq |x_i(t) - x_j(t)| < \xi \cdot \sqrt{n \log n}$, *then*

$$\Pr[|X_i(t+1) - X_j(t+1)| \geq (5/4) \cdot (x_i(t) - x_j(t))] \geq 1 - 14 \cdot e^{-\varepsilon^* \cdot \frac{|x_i(t) - x_j(t)|^2}{16n}}.$$

Proof. We track the difference between two strong opinions, i and j, throughout a single phase. We start with the first statement and assume, w.l.o.g., that $x_i \geq x_j$. The idea is to apply an anti-concentration result during the decision part to establish a sufficiently large difference between the support of both opinions. Additionally, we want to roughly maintain such difference throughout the following boosting part.
At first we analyze the decision part and consider $Y_i(t), Y_j(t)$ and $\|\mathbf{Y}(t)\|_1$. Recall (see Observation 2.3) that $Y_i(t) \sim \mathrm{Bin}(x_i, x_i/n)$, $Y_j(t) \sim \mathrm{Bin}(x_j, x_j/n)$ and $\|\mathbf{Y}(t)\|_1$ can be modeled as a sequence of n Poisson trials. We define the event $\mathcal{E}$ as follows

$$\mathcal{E} = \Bigg\{ Y_i(t) \geq \frac{x_i^2}{n} + \frac{10}{7} \cdot c_a \cdot \frac{x_i}{\sqrt{n}} \text{ and } \left(1 - \frac{2}{\sqrt{\log n}}\right) \cdot \frac{x_j^2}{n} \leq Y_j(t) \leq \frac{x_j^2}{n} - \frac{10}{7} \cdot c_a \cdot \frac{x_j}{\sqrt{n}}$$
$$\text{and } \|\mathbf{Y}(t)\|_1 = \psi \cdot \left(1 \pm \frac{6}{\sqrt{\log n}}\right) \Bigg\}.$$

Now, we bound the probability of $\bar{\mathcal{E}}$. We apply Lemma A.15 to both opinions i and j with $\delta_i = ((10/7) \cdot c_a \sqrt{n})/x_i$ and $\delta_j = ((10/7) \cdot c_a \sqrt{n})/x_j$ and yield

$$\Pr\left[Y_i(t) \geq \frac{x_i^2}{n} + \frac{10}{7} \cdot c_a \cdot \frac{x_i}{\sqrt{n}}\right] \geq \exp\left(-9 \cdot ((10/7) \cdot c_a)^2\right) \geq \exp\left(-20 \cdot c_a^2\right) \tag{2.12}$$

$$\Pr\left[Y_j(t) \leq \frac{x_j^2}{n} - \frac{10}{7} \cdot c_a \cdot \frac{x_j}{\sqrt{n}}\right] \geq \exp\left(-9 \cdot ((10/7) \cdot c_a)^2\right) \geq \exp\left(-20 \cdot c_a^2\right) \tag{2.13}$$

Next, we apply Chernoff Bound(Theorem A.1) to $\|\mathbf{Y}(t)\|_1$ and $Y_j(t)$ where in the latter case, we only need an additional lower bound. Note that $\mathbb{E}[\|\mathbf{Y}(t)\|_1] \geq x_{max}^2/n = \Omega(\log^2 n)$ due to $x_{max} \geq n/k$ and $k \leq \sqrt{n}/\log n$. Hence, for $\delta' = 6/\sqrt{\log n}$ we get

$$\Pr\Big[\|\mathbf{Y}(t)\|_1 \leq \psi \cdot \big(1 - \delta'\big)\Big] \leq n^{-2} \text{ and } \Pr\Big[\|\mathbf{Y}(t)\|_1 \geq \psi \cdot \big(1 + \delta'\big)\Big] \leq n^{-2}.$$

In the case of $Y_j(t)$, we use the fact that opinion j is strong and hence, similar to the previous case, $x_j \geq 0.9 \cdot x_{max} \geq 0.9 \cdot \sqrt{n} \log n$. Again we apply Chernoff Bound(Theorem A.1) for $\delta' = \sqrt{2 \cdot n \cdot \log n}/x_j \leq 4/\sqrt{\log n}$ and yield

$$\Pr\left[Y_j(t) \leq \left(1 - \frac{4}{\sqrt{\log n}}\right) \cdot \frac{x_j^2}{n}\right] \leq \Pr\left[Y_j(t) \leq \left(1 - \frac{\sqrt{2 \cdot n \cdot \log n}}{x_j}\right) \cdot \frac{x_j^2}{n}\right] \leq n^{-2}.$$

An application of the union bound on the bounds of $Y_j(t)$ yields

$$\begin{aligned}\Pr\left[Y_j(t)\notin\left(\left(1-\frac{4}{\sqrt{\log n}}\right)\cdot\frac{x_j^2}{n},\frac{x_j^2}{n}-\frac{10}{7}\cdot c_a\cdot\frac{x_j}{\sqrt{n}}\right)\right]&\leq 1-\exp\left(-9\cdot((10/7)\cdot c_a)^2\right)+n^{-2}\\&\leq 1-\left(\exp\left(-10\cdot c_a^2\right)-n^{-2}\right).\end{aligned}$$

Since $Y_i(t)$ and $Y_j(t)$ are independent, another application of the Union bound yields

$$\Pr\left[\bar{\mathcal{E}}\right]\leq 1-\left(\left(\exp\left(-10\cdot c_a^2\right)-n^{-2}\right)\cdot\exp\left(-20\cdot c_a^2\right)\right)+2n^{-2}=p<1 \tag{2.14}$$

where $p<1$ is a constant probability.
Now, we deal with the outcome of the boosting part conditioned on the event $\mathcal{E}$. We fix $Y_i(t)=y_i$, $Y_j(t)=y_j$ and define $d=\|\mathbf{Y}(t)\|_1$. Again, recall (see Observation 2.3 and Observation 2.4) that $X_i(t+1)\sim\mathrm{PE}(y_i,d-y_i,n-d)$ and $X_j(t+1)\sim\mathrm{PE}(y_j,d-y_j,n-d)$. We apply the tail bound for this Pólya Eggenberger distribution from Theorem A.11. Note that conditioned on the event $\mathcal{E}$ and the previous observations about strong opinions we have $\sqrt{y_i}\geq\sqrt{y_j}\geq 0.9\cdot\sqrt{1-4/\sqrt{\log n}}\cdot\log n$. Thus, clearly $\delta=(2\cdot c_a)/7<\sqrt{y_j}$ and hence we get

$$\Pr\left[X_i(t+1)<y_i\cdot\frac{n}{d}-\sqrt{y_i}\cdot\delta\cdot\frac{n}{d}\;\middle|\;\mathcal{E}\right]\leq 4\exp\left(-\varepsilon_p\cdot\delta^2\right)\leq 4\exp\left(-4\cdot\varepsilon_p\cdot c_a^2/49\right) \tag{2.15}$$

$$\Pr\left[X_j(t+1)>y_j\cdot\frac{n}{d}+\sqrt{y_j}\cdot\delta\cdot\frac{n}{d}\;\middle|\;\mathcal{E}\right]\leq 4\exp\left(-\varepsilon_p\cdot\delta^2\right)\leq 4\exp\left(-4\cdot\varepsilon_p\cdot c_a^2/49\right) \tag{2.16}$$

Now we show that the difference between $X_i(t+1)$ and $X_j(t+1)$ is still sufficiently large. Conditioned on the event $\mathcal{E}$ we get that

$$\begin{aligned}X_i(t+1)-X_j(t+1)&\geq\frac{n}{d}\cdot\left(y_i-y_j-\delta\cdot(\sqrt{y_i}+\sqrt{y_j})\right)\\&\geq\frac{n}{d}\cdot\left(\frac{x_i^2-x_j^2}{n}+\frac{10}{7}\cdot c_a\cdot\frac{x_i+x_j}{\sqrt{n}}\right.\\&\qquad\left.-\delta\cdot\frac{x_i}{\sqrt{n}}\cdot\left(\sqrt{1+\frac{10\cdot c_a\sqrt{n}}{7\cdot x_i}}+\sqrt{1-\frac{10\cdot c_a\sqrt{n}}{7\cdot x_j}}\right)\right)\end{aligned}$$

By the definition of strong opinions and the inequality $\sqrt{1-z}+\sqrt{1+z}\leq 2$ for $z\in(-1,+1), z\neq 1$ it follows that

$$\begin{aligned}X_i(t+1)-X_j(t+1)&\geq\frac{n}{d}\cdot\left(\frac{10}{7}\cdot c_a\cdot\frac{x_i+x_j}{\sqrt{n}}-2\cdot\delta\cdot\frac{x_i}{\sqrt{n}}\right)\\&\geq\sqrt{n}\cdot\frac{x_{max}}{d}\cdot\left(2\cdot\frac{9}{10}\cdot(10/7)\cdot c_a-2\cdot\delta\right)\\&\geq 2\cdot\sqrt{n}\cdot\frac{x_{max}}{d}\cdot\left(\frac{9}{7}\cdot c_a-\delta\right)\end{aligned}$$

Again by the event $\mathcal{E}$ and the fact that $\psi \leq x_{max}$ (see Observation 2.3) it follows that

$$X_i(t+1) - X_j(t+1) \geq \sqrt{n} \cdot \frac{2 \cdot x_{max}}{(1+6/\sqrt{\log n}) \cdot x_{max}} \cdot \left(\frac{9}{7} \cdot c_a - \frac{2}{7} \cdot c_a\right) \geq c_a \cdot \sqrt{n}.$$

An application of the Union bound yields that the difference between $X_i(t+1)$ and $X_j(t+1)$ holds with probability $1 - 8\exp(-4 \cdot \varepsilon_p \cdot c_a^2/49)$. At last, we combine this with (Eq. (2.14)) by applying the law of total probability to deduce that the first statement follows with constant probability. Due to the choice of c_a, it holds, with at least constant probability, that

$$\begin{aligned}
&\Pr\left[|X_i(t+1) - X_j(t+1)| \geq c_a \cdot \sqrt{n}\right] \\
&\quad = \Pr\left[|X_i(t+1) - X_j(t+1)| \geq c_a \cdot \sqrt{n} \;\middle|\; \mathcal{E}\right] \cdot \Pr[\mathcal{E}] \\
&\qquad + \Pr\left[|X_i(t+1) - X_j(t+1)| \geq c_a \cdot \sqrt{n} \;\middle|\; \bar{\mathcal{E}}\right] \cdot \Pr[\bar{\mathcal{E}}] \\
&\quad \geq \Pr\left[|X_i(t+1) - X_j(t+1)| \geq c_a \cdot \sqrt{n} \;\middle|\; \mathcal{E}\right] \cdot \Pr[\mathcal{E}] \\
&\quad \geq \left(1 - 8 \cdot \exp\left(-4 \cdot \varepsilon_p \cdot c_a^2/49\right)\right) \cdot \left(\left(\left(\exp\left(-10 \cdot c_a^2\right) - n^{-2}\right) \cdot \exp\left(-20 \cdot c_a^2\right)\right) - 2n^{-2}\right)
\end{aligned}$$

We continue with the proof of the second statement and assume $x_i - x_j \geq c_a \cdot \sqrt{n}$. We apply Lemma 2.7 to both opinions with $\delta = (x_i - x_j)/4 \cdot \sqrt{n}$ and yield

$$\Pr\left[X_i(t+1) \leq \frac{x_i^2}{\psi} - \frac{x_i}{\psi} \cdot \sqrt{n} \cdot \delta\right] \leq 7 \cdot \exp\left(-\varepsilon^* \cdot \delta^2\right) = 7 \cdot \exp\left(-\varepsilon^* \cdot \frac{(x_i - x_j)^2}{16n}\right)$$
$$\Pr\left[X_j(t+1) \geq \frac{x_j^2}{\psi} + \frac{x_j}{\psi} \cdot \sqrt{n} \cdot \delta\right] \leq 7 \cdot \exp\left(-\varepsilon^* \cdot \delta^2\right) = 7 \cdot \exp\left(-\varepsilon^* \cdot \frac{(x_i - x_j)^2}{16n}\right)$$

As i and j are assumed to be strong opinions, it follows that $(x_i + x_j)/\psi > 9/5$. An application of the union bound yields the second statement, with probability at least $1 - 14 \cdot \exp\left(-\varepsilon^* \cdot (|x_i - x_j|)^2/4n\right)$, due to

$$X_i(t+1) - X_j(t+1) \geq \frac{x_i + x_j}{\psi} \cdot \left(1 - \frac{\sqrt{n} \cdot \delta}{x_i - x_j}\right) \cdot (x_i - x_j) \geq \frac{5}{4} \cdot (x_i - x_j). \qquad \square$$

Lemma 2.16. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$ and an opinion j. If opinion j is weak, then it will become super-weak in $\mathrm{O}(\log\log n)$ phases with probability at least $1 - \mathrm{O}(n^{-1.9})$. If opinion j is super-weak, then it will remain super-weak in at least $\Omega(\log^2 n)$ following phases with probability more than $1 - \mathrm{O}(n^{-1.9})$.*

Proof. First, we show that a weak opinion will become super-weak. Let opinion j be weak but not super-weak opinion in $\mathbf{x}$, i.e., $c_w \cdot \sqrt{n \log n} < x_j < 0.9 \cdot x_{max}$. As opinion j is weak, it follows that the difference $x_{max} - x_j \geq n/(10 \cdot k) = \omega(\sqrt{n \log n})$ is large enough, and we can apply Lemma 2.9. This yields w.h.p. that

$$\frac{X_{max}(t+1)}{X_j(t+1)} \geq \left(\frac{x_{max}}{x_j}\right)^{1.5}. \tag{2.17}$$

This result has two implications. For one, it states that the ratio between the largest opinion and opinion j grows substantially. Second, it implies that opinion j cannot become strong in $\mathbf{X}(t+1)$. To see this, we combine (Eq. (2.17)) together with the fact that j is weak in $\mathbf{x}$

$$X_j(t+1) \overset{(Eq.\ (2.17))}{\leq} X_{max}(t+1) \cdot \left(\frac{x_j}{x_{max}}\right)^{1.5} < X_{max}(t+1) \cdot \left(\frac{9}{10}\right)^{1.5} < X_{max}(t+1) \cdot 0.9.$$

Hence, j is either weak or super-weak in phase $t+1$. If j is super-weak, we are done. Otherwise, we may again apply Lemma 2.9. We follow this approach for $t' = \log_{1.5} \log_{9/10} n$ phases, at which point, either (i) j already became super-weak in some phase $< t + t'$, or (ii) it follows by the growth of the ratio that with probability at least $1 - \mathrm{O}\big(n^{-1.9}\big)$

$$\frac{X_{max}(t+t')}{X_i(t+t')} \geq \left(\frac{10}{9}\right)^{1.5^{t'}} \geq n.$$

This implies that opinion j must have already become super-weak.

In this second part of the proof, we show that a super-weak opinion j in $\mathbf{x}$ remains super-weak. We apply Lemma 2.8 on this opinion j and yield

$$\Pr\left[X_j(t+1) > \frac{n}{\psi} \cdot (12c_w^2 + 74c_p) \log n\right] < 4n^{-2}. \tag{2.18}$$

Using that $\psi = \mathbb{E}\left[\|\mathbf{Y}(t)\|_1\right] \geq n/k \geq \sqrt{n} \log n$ then implies

$$\Pr\left[X_j(t+1) > (12c_w^2 + 74c_p) \cdot \sqrt{n}\right] < 4n^{-2}.$$

As $(12c_w^2 + 74c_p) \cdot \sqrt{n} = \mathrm{o}\big(\sqrt{n \log n}\big)$ it follows that opinion j remains super-weak at the start phase $t+1$. A repetition of this argument, together with a union bound application, yields that opinion j will remain super-weak for at least $\Omega\big(\log^2 n\big)$ phases with probability more than $1 - \mathrm{O}\big(n^{-1.9}\big)$. $\square$

Lemma 2.17. *Fix $\mathbf{X}(t) = \mathbf{x}(t)$. If all but a single opinion i are super-weak, then only opinion i remains in phase $t+2$ w.h.p.*

Proof. Let opinion i be the only non super-weak opinion. It follows that $x_i \geq n-(k-1) \cdot c_w \cdot \sqrt{n \log n}$. As we consider $k \leq \sqrt{n}/\log n$, this implies that $x_i \geq n - n \cdot (c_w/\sqrt{\log n}) = n \cdot (1 - \mathrm{o}\big(1\big))$. Furthermore, as $\psi = \mathbb{E}\left[\|\mathbf{Y}(t)\|_1\right] = \sum_{j=1}^{k} x_j^2/n \geq x_i^2/n$ this also implies that $\psi = \Omega\big(n\big)$.

Now we fix some opinion j that is super-weak. We apply Lemma 2.8 and use that $\psi = \Omega\big(n\big)$ which immediately yields

$$\Pr\left[X_j(t+1) = \omega(\log n)\right] < 4n^{-2}.$$

Such an opinion j will, w.h.p., not have a single decided agent at the start of boosting part $t+1$ because

$$\begin{aligned}\Pr\left[Y_j(t+1) > 0 \;\middle|\; X_j(t+1) = \mathrm{O}\big(\log n\big)\right] &< 1 - \left(1 - \frac{\mathrm{O}\big(\log n\big)}{n}\right)^{\mathrm{O}\left(\log n\right)} \\ &< \frac{\operatorname{polylog} n}{n}.\end{aligned}$$

Therefore, j will vanish before phase $t+2$ w.h.p. A simple union bound argument now yields that, w.h.p., *none* of the super-weak opinions at the beginning of phase t will survive until phase $t+2$. □

2.5. Analysis of the Biased Case

The proof of Theorem 2.2 closely resembles the proofs conducted in [29, 53]. Unfortunately, it is not sufficient to adapt the proof of Theorem 2.1 since the number of phases for the biased case is $\mathrm{O}\big(\log n\big)$, which is too high. Roughly speaking, we track the ratio between the largest and second-largest opinions over the phases. For this, we rely on the following two technical results (Lemma 2.18 and Lemma 2.22).

Lemma 2.18. *Fix $\mathbf{X}(0) = \mathbf{x}$ and let $i^* := \arg\max_i\{x_i\}$. Assume $\mathbf{x}$ has an additive bias of at least $\xi\sqrt{n\log n}$ and a multiplicative bias of α.*

1. *If $k \le \sqrt{n}/\log n$, then all agents agree on opinion i^* in $\mathrm{O}\big(\log\log_\alpha n\big)$ phases, w.h.p.*
2. *If $k > \sqrt{n}/\log n$, then for $t^* = \mathrm{O}\big(\log\log_\alpha n + \log\log n\big)$ we have $X_{i^*}(t^*) > (5/8)\cdot n$ w.h.p.*

Proof. In this proof, we will use the following notions for some configuration $\mathbf{X}(t)$: we let $X_{max}(t)$ and $X_{sec}(t)$ denote the largest and second-largest opinion in $\mathbf{X}(t)$. Furthermore, we define

$$\alpha(t) := X_{max}(t)/X_{sec}(t) \quad \text{and} \quad \gamma(t) := \min\Big\{\alpha(t)\ ,\ X_{max}(t)/(c_w\sqrt{n\log n}\Big\}.$$

We first show three intermediate results to facilitate the proof: Claim 2.19, Claim 2.20, and Claim 2.21.

Claim 2.19. *Fix $\mathbf{X}(t) = \mathbf{x}$. If $x_{max} - x_{sec} > \xi\sqrt{n\log n}$ then*

$$\alpha(t+1) > \gamma(t)^{3/2} \tag{2.19}$$

Proof. We start by considering a fixed opinion j with $x_{max} > x_j \ge c_w\sqrt{n\log n}$. An application of Lemma 2.9 immediately yields with probability $1 - n^{-2}$ that

$$\frac{X_{max}(t+1)}{X_j(t+1)} \overset{\text{Lemma 2.9}}{\ge} \left(\frac{x_{max}(t)}{x_j(t)}\right)^{1.5} \ge \alpha(t)^{1.5} \ge \gamma(t)^{3/2}. \tag{2.20}$$

On the other hand, consider now some fixed opinion j with $x_j < c_w\sqrt{n\log n}$. We lower and upper bound the support of the largest opinion and j, respectively. We define $c_1 := \sqrt{(\ln 7 + 2\log n)/(\varepsilon^* \log n)}$ and apply Lemma 2.7 with $\delta = c_1\sqrt{\log n}$ and Lemma 2.8 with $c = c_w$ we get that with probability at least $1 - 5n^{-2}$

$$
\begin{aligned}
\frac{X_{max}(t+1)}{X_j(t+1)} &\geq \frac{\frac{x_{max}^2(t)}{\psi} \cdot \left(1 - \frac{c_1 \cdot \sqrt{n\log n}}{x_{max}(t)}\right)}{\frac{n}{\psi} \cdot (12c_w^2 + 74c_p) \cdot \log n} \\
&= \left(\frac{x_{max}(t)}{\sqrt{(12c_w^2 + 74c_p) \cdot n\log n}}\right)^2 \cdot \left(1 - \frac{c_1 \cdot \sqrt{n\log n}}{x_{max}(t)}\right) \\
&= \left(\frac{x_{max}(t)}{\sqrt{n\log n}}\right)^{3/2} \cdot \left(\frac{x_{max}(t)}{\sqrt{n\log n}}\right)^{1/2} \cdot \frac{1}{(12c_w^2 + 74c_p)} \cdot \left(1 - \frac{c_1 \cdot \sqrt{n\log n}}{x_{max}(t)}\right) \\
&\overset{(a)}{\geq} \left(\frac{x_{max}(t)}{\sqrt{n\log n}}\right)^{3/2} \cdot \sqrt{\xi} \cdot \frac{1}{(12c_w^2 + 74c_p)} \cdot \left(1 - \frac{c_1}{\xi}\right) \\
&\overset{(b)}{\geq} \left(\frac{x_{max}(t)}{\sqrt{n\log n}}\right)^{3/2} \cdot \frac{\sqrt{\xi}}{24c_w^2 + 148c_p} \overset{(c)}{\geq} \left(\frac{x_{max}(t)}{\sqrt{n\log n}}\right)^{3/2} \cdot \left(\frac{1}{c_w}\right)^{3/2} \geq \gamma(t)^{3/2}.
\end{aligned}
$$

in which for (a) we use $x_{max}(t)/(c_w\sqrt{n\log n}) < x_{max}(t)/x_j(t)$ and $x_{max}(t) \geq \xi\sqrt{n\log n}$, for (b) we utilize $1 - c_1/\xi > 1/2$ and for (c) we consider $\xi = (160 \cdot c_w)^2 + (148 \cdot c_p)^2 \geq ((24c_w^2 + 148c_p)/c_w^{3/2})^2$ (see Definition 2.5). Summarizing, we show for any fixed opinion j with $x_j < x_{max}$ we have with probability at least $1 - \mathrm{O}(n^{-2})$.

$$
\frac{X_{max}(t+1)}{X_j(t+1)} \geq \gamma(t)^{3/2}.
$$

A union bound over all opinions j yields that, w.h.p., $\alpha(t+1) > \gamma(t)^{3/2}$ as desired. □

Claim 2.20. *Fix $\mathbf{X}(t) = \mathbf{x}$. If (i) $x_{max} - x_{sec} \geq \xi \cdot \sqrt{n\log n}$, (ii) $\gamma(t) \geq \xi/c_w = 1 + \Omega(1)$, and (iii) $x_{max}^2/n < x_{sec}/3$, then w.h.p.*

$$
\gamma(t+1) \geq \gamma(t)^{5/4}.
$$

Proof. We again fix a configuration $\mathbf{X}(t) = \mathbf{x}$ at the start of some phase t and assume that $\mathbf{x}$ fulfills the requirements of the claim. We will show that w.h.p.,

$$
\frac{X_{max}(t+1)}{c_w \cdot \sqrt{n\log n}} \geq \gamma(t)^{5/4}.
$$

Note that, together with the result of Claim 2.19 this also implies $\gamma(t+1) \geq \gamma(t)^{5/4}$.

To show this, we distinguish between two cases. First, consider the case of $x_{sec} < c_w \cdot \sqrt{n\log n}$. For $c_1 = \sqrt{(\ln 7 + 2\log n)/(\varepsilon^* \log n)}$ we apply Lemma 2.7 with $\delta = c_1 \cdot \sqrt{n\log n}$. This yields for $\psi = \sum_{j=1}^{k} x_j^2/n$ and with probability $1 - 5n^{-2}$ that

$$
X_{max}(t+1) \overset{\text{Lemma 2.7}}{>} \frac{x_{max}^2}{\psi} \cdot \left(1 - \frac{c_1\sqrt{n\log n}}{x_{max}}\right) > \frac{x_{max}^2}{\psi}\left(1 - \frac{c_1}{\xi}\right) > \frac{x_{max}^2}{\psi}\left(1 - \frac{1}{100}\right).
$$

Here we used in the second step that $x_{max} > \xi\sqrt{n \log n}$ as per assumption (i), and the last step follows from the definition of the constants c_1 and ξ (see Definition 2.5). To further simplify the above result, we note the following. First, $\psi = \sum_{j=1}^{k} x_j^2/n \leq x_{max}^2/n + x_{sec}$ is true for any configuration $\mathbf{x}$. Second, because we additionally assume $x_{sec} < c_w\sqrt{n \log n}$ and $x_{max}^2/n < x_{sec}/3$ this further implies $\psi \leq (1+1/3)c_w \cdot \sqrt{n \log n}$. When using this, we get

$$X_{max}(t+1) \geq \frac{x_{max}^2}{(1+1/3) \cdot c_w \cdot \sqrt{n \log n}} \cdot \left(1 - \frac{1}{100}\right) > \frac{x_{max}^2}{c_w \cdot \sqrt{n \log n}} \cdot \frac{5}{8} = x_{max} \cdot \gamma(t) \cdot \frac{5}{8}.$$

When dividing by $c_w \cdot \sqrt{n \log n}$ on both ends of this inequality chain, we get

$$\frac{X_{max}(t+1)}{c_w \cdot \sqrt{n \log n}} \geq \frac{x_{max}}{c_w \cdot \sqrt{n \log n}} \cdot \gamma(t) \cdot \frac{5}{8} \geq \gamma(t)^2 \cdot \frac{5}{8} > \gamma(t)^{5/4}.$$

In the last step, we rely on assumption (ii), which implies that $\gamma(t) > \xi/c_w \gg 100$.

We must consider the remaining case $x_{sec} \geq c_w\sqrt{n \log n}$. Let j denote an opinion with $x_j = x_{sec}$. First, we argue that $X_j(t+1) > x_j/2$. To that end we again apply Lemma 2.7 with $\delta = c_1 \cdot \sqrt{n \log n}$ which yields with probability $1 - 5n^{-2}$ that

$$X_j(t+1) \overset{Lemma\ 2.7}{>} \frac{x_j^2}{\psi} \cdot \left(1 - \frac{c_1 \cdot \sqrt{n \log n}}{x_j}\right) > \frac{x_j^2}{\psi} \cdot \frac{2}{3}.$$

In the second step, we used $x_j > c_w \cdot \sqrt{n \log n}$ and $c_w > 3c_1$ (see Definition 2.5). Just as before, we argue that $\psi \leq x_{max}^2/n + x_j < (1+1/3) \cdot x_j$ is implied by assumption (i). This implies

$$X_j(t+1) > \frac{x_j^2}{\psi} \cdot \frac{2}{3} \geq x_j \cdot \frac{1}{(1+1/3)} \cdot \frac{2}{3} \geq \frac{x_j}{2}$$

When first using that $x_j \geq c_w \cdot \sqrt{n \log n}$ followed by the above Inequality in the next step, we get

$$\frac{X_{max}(t+1)}{c_w \cdot \sqrt{n \log n}} \geq \frac{X_{max}(t+1)}{x_j} > \frac{1}{2} \cdot \frac{X_{max}(t+1)}{X_j(t+1)}. \tag{2.21}$$

Recall, throughout the proof of Claim 2.19, we established in inequality (Eq. (2.20)) that $X_{max}(t+1)/X_j(t+1) > (x_{max}/x_j)^{1.5}$ in case opinion j has $x_j > c_w \cdot \sqrt{n \log n}$. As this is indeed the case, we have with probability at least $1 - \mathrm{O}(n^{-2})$

$$\frac{X_{max}(t+1)}{X_j(t+1)} \overset{(Eq.\ (2.20))}{>} \left(\frac{x_{max}}{x_j}\right)^{3/2} \geq \gamma(t)^{3/2}$$

Now, when combining this with (Eq. (2.21)) we get with probability at least $1 - \mathrm{O}(n^{-2})$ that

$$\frac{X_{max}(t+1)}{c_w\sqrt{n \log n}} \geq \frac{1}{2} \cdot \left(\frac{X_{max}(t+1)}{X_j(t+1)}\right)^{3/2} = \frac{1}{2}\gamma(t)^{3/2} > \gamma(t)^{5/4}.$$

Where we use in the last step that per assumption (ii) $\gamma(t) > \xi/c_w \gg 100$ is large enough (see Definition 2.5). $\square$

Claim 2.21. *Assume $\mathbf{X}(0)$ has an additive bias of at least $\xi\sqrt{n\log n}$. Then, w.h.p., it holds for all $t < \log^2 n$ that*

1. $X_{max}(t) - X_{sec}(t) > \xi\sqrt{n\log n}$ *and*
2. $\mathcal{S}(\mathbf{X}(t)) = \mathcal{S}(\mathbf{X}(0))$

The first statement implies that the bias does not fall below $\xi\sqrt{n\log n}$ in the first $\log^2 n$ rounds.

Proof. To show this claim, we first consider some fixed phase t and assume that $X_{max}(t) - X_{sec}(t) > \xi\sqrt{n\log n}$. Note that this is equivalent to assuming that $|\mathcal{S}(\mathbf{X}(t))| = 1$, or in other words, assuming that there exists exactly one significant opinion. From Lemma 2.11 it now follows that $\mathcal{S}(\mathbf{X}(t+1)) \subseteq \mathcal{S}(\mathbf{X}(t))$ w.h.p. On the other hand, we know that $\mathcal{S}(\mathbf{X}(t+1)) \neq \emptyset$ as the largest opinion is always significant. When combining these two observations, we therefore get w.h.p. that $\mathcal{S}(\mathbf{X}(t+1)) = \mathcal{S}(\mathbf{X}(t))$. An inductive application of this approach yields that, w.h.p.,

$$\forall 0 < t \leq \log^2 n : \mathcal{S}(\mathbf{X}(t)) = \mathcal{S}(\mathbf{X}(0)).$$

This immediately yields the second result of Claim 2.21. The second result follows as $|\mathcal{S}(\mathbf{X}(t))| = |\mathcal{S}(\mathbf{X}(0))| = 1$, which implies the existence of an additive bias of at least $\xi\sqrt{n\log n}$ in phase t. □

We are ready to start with the proof of Lemma 2.18. We first repeatedly apply the result of Claim 2.19 until we hit a phase t_1 with $\gamma(t_1) < \alpha(t_1)$. From the definition of $\gamma(t_1)$ this is equivalent to $X_{sec}(t_1) < c_w\sqrt{n\log n}$. Recall, Claim 2.19 states that if in some phase t we have $X_{max}(t) - X_{sec}(t) > \varepsilon\sqrt{n\log n}$, then $\alpha(t+1) > \gamma(t)^{3/2}$ w.h.p. In Claim 2.21, we established that this requirement on the additive bias is fulfilled in the first $\log^2 n$ phases w.h.p. Therefore, the result of Claim 2.19 is applicable to the first $t' = \log_{3/2}\log_{\alpha(0)} n = \mathrm{O}(\log n)$ phases w.h.p. That is, w.h.p., we have

$$\forall 0 \leq t < t' : \ \alpha(t+1) \geq \gamma(t)^{3/2}.$$

Now, assume that $t_1 > t'$, i.e., we have $\gamma(t_1) < \alpha(t_1)$ for the first time in some phase after t'. Therefore, in every phase t with $t \leq t'$ we have $\gamma(t) = \alpha(t)$. This implies $\alpha(t') \geq \alpha(0)^{(3/2)^{t'}} > n$. This further implies $X_{sec}(t') < c_w\sqrt{n\log n}$ and from the definition of $\gamma(t')$, it follows that $\gamma(t') < \alpha(t')$. This is a contradiction to our assumption of $t_1 > t'$. Hence $t_1 < \log_{5/4}\log_{\alpha(0)} n = \mathrm{O}(\log n)$ w.h.p.

Starting with phase t_1, we split our analysis into two cases. First assume that $k < \sqrt{n}/\log n$. Per definition of phase t_1 we have $X_{sec}(t_1) < c_w\sqrt{n\log n}$. This implies from the definition of super-weak (see Definition 2.6) that all opinions besides the first are super-weak in $\mathbf{X}(t_1)$. From Lemma 2.17, it follows that in phase $t_1 + 2$, only a single opinion prevails w.h.p. As $t_1 = \mathrm{O}(\log n)$ w.h.p., the second statement of Claim 2.21 implies that this opinion must be the initially significant opinion. The first statement of Lemma 2.18 follows.

The case of $k \geq \sqrt{n}/\log n$ is more involved. In this case, we must follow a different approach from phase t_1 onward. In the following, we define three events. We say that a phase t fulfills

1. event $\mathcal{E}_1(t)$ iff $X_{max}(t) - X_{sec}(t) > \xi\sqrt{n \log n}$
2. event $\mathcal{E}_2(t)$ iff $\gamma(t) \geq \xi/c_w$
3. event $\mathcal{E}_3(t)$ iff $X_{max}(t)^2/n < X_{sec}(t)/3$.

Observe, if t fulfills all three events, then we may apply Claim 2.20 and get $\gamma(t+1) \geq \gamma(t)^{5/4}$. We now show that we will soon arrive at a phase $t_1 + t_2$ such that $\mathcal{E}_3(t)$ does not hold.

Recall that we have $t_1 = \mathrm{O}(\log n)$ w.h.p. It follows from Claim 2.21 that $\mathcal{E}_1(t_1)$ is fulfilled. Additionally, it follows from the definition of $\gamma(t_1)$ that $\gamma(t_1) = \frac{X_{max}(t_1)}{c_w\sqrt{n \log n}} \geq \xi/c_w > 1$. Therefore $\mathcal{E}_2(t_1)$ also holds. Now, in case $\mathcal{E}_3(t_1)$ does not hold, we are finished. Otherwise, we may apply Claim 2.20 and get w.h.p. that $\gamma(t_1+1) \geq \gamma(t_1)^{5/4}$.

Note that, in such case, $\mathcal{E}_2(t_1+1)$ will also be satisfied because $\gamma(t_1+1) \geq \gamma(t_1)^{5/4} \geq \xi/c_w$. Furthermore, Claim 2.21 still guarantees that, w.h.p., $\mathcal{E}_1(t_1+1)$ still holds. In case $\mathcal{E}_3(t_1+1)$ is fulfilled we again apply Claim 2.20 and get $\gamma(t_1+2) \geq \gamma(t_1)^{(5/4)^2}$. We repeat this approach for $t' = \log_{5/4} \log_{\xi/c_w} n = \mathrm{O}(\log\log n)$ phases. Even if there is no t with $0 \leq t < t'$ such that $\mathcal{E}_3(t_1+t)$ is violated, we have w.h.p. that

$$\gamma(t_1+t') \geq \gamma(t_1)^{(5/4)^{t'}} > \left(\frac{\xi}{c_w}\right)^{(5/4)^{t'}} = n.$$

At this point $\mathcal{E}_3(t_1+t')$ must be violated. It follows that $t_2 \leq t' = \mathrm{O}(\log\log n)$ holds w.h.p.

Now fix the configuration $\mathbf{X}(t_1+t_2) = \mathbf{x}$ and assume that $\mathcal{E}_3(t_1+t_2)$ does not hold. This implies that $x_{max}^2/n \geq x_{sec}/3$. As $\psi < x_{max}^2/n + x_{sec}$ this further implies $\psi \leq (x_{max}^2/n) \cdot (1+1/3)$. We now apply Lemma 2.7 with $\delta = c_1 \cdot \sqrt{n \log n}$ which yields that

$$X_{max}(t_1+t_2+1) \geq \frac{x_{max}^2}{\psi}\left(1 - \frac{c_1\sqrt{n\log n}}{x_{max}}\right) \geq \frac{x_{max}^2}{\frac{x_{max}^2}{n}(4/3)}\left(1 - \frac{1}{100}\right) > n \cdot \frac{5}{8}$$

Summarizing, in phase $t_1+t_2+1 = \mathrm{O}(\log\log_{\alpha(0)} n + \log\log n))$ we have $X_{max}(t_1+t_2+1) > n \cdot (5/8)$ w.h.p. Additionally, as in the case of $k \leq \sqrt{n}\log n$, we argue that Claim 2.21 implies that this maximum must be the initial largest opinion. □

Lemma 2.22. *Assume $k > \sqrt{n}/\log n$. Fix $\mathbf{X}(t) = \mathbf{x}$. Assume $x_1(t) \geq (5/8) \cdot n$. Then all agents agree on this opinion within $\mathrm{O}(\log\log n)$ phases, w.h.p.*

Proof. We start by showing that 4 phases following t, at most $\sqrt{n}/\log n$ opinions will have non-zero support. Let i be an opinion which provides $\mathbf{x}_{max}$. From the assumption we have $x_i > (5/8) \cdot n$. Let $L(t)$ denote the set of opinions with support at most $\sqrt{n}\log n$ and fix an opinion $j \in L(t)$. Since $x_i > (5/8) \cdot n$, it follows that $\psi := \mathbb{E}\left[\|\mathbf{Y}(t)\|_1\right] =$

$\sum_{j=1}^{k} x_j^2/n \geq \frac{x_i^2}{n} > (25/64) \cdot n$. We now distinguish two cases, depending on the size of x_j, and show that $X_j(t+1) = \mathrm{O}(\log^2 n)$. First, assume that $x_j < \sqrt{2/\varepsilon^*} \cdot \sqrt{n \log n}$ (remember, the constant $\varepsilon^*$is stated in Definition 2.5). In this case, we apply Lemma 2.8 and use that $\psi = \Theta(n)$ which immediately yields that

$$\Pr\left[X_j(t+1) = \omega(\log n)\right] < 4n^{-2}. \tag{2.22}$$

In the case of $x_j > \sqrt{2/\varepsilon^*} \cdot \sqrt{n \log n}$, we apply Lemma 2.7 together with $\delta = x_j/2\sqrt{n}$ and get

$$\Pr\left[X_j(t+1) > \frac{x_j^2}{\psi} \cdot \frac{3}{2}\right] = \Pr\left[X_j(t+1) > \frac{x_j^2}{\psi}\left(1 + \frac{\delta \cdot \sqrt{n}}{x_j}\right)\right] < 7n^{-2}. \tag{2.23}$$

When using that $\psi > (25/64) \cdot n$ and $x_j < \sqrt{n} \log n$ (remember, $j \in L(t)$) we have $(x_j^2/\psi) \cdot (3/2) \leq 4 \log^2 n$. The inequality in (Eq. (2.23)) then implies that

$$\Pr\left[X_j(t+1) \geq 4 \log^2 n\right] < 7n^{-2}.$$

Hence, in any case, we have that $X_j(t+1) < 4 \log^2 n$ with probability at least $1 - \mathrm{O}(n^{-2})$. By a union bound application, we get that this holds for every $j \in L(t)$ w.h.p. Note that this also implies that $L(t) \subseteq L(t+1)$ w.h.p. Additionally, it follows from Lemma 2.7 that, w.h.p., $X_i(t+1) > x_i(1 - \mathrm{o}(1))$. Now observe that

$$\Pr[Y_j(t+2) = 0 \mid X_j(t+1) < 4 \log^2 n] > \left(1 - \frac{4 \log^2 n}{n}\right)^{4 \log^2 n} > 1 - \frac{\operatorname{polylog} n}{n}.$$

This implies that some fixed opinion $j \in L(t)$ vanishes after decision part $t+2$ w.h.p. As $L(t) \subseteq L(t+1)$ w.h.p. this approach can be repeated, and we deduce that $Y_j(t+3) = 0$ with probability at least $1 - \operatorname{polylog} n/n^2$. By a union bound application it follows w.h.p. for every $j \in L(t)$ that $Y_j(t+3) = 0$, and by a counting argument we have that all but $\sqrt{n}/\log n$ opinions lie in $L(t)$. Therefore, at the start of decision part $t+4$, only $\sqrt{n}/\log n$ opinions remain.

Additionally, consider how the largest opinion i evolves from the decision part of phase t until $t+4$. By Lemma 2.7 we have for $\delta = \sqrt{\log n}$ that, w.h.p.,

$$X_i(t+1) > \frac{x_i^2}{\psi} - \frac{x_i}{\psi}\sqrt{n \log n} \tag{2.24}$$

The inequality $\psi = \sum_{i=1}^{k} x_i^2/n \geq x_{max}$ is true for every configuration $\mathbf{x}$ and additionally we assumed $x_i = x_{max}$. We use this to lower-bound the right-hand side of (Eq. (2.24)) and get that $X_i(t+1) > x_i(1 - \mathrm{o}(1))$ w.h.p. Note, as $x_i \geq (5/8) \cdot n$, this easily implies that $X_i(t+1) > n/2$ and therefore i will remain the opinion with the largest support. When repeating this argument three more times from decision part $t+1$ until $t+4$, we therefore get that $X_i(t+4) > (5/8) \cdot n \cdot (1 - \mathrm{o}(1)) > (24/64) \cdot n$ w.h.p.

Summarizing, we get by a union-bound application that in decision part $t + 4$ (i) less than $\sqrt{n}/\log n$ distinct opinions remain, and (ii) the largest opinion i has at least $(24/64) \cdot n$ support w.h.p. The small amount of remaining opinions allows us the apply results from the analysis in Section 2.4. According to the definition of weak in Section 2.4, all opinions besides i must be weak in phase $t + 4$. Let j with $j \neq i$ be such an opinion. As j is weak, it follows from Lemma 2.16 (which states that weak opinions become super-weak and stay super-weak) that opinion j will be super-weak at some phase $t+4+t'$ with $t' = \mathrm{O}(\log\log n)$ and probability at least $1 - \mathrm{O}(n^{-1.9})$. By a union bound application it follows w.h.p. that *every* such opinion j is super-weak in decision part $t + 4 + t'$. Next we apply Lemma 2.17 which implies that after further 2 phases only opinion i remains and the result follows. □

2.6. Proof of Main Results

We first show that the correctness of the synchronization follows from [1, 74]. There, it is shown that for a polynomial number of phases and any pair of agents u and v, the distance between $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ w.r.t. the circular order modulo $6\tau \log n$ is less than $\tau \log n$, w.h.p. The choice of τ also ensures that every undecided agent can adopt an opinion in the boosting part of a phase, w.h.p.

Proof of Synchronization Properties. In every interaction, every agent is either in the decision or boosting part of a fixed phase. We call an agent u *active in a decision part* as long as $\mathsf{decision}[u] = \textsc{False}$. An agent u is *active in a boosting part* as long as $\mathsf{clock}[u] \leq 5\tau \log n$. Furthermore, we define $P_u(\theta)$ as the number of the phase to which agent u belongs in interaction θ. Intuitively, we aim to show that the leaderless phase clock separates the phases of agents such that no agent is active in the decision part while another agent is simultaneously active in the boosting part. Recall that the leaderless phase clock works as follows. The clock of agent u uses the variable $\mathsf{clock}[u]$, which can take values in $\{0, \dots, 6\tau \log n - 1\}$ for a suitably chosen constant τ. The circular order modulo m, $a \leq_{(m)} b$, is defined as $a \leq_{(m)} b \equiv (a \leq b \text{ XOR } |a - b| > m/2)$, and the distance w.r.t. the circular order modulo m is defined as $\min\{|a - b|,\ m - |a - b|\}$. In every interaction (u, v), the smaller of the two values $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ is increased by one modulo $6\tau \log n$. Smaller refers to the circular order modulo $6\tau \log n$. For the correctness of our protocol, it is sufficient that the following synchronization properties hold for a polynomial number of interactions.

1. For any pair of agents u and v we have $P_u(\theta) = P_v(\theta) \pm 1$.
2. Assume agent u with $P_u(\theta) = t$ is interacting at time θ, and u is active in the decision part of phase t. Then, there exists no agent v that is already active in the boosting part of phase t or is still active in the boosting part of phase $t - 1$.
3. Assume agent u with $P_u(\theta) = t$ is interacting at time θ, and u is active in the boosting part of phase t. Then, no agent v exists that is already active in the decision part of phase $t + 1$ or still active in the decision part of phase t.

4. Let $Z(t)$ be defined as the interval of interactions during which all agents u are together and active in the boosting part of the same phase t, i.e.,

$$Z(t) = \bigcap_u \{\, t \mid P_u(\theta) = t \text{ and } 2\tau \log n \leq \mathsf{clock}[u](t) \leq 5\tau \log n \,\} \,.$$

   Then for each $1 \leq t \leq \operatorname{poly}\log n$ we have $|Z(t)| > n\tau \log n$.

The first condition directly follows from [1, 74]. There, it is shown that for a polynomial number of phases and any pair of agents u and v, the distance between $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ w.r.t. the circular order modulo $6\tau \log n$ is smaller than $\tau \log n$, w.h.p.

To show the second and third conditions, it suffices to show that a) no agent becomes active in the boosting part of phase t while another agent is still active in the decision part of phase t, and b) no agent becomes active in the decision part of a phase $t+1$ while another agent is still active in the boosting part of phase t.

To show a), fix a phase $t \leq \operatorname{poly}\log n$ and let θ be the first interaction in which any agent u has $P_u(\theta) = t$ and $\mathsf{clock}[u] = 2\tau \log n - 1$. As before, observe that $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ differ by less than $\tau \log n$ for any pair of agents u and v at interaction θ w.h.p. Hence, at interaction θ, no other agent v has a clock value $\mathsf{clock}[v] \leq \tau \log n$ w.h.p. It follows that every other agent v has already set $\mathsf{decision}[v] = \textsc{True}$ at interaction θ and thus is not active in the decision part w.h.p. This guarantees (w.h.p.) a clean separation between decision parts and boosting parts. (Technically, it is also necessary that agent v has been at least once the initiator in an interaction pair (v, w). This condition follows from a simple Chernoff bound since every agent was part of at least $\tau \log n$ many interactions, and in each interaction, an involved agent is the initiator with probability $1/2$.)

To show b), fix again a phase $t \leq \operatorname{poly}\log n$ and let θ be the first interaction in which any agent u has $P_u(\theta) = t$ and $\mathsf{clock}[u] = 6\tau \log n - 1$. As before, observe that $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ differ by less than $\tau \log n$ for any pair of agents u and v at interaction θ w.h.p. Hence, at interaction θ, no other agent v has a clock value $\mathsf{clock}[v] \leq 5\tau \log n$ w.h.p., and thus no other agent is active in interaction θ w.h.p. This now guarantees (w.h.p.) the clean separation between boosting and decision parts.

It remains to show the fourth condition. Fix a phase $t \leq \operatorname{poly}\log n$ and let $z_{\min} = \min Z(t)$ and $z_{\max} = \max Z(t)$. At interaction $z_{\min}$, there exists an agent u with $\mathsf{clock}[u] = 2\tau \log n$. Hence no agent can have a clock value larger than or equal to $3\tau \log n$ w.h.p. (since $\mathsf{clock}[u]$ and $\mathsf{clock}[v]$ differ by less than $\tau \log n$ w.h.p., see above). Analogously, at interaction $z_{\max}$, there exists an agent u with $\mathsf{clock}[u] = 5\tau \log n - 1$. As before, no agent can have a clock value smaller than $4\tau \log n$ w.h.p. It takes at least $n\tau \log n$ interactions for all agents to advance their clocks from $3\tau \log n$ to $4\tau \log n$ and hence $|Z(t)| \geq n\tau \log n$ w.h.p.

The fourth condition guarantees that all undecided agents become decided again at the end of a boosting phase. This holds since the interval of interactions during which all agents are in the boosting part of a phase is long enough for a so-called broadcast to succeed, and how agents become decided can be seen as a simple broadcast process. It is folklore that for a sufficiently large constant τ, a broadcast succeeds within $n\tau \log n$ interactions w.h.p. (see, e.g., the notion of "one-way epidemics" in [10]). Note that all agents becoming decided in the boosting part is also crucial to modeling the boosting part by a *Pólya-Eggenberger distribution.* $\square$

Above, we showed that the clocks properly separate the parts of each phase and guarantee w.h.p., long enough phases (of length $O(\log n)$ time) such that Observation 2.3 and Observation 2.4 hold w.h.p. In the rest of the section, we assume that this indeed holds.

Next, we wrap up the unbiased case.

Proof of Theorem 2.1. Proposition 2.13 shows that after $O(\log n)$ phases, only one of the initial significant opinions remains, w.h.p. Then, the theorem follows with the observation that one phase consists of $O(n \log n)$ interactions. $\square$

Finally, we prove the biased case.

Proof of Theorem 2.2. In case $k \leq \sqrt{n}/\log n$, it follows directly from Lemma 2.18 that after $O(\log\log_\alpha n)$ phases all agents agree on the initial majority opinion w.h.p. In case of $k > \sqrt{n}/\log n$, we first rely on Lemma 2.18, where we establish that after $O(\log\log_\alpha n + \log\log n)$ time the initial majority grows to size $(5/8) \cdot n$ w.h.p. Then, we apply Lemma 2.22, which shows that after further $O(\log\log n)$ phases again, all agents agree on this majority opinion w.h.p. The runtime of the result follows as each phase lasts for $O(n \log n)$ interactions. This concludes the proof. $\square$

Chapter 3.

Approximate Plurality Consensus via Undecided State Dynamics

In this part, we analyze the Undecided State Dynamics(USD) in the population protocol model in the high dimensional $k > 2$ regime. Recall the USD is defined as follows. Each agent has either one of k opinions or is undecided, i.e., $Q = \{1, \ldots k, \perp\}$ where $\perp$ stands for undecided. The protocol is given by the transition function

$$(q, q') \to \begin{cases} (\perp, q') \text{ if } q, q' \neq \perp \wedge q \neq q' \\ (q', q') \text{ if } q = \perp, q' \neq \perp \\ (q, q') \text{ otherwise.} \end{cases}$$

Observe that only the responder q changes its state.

It has remained an open problem to analyze the convergence time of the USD for $k > 2$ in this model. [1] In the gossip model with parallel rounds, the USD has been analyzed by Becchetti et al. [20] in the higher dimensional regime $k > 2$. Assuming an initial multiplicative bias and moderate assumption on k, they show the process achieves plurality consensus in $\mathrm{O}\big(\mathrm{md}(\mathbf{x}(0)) \cdot \log n\big)$ rounds, where (assuming x_1 has largest initial support)

$$\mathrm{md}(\mathbf{x}(0)) = \sum_{i \in [k]} \left(\frac{x_i(0)}{x_1(0)}\right)^2 .$$

The so-called monochromatic distance $md(\mathbf{x})$ is an interesting parameter for the convergence time since it captures the whole configuration. We consider a simpler but not necessarily worse parameter: the support of the initial largest opinion.

Results and Methodology As mentioned before, we bound the convergence time in terms of $x_{\max}(0)$, where $x_{\max}(0) = x_1(0)$ is the support of the initially largest opinion. Furthermore, we lose the constraint on the number of initial opinions k and focus solely on the largest opinion. Under a mild assumption on the size of the support ($x_1(0) = \Omega\big(\sqrt{n} \log^2 n\big)$), we show the following result:

[1] Our analysis can also be applied when $k = 2$ and recovers the existing convergence results [9, 35] in this setting.

Theorem 3.1. *Let $c, c_u > 0$ be arbitrary constants and let $\mathbf{x}(0)$ be an initial configuration with $x_1(0) \geq c \cdot \sqrt{n} \log^2(n)$, $u(0) \leq c_u \cdot n$, and $x_1(0) \geq x_i(0)$ for all $i \in [k]$. Then w.h.p. all agents agree on Opinion 1 within*

1. $\mathrm{O}\big(n^2 \log n / x_1(0)\big)$ *interactions if* $\mathbf{x}(0)$ *has an additive bias of at least* $\Omega\big(\sqrt{n} \log n\big)$.
2. $\mathrm{O}\big(n^2 / x_1(0) + n \log n\big)$ *interactions if* $\mathbf{x}(0)$ *has a multiplicative bias of at least* $1 + \epsilon$ *for an arbitrary constant* ϵ.

Without any bias, all agents agree on a significant opinion within $\mathrm{O}\big(n^2 \log n / x_1(0)\big)$ *interactions w.h.p.*

The convergence time heavily depends on the magnitude of the opinion with the largest initial support. Roughly speaking, the convergence time of our result is analogous to that of Becchetti et al. [20] for the gossip model: in that model, plurality consensus is reached within $\mathrm{O}\big(\mathrm{md}(\mathbf{x}(0)) \cdot \log n\big) = \mathrm{O}\big(k \cdot \log n\big)$ *rounds* for $k = \mathrm{O}\big((n/\log n)^{1/3}\big)$. Our result matches their result in this regime by assuming the worst case on $x_{\max}(0) \approx n/k$.

However, we do not require a constant multiplicative bias to reach plurality consensus. We cover the cases with an additive bias of $\Omega\big(\sqrt{n} \log n\big)$ and no initial bias. In the latter, we can only guarantee that one of the significant opinions wins, i.e., one of the initial largest opinions. Note that when the initial configuration does contain a constant multiplicative bias, our analysis gives a faster convergence time than in the additive bias regime. Moreover, our convergence time under a multiplicative bias is faster (when considering its corresponding *parallel time*) than the time given by Becchetti et al. when the support of the initially largest opinion is close to the average opinion support. In this setting, our results for the population protocol model can be viewed as improvements to the analogous results of Becchetti et al. for the gossip model. If there is a large multiplicative bias (larger than $\log n$), Becchetti et al. results give better bounds on the convergence time.

On the other hand, we extend the result by replacing the constraint on k by the support of the initial largest opinion. That is, we originally restricted the initial number of opinions k to fulfill this technical necessity in the analysis. This change further improves the result by dealing with a broader family of configurations, like skewed configurations with many opinions.

Similar to previous analyses in both models [9, 20, 21, 35] our analysis requires carefully defining a sequence of *phases* throughout which the (qualitative and quantitative) behavior of the process varies. The straightforward approach in the analysis of consensus processes is to track the growth of the support of the plurality Opinion 1 via change of the ratio $x_1(t)/x_i(t)$ over time t. Unfortunately, the change in support for a single opinion depends on the entire configuration: the support of all other opinions and the number of undecided agents. Let us fix two opinions i and j with $x_i > x_j$. Then, depending on the number of undecided agents, the support of Opinion j may grow faster than the support of Opinion i and vice versa. Hence, to track the progress of the plurality opinion, one has to examine the number of undecided agents closely. This, in turn, is heavily influenced by the support of all opinions. To cope with this "nonlinearity" we use the potential function

$Z_\alpha(t) = n - 2u(t) - \alpha \cdot x_{\max}(t)$, where we use different values of α for different phases. We analyze the drift of $Z_\alpha(t)$, which allows us to show that the number of undecided agents quickly approaches an "unstable equilibrium" u^*. Whenever the process is close to equilibrium (which changes over time), we can perform a "classical" analysis and show, for example, that bias between two agents doubles in a certain number of interactions. We also deal with a large initial fraction of undecided agents. In this case, we essentially use $-Z_0(t)$ to show the number of undecided agents quickly reaches $n/2$.

As mentioned before, our analysis also handles the case when there is no bias. For this, we proceed in two steps. First, we show that the support difference between two arbitrary but fixed large opinions quickly reaches a value of $\sqrt{n}$ via an anti-concentration bound. From there, we bound the probability that the opinions drift apart further. In our analysis, we rely heavily on existing concentration bounds for the hitting times of one-dimensional random walks with drift, which we can use after establishing the appropriate reductions and potential functions in each phase of the process. The analysis is divided into five parts corresponding to different *phases* of the process. The phases are listed in the following table:

Phase	Section	End Condition	Running Time	Main Lemma
1	Section 3.1	$u \in [(n - x_{\max})/2, n/2]$	$O(n \log n)$	Lemma 3.2
2	Section 3.2	$\forall i : x_{\max} \geq x_i + \Omega(\sqrt{n} \log n)$	$O(n^2 \log n / x_{\max})$	Lemma 3.12
3	Section 3.3	$\forall i : x_{\max} \geq 2x_i$	$O(n^2 \log n / x_{\max})$	Lemma 3.16
4	Section 3.4	$x_{\max} \geq 2n/3$	$O(n^2/x_{\max} + n \log n)$	Lemma 3.21
5	Section 3.5	$x_{\max} = n$	$O(n \log n)$	Lemma 3.25

Note that the process does not have to pass through all five phases. For example, the second phase is not needed if there is a large bias in the initial configuration. Our analysis shows that the identity of the majority opinion does not change after the end of the second phase (or not at all if a large enough additive bias is present from the beginning).

3.1. Evolution of the Undecided (Phase 1)

In this section, we analyze the running time of Phase 1, which ends as soon as we have an appropriate amount of undecided agents (Lemma 3.2). Additionally, we show that the largest opinion roughly maintains its initial support w.h.p. (Lemma 3.3). Furthermore, an additive and multiplicative bias is preserved as long as $\mathbf{x}(0)$ is an initial configuration with bias. At the end of this section, we show an upper bound on the number of undecided agents that hold during the whole running time of the process (Lemma 3.4). This lemma will be used to estimate the running time of the remaining phases.

In the analysis of Lemma 3.2 we use the potential functions

$$Z_1(t) = n - 2u(t) - x_{\max}(t) \quad \text{and} \quad Z_2(t) = 2u(t) - n$$

to track the evolution of the undecided agents. Essentially, the first one $Z_1(t)$ covers the case with an initial small amount of undecided agents while $Z_2(t)$ deals with the relative large case. Observe that Phase 1 ends as soon as $Z_1(t) \leq 0$ and $Z_2(t) \leq 0$, since in this case $u(t) \in [n/2 - x_{\max}(t)/2, n/2]$.

Lemma 3.2. *Let* $T_1 = \inf\{t \geq 0 \mid u(t) \in [n/2 - x_{\max}(t)/2, n/2]\}$. *Then*

$$\Pr\left[T_1 \leq \lceil 7n \ln n \rceil\right] \geq 1 - 2n^{-3}.$$

Proof. To show the lemma, we calculate the expected change in $Z_1(t)$ and $Z_2(t)$ for $Z_1(t) \geq 0$ and $Z_2(t) \geq 0$ and apply a drift theorem from [63]. Note that by definition of both potential functions, it is not possible that both are positive, i.e., at least one of them is trivially fulfilled. We do a case distinction by the initial support of undecided agents. We start with $u(0) \leq (n - x_{\max})/2$ and the expected change of $Z_1(t)$. Considering all possible interactions, there are three cases to be considered. First we consider the case $U(t+1) = u(t) - 1$. In this case, a decided agent interacts with an undecided agent, who adopts the decided agent's opinion Let $M(t) = \{\, i \in [k] \mid x_i(t) = x_{\max}(t) \,\}$ be the set of all opinions with maximum support at time t. For each Opinion i, an undecided initiator interacts with a responder of Opinion i with probability $x_i(t) \cdot u/n^2$. If $i \in M(t)$, then $Z(t)$ increases by 1. Otherwise $Z(t)$ increases by 2.

Next, we consider the case $U(t+1) = u(t)+1$. In this case, a decided initiator interacts with a responder who has a different opinion and becomes undecided. For each Opinion i, this happens with probability $x_i(t) \cdot (n - u(t) - x_i(t))/n^2$. If $i \in M(t)$, then $Z(t)$ decreases by 1 or 2. Otherwise $Z(t)$ decreases by 2.

With the remaining probability, a step is unproductive, and $Z(t)$ does not change. Using these cases, we bound the expected drift of $Z(t)$ as

$$\begin{aligned}
&\mathbb{E}\left[Z(t) - Z(t+1) | \mathbf{X}(t) = \mathbf{x}\right] \\
&\quad \geq - \sum_{i \in M(t)} \frac{x_i \cdot u}{n^2} - 2 \sum_{i \notin M(t)} \frac{x_i \cdot u}{n^2} + \sum_{i \in M(t)} \frac{x_i(n - u - x_i)}{n^2} \\
&\qquad + 2 \sum_{i \notin M(t)} \frac{x_i(n - u - x_i)}{n^2} \\
&\quad \geq \sum_{i \in [k]} \frac{x_i(n - 2u - x_{\max})}{n^2} + \sum_{i \notin M(t)} \frac{x_i(n - 2u - x_{\max})}{n^2} \\
&\quad \geq \frac{(n-u)(n - 2u - x_{\max})}{n^2} \geq \frac{Z(t)}{2n},
\end{aligned}$$

where we used that $x_i \leq x_{\max}$, $Z(t) = n - 2u - x_{\max} \geq 0$, and $u < n/2$ by definition of Phase 1. We now apply the multiplicative drift result of [63] with $r = 3 \ln n$, $s_0 = n - 2u(0) - x_{\max}(0) \leq n$, $s_{min} = 1$, $\delta = 1/(2n)$ and get

$$\begin{aligned}
\Pr\left[T_1 > \lceil 7n \ln n \rceil\right] &\leq \Pr\left[T_1 > \left\lceil \frac{6 \ln n + \ln(n - 2u(0) - x_{\max}(0))}{1/(2n)} \right\rceil\right] \\
&\leq e^{-3 \cdot \ln(n)} = n^{-3}\,.
\end{aligned}$$

Now we deal with the other case where we initially start with $u(0) = c \cdot n > n/2$ and the expected change of $Z_2(t)$. Before we do that, we require an upper bound on the undecided agents during the first $\Omega(n \log n)$ interactions to compute the expected change. It follows from Lemma 3.5 that $d(t) = n - u(t) \geq (1-c)n - \Delta$ for all $t \in [0, \omega(n \log n)]$.

Similar to the first part, we calculate the expected change in $Z_2(t)$ by considering all possible interactions.

$$\begin{aligned}
&\mathbb{E}[Z_2(t+1) - Z_2(t) | \mathbf{X}(t) = \mathbf{x}] \\
&= \sum_{i=1}^{k} \frac{x_i(n-u-x_i)}{n^2} \cdot (Z_2(t)+2) + \sum_{i=1}^{k} \frac{ux_i}{n^2} \cdot (Z_2(t)-2) + \sum_{i=1}^{k} \frac{x_i^2 + x_i u}{n^2} \cdot Z_2(t) - Z_2(t) \\
&= \frac{2}{n^2} \cdot ((n-u)^2 - r^2) - \frac{2}{n^2} \cdot u(n-u) \\
&= -\frac{2}{n^2} \cdot (u(n-u) - (n-u)^2 + r^2) \\
&\leq \frac{2(n-u)(2u-n)}{n^2} = \frac{2(n-u)}{n} \cdot \frac{Z_2(t)}{n}
\end{aligned}$$

where we used that $r^2 = \sum x_i^2 \geq 0$ and $d = \Omega(n)$. We now apply the multiplicative drift result of [63] with $r = 3 \ln n$, $s_0 = n - 2u(0) - x_{\max}(0) \leq n$, $s_{min} = 1$, $\delta = 1/(2n)$ and get

$$\begin{aligned}
\Pr[T_1 > \lceil 7n \ln n \rceil] &\leq \Pr\left[T_1 > \left\lceil \frac{6 \ln n + \ln(n - 2u(0) - x_{\max}(0))}{1/(2n)} \right\rceil\right] \\
&\leq e^{-3 \cdot \ln(n)} = n^{-3} .
\end{aligned}$$

The statement follows by the union bound. □

Given the bound on T_1, we proceed to show that both the support of the most popular opinion and the bias of the initial configuration do not decrease too much until time T_1. Recall that initially, Opinion 1 has the largest support.

Lemma 3.3. *Let $\alpha, \varepsilon > 0$ be arbitrary constants. Then, each of the following statements holds with probability at least $1 - 4n^{-3}$:*

1. *If $x_1(0) - x_i(0) \geq \alpha \cdot \sqrt{n} \log n$, then $X_1(T_1) - X_i(T_1) \geq \alpha/3 \cdot \sqrt{n} \log n$.*
2. *If $x_1(0) \geq (1+\varepsilon) \cdot x_i(0)$, then $X_1(T_1) \geq 1 + \varepsilon/2 \cdot X_i(T_1)$.*
3. *For the largest opinion we have $X_1(T_1) \geq x_1(0)/3$.*

Proof. We start with the proof of the first statement. Fix an Opinion $i \neq 1$. We consider

$$\Psi_t = \frac{x_1(t) - x_i(t)}{n - u(t)}$$

and show via a version of the Hoeffding bound (see Lemma A.7) that this quantity does not decrease significantly throughout the first phase. Recall that T_1 is defined as the first time t where $u(t) \in [n/2 - x_{\max}(t)/2, n/2]$ and that by definition $x_{\max}(0) = x_1(0)$. Let $\hat{T} = \inf\{\, t \geq 0 \mid u(t) \geq c_u \cdot n + \Delta \,\}$ be a stopping time and let $(\hat{\mathbf{X}}(t))_{t\in\mathbb{N}}$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. From Lemma 3.5 it follows $\hat{T} \geq n^3$ with probability at least $1 - n^{-3}$. Also, from Lemma 3.2 it follows that $T_1 \leq 7 \cdot n \ln n$ with probability $1 - n^{-3}$. Thus, w.h.p. $(\mathbf{X}(t))_{t\in\mathbb{N}}$ and $(\hat{\mathbf{X}}(t))_{t\in\mathbb{N}}$ behave the same during the first phase. For readability, we slightly abuse the notation and continue the proof with Z_t and $\mathbf{X}(t)$ instead of $\hat{Z}_t$ and $\hat{\mathbf{X}}(t)$. In the following we define $Z_{t+1} = \Psi_{t+1} - \Psi_t$ and $\mu_{t+1} = \mathbb{E}\left[Z_{t+1}|\mathbf{X}_{<t+1}\right]$. Our goal is to apply Lemma A.7 to $Z_1, Z_2, \ldots$. First we calculate $\mu_{t+1} = \mathbb{E}\left[Z_{t+1}|\mathbf{X}_{<t+1}\right]$. Considering all possible interactions yields

$$
\begin{aligned}
&\mathbb{E}\left[Z_{t+1} \mid \mathbf{X}_{<t+1}\right] \\
&\quad= \frac{x_1(n-u-x_1)}{n^2}\cdot\left(\frac{x_1-x_i-1}{n-u-1}\right)+\frac{x_i(n-u-x_i)}{n^2}\cdot\left(\frac{x_1-x_i+1}{n-u-1}\right) \\
&\qquad+\sum_{j\neq 1,i}\frac{x_j(n-u-x_j)}{n^2}\cdot\left(\frac{x_1-x_i}{n-u-1}\right)+\frac{ux_1}{n^2}\cdot\left(\frac{x_1-x_i+1}{n-u+1}\right)+\frac{ux_i}{n^2}\cdot\left(\frac{x_1-x_i-1}{n-u+1}\right) \\
&\qquad+\sum_{i\neq 1,j}\frac{ux_j}{n^2}\cdot\left(\frac{x_1-x_i}{n-u+1}\right)+\frac{r^2+nu}{n^2}\cdot\Psi_t-\Psi_t \\
&\quad= \frac{x_1(n-u-x_1)}{n^2}\cdot\left(\frac{\Psi_t-1}{n-u-1}\right)+\frac{x_i(n-u-x_i)}{n^2}\cdot\left(\frac{\Psi_t+1}{n-u-1}\right) \\
&\qquad+\sum_{j\neq 1,i}\frac{x_j(n-u-x_j)}{n^2}\cdot\left(\frac{\Psi_t}{n-u-1}\right)+\frac{ux_1}{n^2}\cdot\left(\frac{-\Psi_t+1}{n-u+1}\right)+\frac{ux_i}{n^2}\cdot\left(\frac{-\Psi_t-1}{n-u+1}\right) \\
&\qquad+\sum_{i\neq 1,j}\frac{ux_j}{n^2}\cdot\left(\frac{-\Psi_t}{n-u+1}\right) \\
&\quad= \frac{1}{n^2(n-u-1)}\cdot\left(((n-u)^2-r^2)\Psi_t-x_1(n-u-x_1)+x_i(n-u-x_i)\right) \\
&\qquad-\frac{u}{n^2(n-u+1)}\left((n-u)\Psi_t-(x_1-x_i)\right) \\
&\quad= \frac{x_1-x_i}{n^2(n-u)(n-u-1)}\cdot\left((x_1+x_i)(n-u)-r^2\right) \\
&\quad= \frac{x_1-x_i}{n^2(n-u)(n-u-1)}\cdot\left(\sum_{j=1}^{k}(x_j(x_1-x_j))+x_i\cdot(n-u)\right)\geq 0
\end{aligned}
$$

Now we show that $|Z_{t+1}| \leq a$ with $a = 4/(n-2)$ for all $t < \tau$. To do so, we consider every possible outcome of

$$\frac{x_1(t+1)-x_i(t+1)}{n-u(t+1)} - \frac{x_1(t)-x_i(t)}{n-u(t)},$$

i.e., either the number of undecided agents u is increased and decreased by one, or x_i or x_1 are increased by one.

It is easy to see that

$$|Z_{t+1}| \leq \frac{x_1 - x_i}{(n-u)(n-u-1)} \leq \frac{n-u+(x_1-x_i)}{(n-u)(n-u-1)} \leq \frac{(n-u)+(n-u)}{(n-u)(n-u-1)} = \frac{2}{n-u-1}$$

By definition of $\hat{T}$ we know $u(t) < c_u \cdot n + \Delta$ and thus $|Z_{t+1}| < 4/((1-c_u)n)$. Observe that for every t we have

$$\sum_{i=1}^{t} Z_i = \sum_{i=1}^{t} (\Psi_i - \Psi_{i-1}) = \Psi_t - \Psi_0.$$

Finally, the application of Lemma A.7 with $|Z_{t+1}| \leq 4/((1-c_u)n)$, $\tau = 7n\ln n$, $\lambda = \sum_{i=1}^{\tau} \mu_i + c' \log n/\sqrt{n}$ and $c' = 15 \cdot \sqrt{3} \cdot \ln 2$ yields

$$\Pr\left[\Psi_\tau - \Psi_0 < -c' \cdot \frac{\log n}{\sqrt{n}}\right] = \Pr\left[\sum_{i=1}^{\tau} Z_i - \sum_{i=1}^{\tau} \mu_i < -\lambda\right] \leq n^{-3}.$$

Recall that from Lemma 3.2 it follows that $T_1 \leq 7 \cdot n \ln n$ with probability $1 - n^{-3}$. Thus $\Psi_\tau > \Psi_0 - c' \cdot \log n/\sqrt{n}$ with probability at least $1 - n^{-3}$. Again, since $u(t) < n/2$, we have

$$\begin{aligned}
& \frac{x_1(T_1) - x_i(T_1)}{n - u(T_1)} = \Psi_{T_1} > \Psi_0 - c' \cdot \log n/\sqrt{n} = \frac{x_1(0) - x_i(0)}{n - u(0)} - c' \cdot \log n/\sqrt{n} \\
\Leftrightarrow \quad & x_1(T_1) - x_i(T_1) > (x_1(0) - x_i(0))\frac{n - u(T_1)}{n - u(0)} - (n-u)(T_1)c' \cdot \log n/\sqrt{n} \\
\Longrightarrow \quad & x_1(T_1) - x_i(T_1) > \frac{x_1(0) - x_i(0)}{2} - c'\sqrt{n}\log n
\end{aligned}$$

Due to the assumption $x_1(0) - x_i(0) \geq c\sqrt{n}\log n$ with $c = 6c'$ we have $x_1(T_1) - x_i(T_1) > (x_1(0) - x_i(0))/3$. By application of the union bound over all opinions $i \neq 1$, the first statement holds with probability at least $1 - n^{-2}$.

For the second statement, we consider $\Phi(t) = x_1(t)/x_i(t)$ for an opinion $i \neq 1$ and show via a concentration result that this quantity does not decrease significantly throughout the first phase. To do that we define $W_{t+1} = \Phi(t+1) - \Phi(t)$ and $\mu_{t+1} = \mathbb{E}[W_{t+1} \mid \mathcal{F}_t]$.

First we calculate μ_{t+1}.

$$
\begin{aligned}
&\mathbb{E}\left[\Phi(t+1)-\Phi(t)\mid\mathcal{F}_t\right]\\
&\quad=\frac{x_1(n-u-x_1)}{n^2}\left(\frac{x_1-1}{x_i}\right)+\frac{x_i(n-u-x_i)}{n^2}\left(\frac{x_1}{x_i-1}\right)+\sum_{j\neq 1,i}\frac{x_j(n-u-x_j)}{n^2}\Phi(t)\\
&\quad+\frac{ux_1}{n^2}\left(\frac{x_1+1}{x_i}\right)+\frac{ux_i}{n^2}\left(\frac{x_1}{x_i+1}\right)+\frac{u(n-u-x_1-x_i)}{n^2}\Phi(t)\\
&\quad+\frac{(n-u)^2-\sum x_i^2+un}{n^2}\Phi(t)-\Phi(t)\\
&\quad=-\frac{(n-u-x_1)\Phi(t)}{n^2}+\frac{x_i(n-u-x_i)\Phi(t)}{n^2(x_i-1)}+\frac{u\Phi(t)}{n^2}-\frac{ux_i\Phi(t)}{n^2(x_i+1)}\\
&\quad=\frac{\Phi(t)}{n^2}\left(x_1-x_i+\frac{n-u-x_i}{x_i-1}+\frac{u}{x_i+1}\right)\geq 0
\end{aligned}
$$

Next, by considering every possible outcome it follows that $|W_{t+1}|\leq 4/n$ for all $t<\tau$. Observe that $\sum W_t=\Phi(t)-\Phi(0)$. Finally, the application of Lemma A.7 with $|W_{t+1}|\leq 4/n$, $\tau=cn\log n$, $\lambda=\sum\mu_i+c\log n/\sqrt{n}$ yields

$$\Pr\left[\Phi(\tau)<\Phi(0)-c\log n/\sqrt{n}\right]=\Pr\left[W_\tau-\mu<-\lambda\right]\leq e^{-\frac{2\lambda^2}{\tau\cdot 16/n^2}}\leq e^{-\frac{-c^2\log n}{8}}\leq n^{-\frac{c^2}{8}}$$

Thus,

$$\Phi(\tau)=\frac{x_1(\tau)}{x_i(\tau)}\geq\frac{x_1(0)}{x_i(0)}-\frac{c\log n}{\sqrt{n}}\geq(1+\varepsilon)-\frac{c\log n}{\sqrt{n}}\geq 1+\varepsilon/2$$

For the third statement, we consider 1-productive interactions. The proof follows the analysis of the classical Gambler's ruin problem by tracking the evolution of $x_1(t)$ over time. An interaction is 1-productive with probability at least

$$\frac{x_1(n-x_1)}{n^2}$$

Thus, an application of Chernoff bounds provides at most τ productive interactions in T interactions w.h.p.

For $1\leq i\leq\tau$ we define t_i as the i-th 1-productive interaction in $[0,T]$. Then for an arbitrary $i\in[1,\tau]$ we have

Conditioned on a 1-productive interaction, the probability that the support of the largest opinion increases by one is at least

$$\Pr\left[x_1(t_i+1)=x_1+1|\mathbf{X}(t_i)=\mathbf{x}\right]=\frac{1}{2}+\frac{2u-n+x_1}{2(n-x_1)}\geq\frac{1}{2}+\frac{x_1}{2n}$$

where we use $u\geq n/2$.

Observe that starting at time t_0 with $\Delta=x_1(t_0)$ as long as $x_1(t_i)-\geq\Delta/2$ for the first $i\leq\tau$ many $(1,i)$-productive interactions in $[0,T]$ the evolution of $x_1(t_i)$ can be viewed as a biased random walk on the line starting at Δ where a "right step" happens with probability $p = 1/2 + \Delta/4n$ and "left step" with probability $1-p$, otherwise. The correctness follows from a standard coupling argument between two biased coins. Formally let $T_{min} = \inf\{t' \geq t_0 \mid x_1(t') = (1/2)\cdot x_1(t_0)\}$ We bound $\Pr[T_{min} > \tau]$. It follows from Lemma A.18 the probability of ever having an excess of $(5/12)\cdot\Delta$ "left steps" to "right steps" is at most

$$\left(\frac{1-p}{p}\right)^{(5/12)\cdot\Delta} \leq n^{-5}$$

In the other case,the third statement is a direct consequence of the first statement. To see this, consider n agents with k opinions and let $\mathbf{x}$ be a vector of size $k+1$ that denotes the configuration (x_{k+1} is the number of undecided agents). Let $\mathbf{y}$ be a vector of size $k+2$ where $y_{k+2} = x_{k+1}$, $y_{k+1} = 0$ and otherwise $y_i = x_i$. The process does not depend on whether we use $\mathbf{x}$ or $\mathbf{y}$ to describe the configuration, and thus $y_{k+1}(t) = 0$ for all $t \geq 0$. Therefore, any statement that holds for $\mathbf{x}$ also holds for $\mathbf{y}$ (using $k+1$ instead of k). The second statement assumes that $y_1(0) = x_1(0) = \Omega\big(\sqrt{n}\log^2 n\big)$. Hence, we can apply the first statement to $\mathbf{y}$ and get $x_1(T_1) = y_1(T_1) - 0 = y_1(T_1) - y_{k+1}(T_1) \geq (y_1(0) - y_{k+1}(0))/3 = y_1(0)/3 = x_1(0)/3$. $\square$

Next, we prove the upper bound on the number of undecided agents. The lemma shows that the number of undecided agents stays close to a threshold value $u^* = n\cdot(k-1)/(2k-1) \approx n/2$. Intuitively, this threshold u^* can be regarded as an (unstable) equilibrium for the number of undecided agents: in configurations with more than u^* undecided agents, it is more likely that an undecided agent becomes decided than vice versa, whereas in configurations with less than u^* undecided agents, it is more likely that a decided agent becomes undecided than vice versa.

Lemma 3.4.

$$\Pr\left[\forall t \in [T_1, T_1 + n^3] \colon u(t) \leq \frac{n}{2} + 2\Delta_u\right] > 1 - n^{-3}.$$

Proof. We model the number of undecided agents over time t as a random walk $U(t)$ with state space $\{0,\ldots,n-1\}$ and denote the corresponding non-lazy random walk $Z(r)$. The transition probabilities of $Z(r)$ are denoted as

$$\Pr[Z(r+1) = z(r)+1 | Z(r) = z(r)] = \tilde{p}_+(r)$$
$$\Pr[Z(r+1) = z(r)-1 | Z(r) = z(r)] = 1 - \tilde{p}_+(r).$$

Unfortunately the transition probabilities of $Z(r)$ depend on the configuration at time r and thus the random walk is not time-homogeneous. However, we can bound the transition probabilities as follows. From calculations it follows that $\tilde{p}_+(r) \leq 1/2 - \Delta_u/n$

if $Z(r) \geq n/2 + \Delta_u$. That is, conditioned on a productive interaction at time t, it holds

$$\begin{aligned}\Pr\left[U(t+1) = u(t+1)+1 \mid X(t) = x\right] &= \frac{1}{2} - \frac{u(n-u)+r^2-(n-u)^2}{2(u(n-u)+(n-u)^2-r^2)} \\ &\leq \frac{1}{2} - \frac{u(n-u)-(n-u)^2}{2(u(n-u)+(n-u)^2)} \\ &= \frac{1}{2} - \frac{2u-n}{2n} \leq \frac{1}{2} - \frac{\Delta_u}{n}\end{aligned}$$

where in the last inequality, we use $u \geq n/2 + \Delta_u$.

For $Z(r)$ we know that we have a drift "in the right direction" between $n/2 + \Delta_u$ and $n/2 + 2\Delta_u$. Therefore, it suffices to bound the probability that a random walk traverses from $n/2+\Delta_u$ to $n/2+2\Delta_u$. To do so, we define a random walk $W(r)$ on the non-negative integers with a reflecting barrier at 0 and otherwise transition probabilities

$$\begin{aligned}\Pr\left[W(r+1) = w(r)+1 | W(r) = w(r)\right] = p = 1/2 - \Delta_u/n \\ \Pr\left[W(r+1) = w(r)-1 | W(r) = w(r)\right] = q = 1/2 + \Delta_u/n.\end{aligned}$$

To show the statement we now define $Z'(r) = Z(r) - (n/2+\Delta_u)$ and couple $Z'(r)$ with $W(r)$. From the definitions of the random walks $Z'(r)$ and $W(r)$ the following two statements follow. If $Z'(r) < 0$, then $Z'(r+1) \leq W(r+1)$ since $W(r)$ has a reflecting barrier at 0. Otherwise, $Z'(r+1) \leq W(r+1)$ follows from the coupling between $Z'(r)$ and $W(r)$ since $\Pr\left[Z'(r+1) = z(r)+1 | Z'(r) = z(r)\right] \leq \Pr\left[W(r+1) = w(r)+1 | W(r) = w(r)\right]$ for any $z(r)$ and $w(r)$. It therefore follows that, deterministically, $Z(r) \leq n/2 + \Delta_u + W(r)$.

We now proceed to prove that $W(r) \leq \Delta_u$ w.h.p. It is straightforward to verify that the stationary distribution W of $W(r)$ is given by $\Pr\left[W = n\right] = (p/q)^n \cdot (1-p/q)$. Therefore $\Pr\left[W \geq n\right] = (p/q)^n$. When we start with $W(0) = 0$, it holds from a union bound over n^3 steps that $\Pr\left[\exists t \in [n^3] \colon W(t) \geq m\right] \leq n^3 \cdot \Pr\left[W \geq m\right] \leq n^3 \cdot (p/q)^m$ for some value $m > 0$. It yields

$$\Pr\left[\exists t \in [n^3] : W(t) \geq \Delta_u\right] \leq n^3 \cdot \left(\frac{1/2 - \Delta_u/n}{1/2 + \Delta_u/n}\right)^{\Delta_u} \leq n^3 \cdot e^{-\frac{4\Delta_u^2}{n+2\Delta_u}} < n^{-3}.$$

□

Lemma 3.5. *Let* $\mathbf{x}(0)$ *be an initial configuration with* $u(0) \leq c_u \cdot n$ *for an abitrary constant* $c_u \in (1/2, 1)$*. Then*

$$\Pr\left[\text{ for all } t \in [n^3] : u(t) \leq c_u \cdot n + \Delta\right] \geq 1 - n^{-3}$$

Proof. The idea is to track the evolution of undecided agents over time via a biased random walk. Whenever the support of the undecided agents exceeds a certain upper threshold, it tends to decrease again. That is, whenever $u(t) \geq cn$, i.e., crosses the threshold from below, it does not reach $cn+\Delta$ in the next polynomial many interactions.

Let t_i denote the time whenever $u(t_i) \geq cn$ crosses for the i-th time, i.e., $u(t_i - 1) < cn$. Conditioned on a productive interaction at time $t \geq t_i$, it holds

$$\begin{aligned}
\Pr\left[U(t+1) = u(t+1) + 1 \mid X(t) = x\right] &= \frac{1}{2} - \frac{u(n-u) + r^2 - (n-u)^2}{2(u(n-u) + (n-u)^2 - r^2)} \\
&\leq \frac{1}{2} - \frac{u(n-u) - (n-u)^2}{2(u(n-u) + (n-u)^2)} \\
&= \frac{1}{2} - \frac{2u-n}{2n} = 1 - \frac{u}{n} < 1/2
\end{aligned}$$

where in the last inequality we use $u \geq cn$. Similar to Lemma 3.4, the statement follows w.h.p. by a known Gambler's ruin result and the union bound. □

3.2. Generation of an Additive Bias (Phase 2)

Recall that T_1 is defined as the end of Phase 1. This section considers configurations at time T_1 without any additive bias. These configurations will have several significant opinions. We define T_2 as the first time $t \geq T_1$ where $\mathbf{x}(t)$ has only one opinion left, which is significant.

Note that $x_{\max}(t)) \geq x_{\max}(0)/2 = \Omega\big(\sqrt{n} \cdot \log^2(n)\big)$ for each interaction t in this phase. This follows from Lemma 3.4 together with the pigeonhole principle. In Lemma 3.12 we show that w.h.p. the running time of this phase is $\mathrm{O}\big(n^2 \cdot \log n / x_{\max}(T_1)\big)$. To show that result, we first need a lower bound (as opposed to the upper bound of Lemma 3.4) on the number of undecided agents. Again, this bound holds until the end of the process.

Lemma 3.6. $\Pr\Big[$*for all* $t \in [T_1, n^3]\colon u(t) \geq n/2 - x_{\max}(t)/2 - 8\sqrt{n \cdot \ln n}\Big] \geq 1 - n^{-5}$.

Proof. Recall that we defined $Z(t) = n - 2u(t) - x_{\max}(t)$ and that $Z(T_1) \leq 0$. In the following, we show that w.h.p. $Z(t) \leq c\sqrt{n \cdot \log n}$ for all $T_1 \leq t \leq n^3$.

We follow the proof idea of Theorem 6 in [64]. We define a new set of random variables with $Y(t) = \exp(\eta \cdot Z(t))$ for $t \geq T_1$ and $\eta = \sqrt{\ln n / n}$ and let $z_0 = 4\eta \cdot n$.

Fix an arbitrary $i \geq 0$. We first give a bound for $\mathbb{E}\left[Y(i+1) - Y(i) \mid Z(i) = z\right]$. Note that $Z(i+1) - Z(i) \in [-2, 2]$. We get

$$\begin{aligned}
\mathbb{E}\left[Y(i+1) - Y(i) \mid Z(i) = z\right] &= \mathbb{E}\left[e^{\eta \cdot Z(i+1)} - e^{\eta \cdot Z(t)} \mid Z(i) = z\right] \\
&= e^{\eta \cdot z} \cdot \mathbb{E}\left[e^{\eta \cdot (Z(i+1) - z)} - 1 \mid Z(i) = z\right] \\
&= e^{\eta \cdot z} \cdot \sum_{j \in [-2,2]} (e^{\eta \cdot j} - 1) \cdot \Pr\left[Z(i+1) - z = j \mid Z(i) = z\right]
\end{aligned}$$

We derive the following bound for $\exp(\eta \cdot j) - 1$. Since $\exp(x) \leq 1 + x + x^2$ for $x \leq 1$ and $\eta \to 0$ for large n, we have $\exp(2\eta) \leq 1 + 2\eta + (2\eta)^2 = 1 + 2\eta + \eta \cdot z_0/n$. For

$j \in [-2, 2]$, we thus have $\exp(\eta j) - 1 \leq \eta j + \eta \cdot z_0/n$. In Lemma 3.2 we calculated $\mathbb{E}[Z(i+1) - Z(i) \mid Z(i) = z] \leq -\frac{z}{n}$.

Thus, for all $z \geq z_0$ we have

$$\begin{aligned}
&\mathbb{E}[Y(i+1) - Y(i) \mid Z(i) = z] \\
&\quad \leq e^{\eta \cdot z} \cdot \sum_{j \in [-2,2]} (\eta \cdot j + \eta \cdot z_0/n) \cdot \Pr[Z(i+1) - z = j | Z(i) = z] \\
&\quad = e^{\eta \cdot z} \cdot \eta \cdot (\mathbb{E}[Z(i+1) - Z(i) | Z(i) = z] + z_0/n) \leq 0.
\end{aligned}$$

In total, we get

$$\mathbb{E}[Y(t)] = \mathbb{E}[Y(0)] + \sum_{i=0}^{t-1} \mathbb{E}[Y(i+1) - Y(i)] \leq 1.$$

We can apply Markov's inequality since $\forall t \geq 0 : Y(t) \geq 0$. Thus,

$$\Pr[Z(t) \geq 2z_0] = \Pr[Y(t) \geq \exp(2\eta z_0)] \leq \frac{\mathbb{E}[Y(t)]}{n^8} \leq n^{-8}.$$

Finally, we apply the union bound over $n^3 - T_1 \leq n^3$ interactions. □

In the following lemma, we show that the support of the largest opinion does not shrink by more than a factor of two during Phase 2.

Lemma 3.7. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \cdot \log n / x_{\max}(T_1)$. Then*

$$\Pr[\textit{for all } t \in [T_1, T_1 + T] \colon x_{\max}(t) \geq x_{\max}(T_1)/2] \geq 1 - n^{-5}.$$

Proof. We show that $x_{\max}(t) \geq x_{\max}(T_1)/2$ for all $t \in [T_1, T_1 + cn^2 \log n / x_{\max}(T_1)]$. First, we bound the number of productive interactions w.r.t. $x_{\max}(t)$ within $cn^2 \log n / x_{\max}(T_1)$ interactions, and then we bound its effect on the support of the largest opinion. Let

$$\hat{T} = \inf \{ \, t \geq T_1 \;\mid\; u(t) \notin [(n - x_{\max}(t'))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u] \, \}$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. From Lemma 3.4 and Lemma 3.6 it follows $\hat{T} - t = \Omega\big(n^2/x_{\max}(t) \cdot \log n\big)$ w.h.p. Thus, $(\mathbf{X}(t))_t$ and $(\hat{\mathbf{X}}(t))_t$ behave the same between time t and $t + \mathrm{O}\big(n^2/x_{\max}(T_1) \cdot \log n\big)$.

As long as $x_{\max}(t) \leq 2 \cdot x_{\max}(T_1)$ an interaction is productive w.r.t. $x_{\max}(t)$ with probability

$$\frac{u \cdot x_{\max} + x_{max} \cdot (n - u - x_{\max})}{n^2} = \frac{x_{\max} \cdot (n - x_{\max})}{n^2} \leq 2 \cdot \frac{x(T_1)}{n}$$

It follows from an application of Chernoff bounds that within a sequence of $c \cdot n^2 \cdot \log n / x_{\max}(T_1)$ interactions, the number of such productive interactions is at most $4 \cdot c \cdot$

$n \log n$ with probability at least $1 - n^{-10}$. Now consider $\tau = 4 \cdot c \cdot n \log n$ such productive interactions and let Z_t denote the change w.r.t. $x_{\max}(t)$, i.e., the support of the largest opinion increase or decrease by one, respectively. That is, assuming the next interaction is such a productive interaction for $x(t)$, we have

$$\Pr[Z_t = 1] = \frac{u \cdot x_{\max}}{u \cdot x_{\max} + x_{\max} \cdot (n - u - x_{\max})} = \frac{u}{(n - x_{\max})}$$
$$\Pr[Z_t = -1] = 1 - \Pr[Z_t = 1]$$

Therefore

$$\mathbb{E}[Z_t] = \frac{u - (n - u - x_{\max})}{n - x_{\max}} = \frac{2 \cdot u + x_{\max} - n}{n - x_{\max}} \geq -48 \cdot \frac{\sqrt{n \ln n}}{n}$$

Let Z be the sum of Z_t for all $t \in [1, \tau]$. Then it follows from Hoeffding bound with $\lambda = x_{\max}(T_1)/2 - 200 \cdot \sqrt{n} \ln^{3/2} n$

$$\Pr\left[Z < -\frac{1}{2} \cdot x_{\max}(T_1)\right] \leq \Pr[Z < \mathbb{E}[Z] - \lambda] \leq e^{-\frac{2\lambda^2}{4\tau}} \leq n^{-10}$$

Note that if (ever) $x_{\max}(t') > 2 \cdot x_{\max}(T_1)$ for some $t' \in [T_1, T_1 + T]$ the statement hold as well by the union bound and the previous part. Thus, starting with $x_{\max}(T_1)$ throughout the next $c \cdot n^2 / x_{\max}(T_1) \cdot \log n$ interactions $x_{\max}(t) \geq x_{\max}(T_1)/2$ with probability at least $1 - n^{-5}$. $\square$

In the following, we define two observations that are frequently used in our proofs. That is, we need some statements about the process , e.g., bounds on the support of undecided agents, that last throughout the phases and hold with high probability. Thus, we formally define another process with such properties and show that both processes behave the same for a sufficiently larger number of interactions.

Observation 3.8. *For a given time $\hat{t}$ and configuration $\mathbf{x}(\hat{t})$ let*

$$\hat{T} = \inf\{\, t \geq \hat{t} \,\} \ \mid u(t) \notin [(n - x_{\max}(t'))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u]$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$.

Observation 3.9. *For a given time T_i and configuration $\mathbf{x}(T_i)$ let*

$$\hat{T} = \inf\{\, t \geq T_i \ \mid \ u(t) \notin [(n - x_{\max}(t))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u] \ \text{or} \ x_{\max}(t) < x_{\max}(\hat{t})/2 \,\}$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. Then $(\hat{X}(t))_t$ and $(X(t))_t$ behave the same for every $t \in [T_i, T_i + \mathrm{O}(n^2 / x_{\max}(T_i) \cdot \log n)]$ w.h.p.

Proof. From Lemma 3.6 it follows that $u(t) \geq (n - x_{\max}(t))/2 - 8 \cdot \sqrt{n \ln n}$ for all $t \in [T_1, n^3]$, w.h.p. From Lemma 3.4 it follows that $u(t) \leq n/2 + \Delta_u$ for all $t \in [T_1, n^3]$, w.h.p. Finally, Lemma 3.7 gives us that $x_{\max}(t) \geq x_{\max}(T_1)/3$ for all $t \in [T_1, T_1 + cn^2 \log n / x_{\max}(T_1)]$, w.h.p. Thus, $\hat{T} - T_1 = \Omega\big(n^2 \cdot \log n / x_{\max}(T_1)\big)$ w.h.p. and we can assume that $(\mathbf{X})_{t \geq T_1}$ and $(\hat{\mathbf{X}})_{t \geq T_1}$ are identical for $t \in [T_1, T_1 + \mathrm{O}\big(n^2 \cdot \log n / x_{\max}(T_1)\big)]$. $\square$

In Lemma 3.10, we first show that "small opinions" remain small (they only double their support). With small opinion, we mean opinions having support which have support at most $20\sqrt{n \log n}$ and are thus at least a polylogarithmic factor smaller compared to $x_{\max}(t)$. Then, in the second part, we show that insignificant opinions remain insignificant. Recall that an Opinion i is insignificant if $x_{\max}(t) - x_i(t) = \Omega\big(\sqrt{n} \log n\big)$.

Lemma 3.10. *Let $c, c' > 0$ be arbitrary constants and define $T = c \cdot n^2 \cdot \log n / x_{\max}(T_1)$. Assume for Opinion j there exists a time $t_0 \in [T_1, T_1 + T]$ with*

1. *$x_j(t_0) \leq 20\sqrt{n \log n}$. Then*

$$\Pr\left[\textit{for all } t \in [t_0, T_1 + T] \colon x_j(t) \leq 40\sqrt{n \log n}\right] \geq 1 - 2n^{-3}.$$

2. *$x_{\max}(t_0) - x_j(t_0) \geq c' \cdot \sqrt{n} \log n$. Then*

$$\Pr\left[\textit{for all } t \in [t_0, T_1 + T] \colon x_{\max}(t) - x_i(t) \geq c'/2 \cdot \sqrt{n} \log n\right] \leq 1 - 2n^{-3}.$$

Proof. We use Observation 3.9 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$. Now, we start with the first statement. First, we bound the number of i-productive interactions in the interval $[t_0, T_1 + T]$. Recall that only i-productive interactions change the support of Opinion i. As long as $x_i(t) \leq 40\sqrt{n \log n}$ for $t \in [t_0, T_1 + T]$ an interaction is i-productive with probability

$$\frac{u \cdot x_i + x_i \cdot (n - u - x_i)}{n^2} = \frac{x_i \cdot (n - x_i)}{n^2} \leq \frac{40\sqrt{n \log n}}{n}$$

It follows from an application of Chernoff bounds that the number of such productive interactions is at most $n / \log^{1/4} n$ w.h.p.

Now consider $\tau = n / \log^{1/4} n$ i-productive interactions and let Z_ℓ denote the respective change of the ℓth i-productive interaction. That is,

$$\Pr[Z_\ell = 1] = \frac{u \cdot x_i}{u \cdot x_i + x_i \cdot (n - u - x_i)} = \frac{u}{(n - x_i)}$$
$$\Pr[Z_\ell = -1] = 1 - \Pr[Z_\ell = 1]$$

Therefore

$$\mathbb{E}[Z_\ell] = \frac{u - (n - u - x_i)}{n - x_i} = \frac{2 \cdot u + x_i - n}{n - x_i} \leq \frac{2\Delta_u + 40\sqrt{n \log n}}{n - 40\sqrt{n \log n}} \leq \frac{c_u\sqrt{n \log n}}{n - 40\sqrt{n \log n}}$$

Let Z be the sum of Z_ℓ for all $\ell \in [1, \tau]$. Then it follows from Hoeffding bound with $\lambda = \sqrt{n \log n} - \mathbb{E}[Z]$

$$\Pr\left[Z > \sqrt{n \log n}\right] \leq \Pr\left[Z > \mathbb{E}[Z] + \lambda\right] \leq e^{-\frac{2\lambda^2}{4\tau}} \leq n^{-6}$$

Thus, starting with $x_i(t_0) \leq 20\sqrt{n \log n}$ it holds that $x_i(t) \leq 40\sqrt{n \log n}$ for all $t \in [t_0, T_1 + \tau]$ w.h.p.

We now show the second statement. Our proof follows the analysis of the classical Gambler's ruin problem on the quantity $x_{max}(t) - x_j(t)$. That is, for T interactions we show that $x_{\max}(t) - x_j(t) \geq (x_{\max}(t_0) - x_j(t_0))/2$ for all $t \in [t_0, T_1 + T]$.

Let $B(t) = \{\, i \mid x_i(t) \geq x_{max}(t) \,\}$ denotes the set of all opinions i with maximum support at time t. Note that $|B(t)| \geq 1$ for all t. Consider an arbitrary time $t \in [t_0, T_1 + T]$. By the definition of USD, it follows that

$$\Pr\left[X'_{max} - X_j = x_{max} - x_j + 1 | \mathbf{X}(t) = \mathbf{x}(t)\right] = \begin{cases} \frac{u \cdot \sum_{i \in B} x_i + x_j \cdot (n-u-x_j)}{n^2} & , |B| > 1 \\ \frac{u \cdot x_{max} + x_j \cdot (n-u-x_j)}{n^2} & , |B| = 1 \end{cases} \tag{3.1}$$

$$\Pr\left[X'_{max} - X_j = x_{max} - x_j - 1 \mid \mathbf{X}(t) = \mathbf{x}(t)\right] = \begin{cases} \frac{u \cdot x_j}{n^2} & , |B| > 1 \\ \frac{x_{max} \cdot (n-u-x_{max}) + u \cdot x_j}{n^2} & , |B| = 1 \end{cases}$$

Let $p_1(t)$ and $p_2(t)$ denote the first and second probability from Eq. (3.1), respectively. Now assuming for $\mathbf{x}(t)$ the next interaction is productive w.r.t. $x_{max}(t) - x_j(t)$ then we have

$$\Pr\left[X'_{max} - X_j = x_{max} - x_j + 1 | \mathbf{X}(t) = \mathbf{x}(t)\right] = \frac{p_1(t)}{p_1(t) + p_2(t)}$$

We consider two cases. In the first case assume $x_j(t) \geq 20 \cdot \sqrt{n \ln n}$ for all $t \in [t_0, T_1 + T]$. Then for $p'_1(t) = (u(t) \cdot x_{max}(t) + x_j(t') \cdot (n - u(t) - x_j(t))) \cdot n^{-2}$ and $p'_2(t) = (x_{max}(t) \cdot (n - u(t) - x_{max}(t)) + u(t) \cdot x_j(t)) \cdot n^{-2}$ we have

$$\begin{aligned} \frac{p_1}{p_1 + p_2} \geq \frac{p'_1}{p'_1 + p'_2} &= \frac{1}{2} + \frac{p'_1 - p'_2}{2(p'_1 + p'_2)} \\ &= \frac{1}{2} + \frac{(2 \cdot u + x_{max} + x_j - n) \cdot (x_{max} - x_j)}{2((x_{max} + x_j) \cdot n - x_{max}^2 - x_j^2)} \\ &\geq \frac{1}{2} + \frac{(x_j - 16\sqrt{n \ln n}) \cdot (x_{\max} - x_j)}{2((x_{\max} + x_j) \cdot n} \\ &\geq \frac{1}{2} + \frac{(x_{max} - x_j)}{25n} \end{aligned}$$

Observe that if ever $x_{\max}(t') - x_j(t') \leq (x_{\max}(t_0) - x_j(t_0))/2$ for some $t' \in [t_0, T_1 + T]$ the probability to increase $x_{\max}(t') - x_j(t')$ by one is at least $1/2 + (x_{\max}(t_0) - x_j(t_0))/50n \geq 1/2 + (c' \cdot \sqrt{n}\log(n))/(50n)$ assuming a productive interaction. Finally an application of Lemma A.18 for $b = (x_{\max}(t_0) - x_j(t_0))/2$ and $p = 1/2 + c\sqrt{n}\log n/50n$ yields that the probability $x_{\max}(t') - x_j(t') \geq (x_{\max}(t) - x_j(t))/2$ for all $t' \in (t, T]$ is ever violated is at most

$$\left(\frac{1-p}{p}\right)^b = \left(\frac{25n - 2c'\sqrt{n}\log n}{25n + 2c'\sqrt{n}\log n}\right)^b = \left(1 - \frac{4c'\sqrt{n}\log n}{25n + 2c'\sqrt{n}\log n}\right)^b \leq n^{-6}.$$

It remains to show the second case. That is, there exists a time $t \in [t_0, T_1 + \tau]$ such that $x_j(t) < 20\sqrt{n\ln n}$. From the first statement it follows that $x(t') \leq 40\sqrt{n\log n}$ for all $t' \in [t, T_1 + \tau]$. Additionally, we know that $x_{\max}(t') \geq \cdot\sqrt{n}\log n$ and hence, the statement follows by a union bound over both cases w.h.p. $\square$

The following lemma constitutes the foundation of applying the drifting result from [46] which will be used in the proof of Lemma 3.12. In the first part of Lemma 3.11, we consider two important opinions with (almost) the same support. We use an anti-concentration result to show that their support difference quickly reaches $\Omega(\sqrt{n})$. In the second part, we again consider two important opinions and give precise bounds on the probability that their support difference increases by a constant factor. Our proof is based on the gambler's ruin problem.

Lemma 3.11. *Fix two opinions i and j and assume there exists $t_0 \geq T_1$ with $x_i(t_0) \geq x_j(t_0) \geq x_{\max}(t_0) - 4\alpha\sqrt{n}\log n$. Let $T = 40 \cdot n^2/x_{\max}(T_1)$. Then*

1. *If $x_i(t_0) - x_j(t_0) < 4\alpha \cdot \sqrt{n}$ then*

$$\Pr\left[X_i(t_0 + T) - X_j(t_0 + T) \geq 4\alpha \cdot \sqrt{n}\right] \geq e^{-\frac{\alpha^2}{16}}.$$

2. *If $x_i(t_0) - x_j(t_0) \geq 4\alpha \cdot \sqrt{n}$ then*

$$\Pr\left[X_i(t_0+T) - X_j(t_0+T) \geq \min\left\{\frac{3(x_i(t_0) - x_j(t_0))}{2}, 4\alpha\sqrt{n}\log n\right\}\right] \geq 1 - e^{-\frac{x_i(t_0)-x_j(t_0)}{\sqrt{n}}}$$

Proof. We use Observation 3.9 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$. We consider a pair of opinions i and j, which are important at time t_0, and track the evolution of the difference between the support of i and j within the next T interactions.

First, we bound the (i,j)-productive interactions in the time interval $[t_0, t_0 + T]$ interactions. Recall that only (i,j)-productive interactions change the support of Opinion i

or Opinion j. An interaction is (i,j)-productive with probability

$$\begin{aligned}
&\frac{u(t)\cdot x_i(t)+x_i(t)\cdot(n-u(t)-x_i(t))+u(t)\cdot x_j(t)+x_j(t)\cdot(n-u(t)-x_j(t))}{n^2}\\
&\quad=\frac{(x_i(t)+x_j(t))\cdot n-x_i(t)^2-x_j(t)^2}{n^2}\\
&\quad=\frac{(x_i(t)+x_j(t))\cdot(n+x_i(t)-x_j(t))-2x_i(t)^2}{n^2}\\
&\quad\geq\frac{(2x_{\max}-8\alpha\sqrt{n}\log n)\cdot n-2x_{\max}^2}{n^2}\geq\frac{2x_{\max}}{n}\cdot\left(1-\frac{x_{\max}}{n}-\frac{8\alpha\sqrt{n}\log n}{2x_{\max}}\right)\\
&\quad\geq\frac{x_{\max}(T_1)}{2n}
\end{aligned}$$

In the first inequality, we used that i and j are both important. Additionally we use $x_{max}(t)\geq x_{max}(T_1)/2$ and $x_{max}(t)\leq 2n/3$. Thus, assuming i and j both remain important during the whole time interval, an application of Chernoff bounds provides at least $16\cdot n$ many (i,j)-productive interactions in $[t_0,t_0+T]$, w.h.p. For $1\leq i\leq 16\cdot n$ we define t_i as the ith (i,j)-productive interaction in $[t_0,t_0+T]$.

Recall that only (i,j)-productive interactions change the quantity $x_i(t)-x_j(t)$, but other interactions may change the remainder of the configuration, e.g., an additional undecided agent is created. If an interaction is not (i,j)-productive then

$$\Pr\left[X_i(t+1)-X_j(t+1)\neq x_i-x_j|\mathbf{X}(t)=\mathbf{x}\right]=0.$$

If an interaction is (i,j)-productive then

$$\begin{aligned}
&\Pr\left[X_i(t+1)-X_j(t+1)=x_i-x_j+1|\mathbf{X}(t)=\mathbf{x}\right]\\
&\quad=\frac{1}{2}+\frac{u\cdot x_i+x_j\cdot(n-u-x_j)-(u\cdot x_j+x_i\cdot(n-u-x_i))}{2(u\cdot x_i+x_j\cdot(n-u-x_j)+(u\cdot x_j+x_i\cdot(n-u-x_i)))}\\
&\quad=\frac{1}{2}+\frac{(x_i+x_j+2u-n)\cdot(x_i-x_j)}{2((x_i+x_j)\cdot n-x_i^2+x_j^2)}
\end{aligned}$$

Now, we consider two cases. In the first case assume $x_i(t_0)-x_j(t_0)<4\alpha\sqrt{n}$. W.l.o.g. we assume for the rest of the proof that $x_i(t)\geq x_j(t)$ (otherwise, we simply switch the roles of i and j). We consider an (arbitrary) (i,j)-productive interaction t_i and refine the probability from above in the following way

$$
\begin{aligned}
&\Pr\left[X_i(t_i+1) - X_j(t_i+1) = x_i - x_j + 1 | \mathbf{X}(t_i) = \mathbf{x}\right] \\
&\quad = \frac{1}{2} + \frac{(x_i + x_j + 2u - n) \cdot (x_i - x_j)}{2((x_i + x_j) \cdot n - (x_i^2 + x_j^2))} \\
&\quad \geq \frac{1}{2} + \frac{(2x_{max} - 8\alpha\sqrt{n}\log n + 2u - n) \cdot (x_i - x_j)}{2((x_i + x_j) \cdot n - (x_i^2 + x_j^2))} \\
&\quad \geq \frac{1}{2} + \frac{(x_{max} - 8\alpha\sqrt{n}\log n - 16\sqrt{n\ln n}) \cdot (x_i - x_j)}{2((x_i + x_j) \cdot n - (x_i^2 + x_j^2))} \\
&\quad \geq \frac{1}{2}
\end{aligned}
$$

where we use that $x_i(t_i), x_j(t_i) \geq x_{max}(t_i) - 4\alpha\sqrt{n}\log n$ and $x_{max}(t_i) \geq 8\alpha\sqrt{n}\log n - 16\sqrt{n\ln n}$. Thus, the evolution of $x_i(t) - x_j(t)$ over a sequence of $16 \cdot n$ many (i,j)-productive interactions can be viewed as tossing biased coins with success probability larger than 1/2 via standard coupling argument between biased and fair coins. Applying Lemma A.15 with $\delta = \alpha/(2\sqrt{n})$ yields

$$
\Pr\left[\mathrm{Bin}(16 \cdot n, 1/2) \geq \frac{n}{8} + 4\alpha\sqrt{n}\right] \geq e^{-9\delta^2 \cdot 8 \cdot n} = e^{-\frac{9 \cdot \alpha^2}{32}}
$$

Hence, the first statement follows by the union bound with the high probability of events from above.

In the second case we assume $x_i(t) - x_j(t) \geq 4\alpha\sqrt{n}$. Similar to the first case, we refine the probability from above assuming a (i,j)-productive interaction occur

$$
\begin{aligned}
&\Pr\left[X_i(t_i+1) - X_j(t_i+1) = x_i - x_j + 1 | \mathbf{X}(t_i) = \mathbf{x}\right] \\
&\quad = \frac{1}{2} + \frac{(x_i + x_j + 2u - n) \cdot (x_i - x_j)}{2((x_i + x_j) \cdot n - (x_i^2 + x_j^2))} \\
&\quad \geq \frac{1}{2} + \frac{(2x_{max} - 8\alpha\sqrt{n}\log n + 2u - n) \cdot (x_i - x_j)}{2((x_i + x_j) \cdot n - (x_i^2 + x_j^2))} \\
&\quad \geq \frac{1}{2} + \frac{(x_{max} - 8\alpha\sqrt{n}\log n - 8 \cdot \sqrt{n\ln n}) \cdot (x_i - x_j)}{4 \cdot x_{max} \cdot n)} \\
&\quad \geq \frac{1}{2} + \left(1 - \frac{4\alpha \cdot \sqrt{n}\log n + 16 \cdot \sqrt{n\ln n}}{x_{\max}(t_0)}\right) \cdot \frac{x_i - x_j}{4n} \\
&\quad \geq \frac{1}{2} + \frac{x_i - x_j}{12n}
\end{aligned}
$$

Thus, the quantity $x_i(t_i) - x(t_i)_j$ increases by 1 with probability at least $1/2 + (x_i(t) - x(t)_j)/(12n)$ and decreases by 1, otherwise. Observe that starting at time t_0 with $\Delta = x_i(t_0) - x_j(t_0)$ as long as $x_i(t_i) - x_j(t_i) \geq (3/4) \cdot \Delta$ for the first $i \leq 16 \cdot n$ (i,j)-productive interactions in $[t_0, t_0 + T]$ the evolution of $x_i(t) - x_j(t)$ can be viewed as a biased random walk on the line starting at Δ with success probability (i.e., "right step") $p = \frac{1}{2} + \frac{\Delta}{16n}$.

Let $T_{min} = \inf\{t' \in [t_1, t_{c_3 \cdot n}] \mid x_i(t') - x(t')_j = (3/4) \cdot (x_i(t_0) - x_j(t_0))\}$ and $T_{max} = \inf\{t' \geq [t_1, t_{16 \cdot n}] \mid x_i(t') - x(t')_j = 2(x_i(t_0) - x_j(t_0))\}$.

First we bound $\Pr[T_{max} > T_{min}]$. It follows from Lemma A.18 the probability of ever having an excess of $\Delta/8$ "left steps" to "right steps" is at most

$$\left(\frac{1-p}{p}\right)^{\Delta/8} = \left(\frac{8n-\Delta}{8n+\Delta}\right)^{\Delta/8} = \left(1 - \frac{2\Delta}{8n+\Delta}\right)^{\Delta/8} \leq e^{-\frac{\Delta^2}{4\cdot(8n+\Delta)}}$$

Next we bound $\Pr[T_{max} > 16 \cdot n]$. Again we use the assumption $x_i(t_i) - x_j(t_i) \geq (3/4) \cdot \Delta$. Now consider $\tau = 16 \cdot n$ independent Poisson trials ($S_i \in \{-1, 1\}$ for all $i \leq 16 \cdot n$) each with success probability $p = 1/2 + \Delta/16n$. Let $S = \sum_{i=1}^{16 \cdot n} S_i$. Using the Hoeffding bound (Theorem A.6) for $\lambda = \Delta$ we get

$$\begin{aligned}
\Pr[T_{max} > 16 \cdot n] &\leq \Pr[S < \Delta] \\
&= \Pr[S - \mathbb{E}[S] < \Delta - \mathbb{E}[S]] \\
&\leq \Pr[|S - \mathbb{E}[S]| > \mathbb{E}[S] - \Delta] \\
&\leq 2 \cdot e^{-\frac{2\Delta^2}{4 \cdot 16 \cdot n}} \\
&\leq 2 \cdot e^{-\frac{\Delta^2}{32 \cdot n}}
\end{aligned}$$

At last we compute the probability of the event $\mathcal{E}$ that if there exists a time $t \in [T_{max}, t_0 + T]$,i.e., $x_i(t) - x(t)_j = 2\Delta$, then $x_i(t') - x'_j \geq 3/2 \cdot \Delta$ for all $t' \in [t, t_0 + T]$. We can similarly compute this probability with Lemma A.18 as we have shown $\Pr[T_{max} > T_{min}]$. In fact we can simply use $\Pr[T_{max} > T_{min}]$ as an upper bound for $\Pr\left[\bar{\mathcal{E}}\right]$.

To conclude the second statement, we have to show that

$$\Pr[T_{max} \leq c_3 \cdot n \wedge T_{max} \leq T_{min} \wedge \mathcal{E}] \geq 1 - e^{-(x_i(t) - x_j(t))/\sqrt{n}} \tag{3.2}$$

it remains to show

$$\Pr[T_{max} > 16 \cdot n] + \Pr[T_{max} > T_{min}] + \Pr\left[\bar{\mathcal{E}}\right] \leq e^{-(x_i(t) - x_j(t))/\sqrt{n}}$$

To do so, recall $\Delta = x_i(t) - x_j(t) \geq 4\alpha\sqrt{n}$. Then, starting from the left-hand side, we have

$$\begin{aligned}
2 \cdot e^{-\frac{\Delta^2}{32n}} + 2 \cdot e^{-\frac{\Delta^2}{4\cdot(8n+\Delta)}} &= 2 \cdot \left(e^{-\frac{\Delta^2}{32n}} + e^{-\frac{\Delta^2}{4\cdot(8n+\Delta)}}\right) \leq 2 \cdot \left(e^{-\frac{\Delta^2}{32n}} + e^{-\frac{\Delta^2}{36n}}\right) \cdot \left(e^{-\frac{\Delta}{\sqrt{n}}} \cdot e^{\frac{\Delta}{\sqrt{n}}}\right) \\
&= 2 \cdot \left(e^{\frac{\Delta}{\sqrt{n}} \cdot \left(1 - \frac{\Delta}{32\sqrt{n}}\right)} + e^{\frac{\Delta}{\sqrt{n}} \cdot \left(1 - \frac{\Delta}{36\sqrt{n}}\right)}\right) \cdot e^{-\frac{\Delta}{\sqrt{n}}} \leq 2 \cdot \left(\frac{1}{10} + \frac{1}{10}\right) \cdot e^{-\frac{\Delta}{\sqrt{n}}} \leq \frac{1}{5} \cdot e^{-\frac{\Delta}{\sqrt{n}}}
\end{aligned}$$

where we use that the constant α (from the definition of the additive bias) is sufficiently large. Hence, the second statement follows by the union bound with the high probability events from above. □

Now, we are ready to analyze the running time of Phase 2.

Lemma 3.12. *Let* $T_2 = \inf\{\, t \geq T_1 \mid \exists i \in [k] : \forall j \neq i : x_i(t) - x_j(t) \geq \alpha\sqrt{n}\log n \,\}$. *Then* $\Pr\left[T_2 - T_1 \leq 40 \cdot c \cdot n^2 \cdot \log n / x_{\max}(T_1)\right] \geq 1 - 2n^{-2}$.

Proof. We use Observation 3.9 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$.

Recall that an Opinion i is *significant* at time t if $x_i(t) > x_{\max}(t) - \alpha\sqrt{n}\log n$. In the following we call an Opinion i *important* at time t if $x_i(t) > x_{\max}(t) - 4 \cdot \alpha\sqrt{n}\log n$. In the following, we will show that for each pair of important opinions, i and j at time T_1, at least one of them becomes unimportant. Furthermore, we show that no unimportant opinion ever becomes significant. From this, it follows that after $\mathrm{O}\left(n^2/x_{max}(T_1) \cdot \log n\right)$, only one significant opinion remains.

First, we consider a pair of opinions i and j which are important at time T_1 and show that w.h.p. at least one of them becomes unimportant within the next $\tau = 40 \cdot cn^2 \cdot \log n / x_{max}(T_1)$ interactions.

We divide the interactions from $[T_1, T_1+\tau]$ into $c_1 \log n$ subphases of length $40 \cdot n^2/x(T_1)$ each. For $1 \leq i \leq c\log n$ we define $\ell_1 = 1$ and $\ell_i = 1 + (i-1) \cdot n^2/x(T_1)$. Then the ith subphase contains interactions ℓ_i to $(\ell_{i+1} - 1)$. Furthermore, we define t_i as the first interaction in subphase i.

Now, we fix an arbitrary subphase i and consider two cases. If $x_i(t_i) - x_j(t_i) < 4\alpha\sqrt{n}$ then it follows from Lemma 3.11

$$\Pr\left[X_i(t_{i+1}) - X_j(t_{i+1}) \geq 4\alpha\sqrt{n}\right] \geq e^{-\frac{\alpha^2}{16}} \tag{3.3}$$

Otherwise, if $x_i(t_i) - x_j(t_i) \geq 4\alpha\sqrt{n}$ then

$$\Pr\left[X_i(t_{i+1}) - X_j(t_{i+1}) \geq \min\{\,(3/2)\cdot(x_i(t_i) - x_j(t_i)), 4\alpha\sqrt{n}\log n\,\}\right] \geq 1 - e^{-\frac{x_i(t_i) - x_j(t_i)}{\sqrt{n}}} \tag{3.4}$$

In either case, we call such a subphase successful.

In the following, we show that in the interval $[T_1, T_1+\tau]$, there is a sufficient amount of consecutive successful subphases such that at least one of the two opinions becomes unimportant. To do so, we define a function $f : [1, c_1 \log n] \to [0, \log\log n]$, which counts the consecutive number of successful subphases.

$$f(i) = \begin{cases} 0 & \text{if } |x_i(t_i) - x_j(t_i)| < 4\alpha\sqrt{n} \\ j & \text{if } (3/2)^{j-1} \cdot 4\alpha\sqrt{n} \leq |x_i(t_i) - x_j(t_i)| < (3/2)^j \cdot 4\alpha\sqrt{n} \end{cases}$$

Note that either Opinion i or Opinion j is unimportant at the beginning of subphase i if $f(t_i) = \log\log n$.

We define a random walk W over the state space $[0, \log\log n]$ as follows. W has a reflective state 0 and an absorbing state $\log\log n$. Initially, $W(1) = 0$. For $w \in [0, \log\log n - 1]$ the transition probabilities are defined as follows

$$\begin{aligned} \Pr\left[W(t+1) = 1 | W(t) = 0\right] &= e^{-\frac{\alpha^2}{16}} \\ \Pr\left[W(t+1) = w+1 | W(t) = w\right] &= 1 - e^{-2^w} \\ \Pr\left[W(t+1) = 0 | W(t) = w\right] &= e^{-2^w}. \end{aligned}$$

To show that either Opinion i or Opinion j becomes unimportant, which is equivalent to our function f taking on the value $\log\log n$, we define coupling between $f(i)$ and $W(i)$ such that $f(i) \geq W(i)$ for all $i \in [1, c_1 \log n]$.

For $i = 1$ the claim holds trivially since we have $W(1) = 0$ and $f(1) \geq 0$. Now assume for $i \geq 1$ that $f(i) \geq W(i)$. Now, we consider two cases. In the first case assume $|x_i(t_i) - x_j(t_i)| < 4\alpha\sqrt{n}$. Then we know $f(i) = 0$ and hence, $W(i) = 0$. It follows from Eq. (3.3) and $|x_i(t_i) - x_j(t_i)| \geq 0$

$$\Pr[f(i+1) \geq f(i) + 1 | f(i) = 0] \geq e^{-\frac{\alpha^2}{16}} \text{ and}$$
$$\Pr[f(i+1) \geq 0 | f(i) = 0] < 1 - e^{-\frac{\alpha^2}{16}}$$

Likewise, from the definition of W it follows

$$\Pr[W(i+1) = W(i) + 1 | W(i) = 0] = e^{-\frac{\alpha^2}{16}} \text{ and}$$
$$\Pr[W(i+1) = 0 | W(i) = 0] = 1 - e^{-\frac{\alpha^2}{16}}$$

Hence, we can couple the two processes such that the following holds: whenever $W(i)$ is increased by one, then $f(i)$ is increased, too. Whenever $f(i)$ is decreased, $W(i)$ jumps back to zero.

In the second case, we assume

$$4\alpha\sqrt{n} \leq |x_i(t_i) - x_j(t_i)| < \min\left\{ 2(x_i(t_i) - x_j(t_i)), 4\alpha\sqrt{n}\log n \right\} .$$

Then it follows from Eq. (3.4) and $|x_i(t_i) - x_j(t_i)| \geq 0$

$$\Pr[f(i+1) \geq f(i) + 1 | f(i) = 0] \geq 1 - e^{-(x_i(t_i) - x_j(t_i))/\sqrt{n}} \text{ and}$$
$$\Pr[f(i+1) \geq 0 | f(i) = 0] < e^{-(x_i(t_i) - x_j(t_i))/\sqrt{n}}$$

Likewise, from the definition of W it follows

$$\Pr[W(i+1) = W(i) + 1 | W(i) = m] = 1 - e^{-2^m} \text{ and}$$
$$\Pr[W(i+1) = 0 | W(\ell) = m] = e^{-2^m}$$

Observe that

$$1 - e^{-(x_i(t_i) - x_j(t_i))/\sqrt{n}} \geq 1 - e^{-2^{f(i)}} \geq 1 - e^{-2^m} .$$

Again, we can couple the two processes such that $f(i) \geq W(i)$.

Finally an application of Lemma A.10 that w.h.p. there exists $i \in [1, c_1 \log n]$ such that $W(i) = \log\log n$. From this follows that there exists a time $t' \leq [T_1, T_1 + \tau]$ such that $x_i(t') - x_j(t') \geq 4\alpha\sqrt{n}\log n$. This implies, in turn, that at least Opinion j is unimportant. From Statement 2 in Lemma 3.10 it follows that $x_{max}(t) - x_j(t) \geq 2\alpha\sqrt{n}\log n$ for all $t \in [t', T_1 + \tau]$ w.h.p. Hence, the Opinion j does not become significant during the interval. Finally, a union bound over all pairs of initial important opinions at time T_1

yields that all but a single opinion of those important opinions becomes insignificant in the time interval w.h.p.

Now we show that none of the unimportant opinions at time T_1 ever becomes significant during $[T_1, T_1 + \tau]$. First, we fix an Opinion j, which is unimportant at time T_1. Again from Statement 2 in Lemma 3.10 it follows that $x_{max}(t) - x_j(t) \geq 2\alpha\sqrt{n}\log n$ for all $t \in [T_1, T_1 + \tau]$ w.h.p. Hence, all unimportant opinions at time T_1 do not become significant during the time interval by a union bound. At last, the statement follows because all but a single opinion becomes insignificant, hence, $T_2 - T_1 \leq \tau$. □

3.3. From Additive to Multiplicative Bias (Phase 3)

Recall that T_2 is defined as the end of Phase 2, and $\mathbf{x}(T_2)$ is a configuration with an additive bias of $\Omega\big(\sqrt{n}\log n\big)$. In the following we assume w.l.o.g. that $x_1(T_2) \geq x_2(T_2) \ldots \geq x_k(T_2)$.

We start our analysis of Phase 3 with Lemma 3.13 where we show that the support of the largest opinion does not shrink by more than a factor of two. The lemma is equivalent to Statement 2 of Lemma 3.10 from Phase 2.

Lemma 3.13. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \cdot \log n / x_1(T_2)$. Then*

$$\Pr\left[\text{for all } t \in [T_2, T_2 + T] \colon x_1(t) \geq x_1(T_1)/2\right] \geq 1 - n^{-5}.$$

Observation 3.14. *For a given time T_i and configuration $\mathbf{x}(T_i)$ let*

$$\hat{T} = \inf\{t \geq T_i \mid u(t) \notin [(n - x_{\max}(t))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u] \text{ or } x_{\max}(t) < x_{\max}(\hat{t})/2\}$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. Then $(\hat{X}(t))_t$ and $(X(t))_t$ behave the same for every $t \in [T_i, T_i + \mathrm{O}\big(n^2/x_{\max}(T_i) \cdot \log n\big)]$ w.h.p.

We proceed to show that the support difference between Opinion 1 and each other opinion doubles every $\mathrm{O}\big(n^2/x_1(T_2)\big)$ interaction until the ratio between the support of both opinions is sufficiently large. This will be used in Lemma 3.16 to show that after $\mathrm{O}\big(\log n \cdot n^2/x_1(T_2)\big)$ interactions we reach w.h.p. a configuration with a constant factor multiplicative bias.

Lemma 3.15. *Fix an Opinion $i \neq 1$ and assume there exists $t_0 \geq T_2$ with $x_i(t_0) \geq 20\sqrt{n \log n}$ and $x_1(t_0) - x_i(t_0) \geq \alpha\sqrt{n}\log n$. Let $T = 420 \cdot n^2/x_1(T_2)$ and let $\Delta_0 = x_1(t_0) - x_i(t_0)$. Then*

$$\Pr\left[\exists t \in [t_0, t_0 + T] \colon x_1(t) - x_i(t) \geq \min\{2 \cdot \Delta_0,\ 3 \cdot x_i(t)\} \text{ or } x_i(t) < 20\sqrt{n \log n}\right] \geq 1 - 2n^{-3}.$$

Proof. Our proof follows the analysis of the classical Gambler's ruin problem that within $\mathrm{O}\big(n^2/x_1(t_0)\big)$ interactions we track the evolution of $x_1(t) - x_i(t)$ and show it reaches $2(x_1(t_0) - x_i(t_0))$ before $(x_1(t_0) - x_i(t_0))/2$ as long as $x_i(t)$ remains larger than $20\sqrt{n \log n}$.

We use Observation 3.14 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$. First we bound the number of $(1,i)$-productive interactions in the interval $[t_0, t_0+T]$. Assume for the remainder of the proof that $x_i(t) \geq 20\sqrt{n \log n}$ for all $t \in [t_0, t_0+T]$ (otherwise the statement follows immediately). Recall that only $(1,i)$-productive interactions change the quantity $x_1(t) - x_i(t)$. Still, other interactions may change the remainder of the configuration, e.g., an additional undecided agent is created. An interaction is $(1,i)$-productive with probability

$$\begin{aligned}
&\frac{u(t)\cdot x_1(t) + x_1(t)\cdot(n-u(t)-x_1(t)) + u(t)\cdot x_i(t) + x_i(t)(n-u(t)-x_i(t))}{n^2} \\
&\qquad = \frac{(x_1+x_i)\cdot n - x_1^2 - x_i^2}{n^2} = \frac{x_1\cdot(n-x_1) + x_i\cdot(n-x_i)}{n^2} \geq \frac{x_1\cdot(n-x_1)}{n^2} \geq \frac{x_1(T_2)}{6n}
\end{aligned}$$

where we use $x_1(T_2)/2 \leq x_1(t) \leq 2n/3$. Thus, an application of Chernoff bounds provides for $c_1 = c/7$ at least $c_1 \cdot n$ many $(1,i)$-productive interactions in $[t_0, t_0+T]$ w.h.p.

For $1 \leq i \leq c_1 \cdot n$ we define t_i as the ith $(1,i)$-productive interaction in $[t_0, t_0+\tau]$. Then for an arbitrary $i \in [1, c_1\cdot n]$ we have

$$\begin{aligned}
&\Pr\left[X_1(t_i+1) - X_i(t_i+1) = x_1 - x_i + 1 | \mathbf{X}(t_i) = \mathbf{x}\right] \\
&\quad = \frac{1}{2} + \frac{u\cdot x_1 + x_i\cdot(n-u-x_i) - (u\cdot x_i + x_1\cdot(n-u-x_1))}{2(u\cdot x_1 + x_i\cdot(n-u-x_i) + (u\cdot x_i + x_1\cdot(n-u-x_1)))} \\
&\quad = \frac{1}{2} + \frac{(x_1+x_i+2u-n)\cdot(x_1-x_i)}{2((x_1+x_i)\cdot n - (x_1^2+x_i^2))} \\
&\quad \geq \frac{1}{2} + \frac{(x_i - 16\sqrt{n\ln n})(x_1-x_i)}{2(n(x_1+x_i) - (x_1^2+x_i^2))} \\
&\quad \geq \frac{1}{2} + \frac{(x_i - 16\sqrt{n\ln n})(x_1-x_i)}{2n(4x_i+x_i)} \\
&\quad = \frac{1}{2} + \left(1 - \frac{16\sqrt{n\ln n}}{x_i}\right)\cdot\frac{(x_1-x_i)}{10n} \\
&\quad \geq \frac{1}{2} + \frac{x_1(t_0) - x_i(t_0)}{60n}
\end{aligned}$$

where we use $x_1 < 4\cdot x_i$ and $x_i > 20\sqrt{n\log n}$ (otherwise the statement follows immediately). Additionally note that the last inequality holds as long as $x_1(t) - x_i(t) \geq (5/6)\cdot(x_1(t_0) - x_i(t_0))$.

Thus, the quantity $x_1(t_i) - x_i(t_i)$ increases by 1 with probability at least $p = 1/2 + (x_1(t_0) - x_i(t_0))/(60n)$ and decreases by 1, otherwise. Observe that starting at time t_0 with $\Delta = x_1(t_0) - x_i(t_0)$ as long as $x_1(t_i) - x_i(t_i) \geq \Delta/2$ for the first $i \leq c_1\cdot n$ many $(1,i)$-productive interactions in $[t_0, t_0+T]$ the evolution of $x_1(t_i) - x_i(t_i)$ can be viewed as a biased random walk on the line starting at Δ where a "right step" happens with probability p and "left step" with probability $1-p$, otherwise. The correctness follows from a standard coupling argument between two biased coins. Formally let $T_{min} = \inf\{t' \geq t_0 \mid x_1(t') - x_i(t') = (5/6)\cdot(x_1(t_0) - x_i(t_0))\}$ and $T_{max} = \inf\{t' \geq t_0 \mid x_1(t') - x_i(t') =$

$2(x_1(t_0) - x_i(t_0))\}$. First we bound $\Pr[T_{max} > T_{min}]$. It follows from Lemma A.18 the probability of ever having an excess of $(5/12) \cdot \Delta$ "left steps" to "right steps" is at most

$$\left(\frac{1-p}{p}\right)^{(5/12)\cdot\Delta} = \left(\frac{30n-\Delta}{30n+\Delta}\right)^{(5/12)\cdot\Delta} = \left(1 - \frac{2\Delta}{30n+\Delta}\right)^{(5/12)\cdot\Delta} \leq e^{-\frac{(5/6)\cdot\Delta^2}{30n+\Delta}} \leq n^{-5}$$

where we use $\Delta \geq \alpha\sqrt{n}\log n$.

Next we bound $\Pr[T_{max} > c_1 \cdot n]$. Again we use the assumption $x_i(t_i) - x_j(t_i) \geq (1/2) \cdot \Delta$. Now consider $\tau = c_1 \cdot n$ independent Poisson trials ($S_i \in \{-1, 1\}$ for all $i \leq c_1 \cdot n$) each with success probability $p = 1/2 + \Delta/60n$. Let $S = \sum_{i=1}^{c_1 \cdot n} S_i$. Using the Hoeffding bound (Theorem A.6) for $\lambda = \Delta$ we get

$$\begin{aligned}
\Pr[T_{max} > c_1 \cdot n] &\leq \Pr[S < \Delta] \\
&= \Pr[S - \mathbb{E}[S] < \Delta - \mathbb{E}[S]] \\
&\leq \Pr[|S - \mathbb{E}[S]| > \mathbb{E}[S] - \Delta] \\
&\leq 2 \cdot e^{-\frac{2\Delta^2}{4\cdot c_1 \cdot n}} \\
&\leq 2 \cdot e^{-\frac{\Delta^2}{2\cdot c_1 \cdot n}} \\
&\leq n^{-5}
\end{aligned}$$

Hence, the statement follows by the union bound over the high probability events from above. □

Now, we are ready to analyze the running time of Phase 3.

Lemma 3.16. *Assume that* $\mathbf{x}(T_2)$ *is a configuration with* $x_1(T_2) - x_i(T_2) \geq \alpha\sqrt{n}\log n$ *for all* $i \neq 1$. *Let* $T_3 = \inf\{\, t \geq T_2 \mid \forall i \neq 1 : x_1(t) \geq 2x_i(t) \,\}$. *Then*

$$\Pr\left[T_3 - T_2 \leq 420 \cdot n^2 \cdot \log n / x_1(T_2)\right] \geq 1 - 2n^{-2}.$$

Proof. The main idea of this proof is to repeatedly apply Lemma 3.15 to each Opinion $i \neq 1$ until either the support of Opinion 1 becomes larger than $2n/3$ or the support of Opinion i becomes less than $20 \cdot \sqrt{n \log n}$. In both cases, it then follows that the ratio between the support of Opinion 1 and Opinion i is larger than two, and there is a time when there is a multiplicative bias between the first opinion and each other opinion. We use Observation 3.14 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$. Let $\tau = 420 \cdot n^2 \cdot \log n / x_{\max}(T_2)$ and fix an Opinion $i \neq 1$ at time T_2 with $x_i(T_2) \geq 20\sqrt{n \log n}$. We divide the interactions from $[T_2, T_2 + \tau]$ into $\log n$ subphases of length $420 \cdot n^2 / x_1(T_2)$ each. For $1 \leq j \leq \log n$ we define $\ell_1 = 1$ and $\ell_j = 1 + (j-1) \cdot 420 \cdot n^2 / x_1(T_2)$. Then the jth subphase contains interactions ℓ_j to $(\ell_{j+1} - 1)$. Furthermore, we define t_j as the first interaction in subphase j. Now fix an arbitrary subphase j. It follows from Lemma 3.15 that there exists a time $t' \in [t_j, t_{j+1}]$ such that w.h.p. either $x_1(t) - x_i(t) \geq \min\{2 \cdot (x_1(t_j) - x_i(t_j)),\ 3 \cdot x_i(t)\}$ or $x_i(t) < 20\sqrt{n \log n}$.

We apply Lemma 3.15 to each subphase. From the union bound over all subphases and all opinions, it follows that after at most $\log n$ subphases w.h.p. there exists for each Opinion i a time $t'_i \in [T_2, T_2 + \tau]$ with either (a) $x_1(t'_i) - x_i(t'_i) \geq 2n/3$ or (b) $x_i(t'_i) < 20\sqrt{n \log n}$ or (c) $x_1(t'_i) \geq 4 \cdot x_i(t'_i)$. In the following, we consider three cases.

Case (a) There exists an Opinion $i \neq 1$ such that $x_1(t'_i) - x_i(t'_i) \geq 2n/3$. Hence, we have at t'_i a constant multiplicative bias between Opinion 1 and all other opinions $i \neq 1$. From this, the statement follows immediately with $T_3 = t'_i$.

Case (b) For Opinion i there exists a t'_i such that $x_i(t'_i) < 20\sqrt{n \log n}$. From Lemma 3.10 it follows that $x_i(t) \leq 40\sqrt{n \log n}$ for all $t \in [t'_i, T_2 + \tau]$ w.h.p. Additionally we know $x_1(t) \geq x_1(T_2)/2 \geq c'\sqrt{n} \log^2 n$ for all $t \in [t'_i, T_2 + \tau]$. Hence, $x_1(t)/x_i(t) \gg 2$ for all $t \in [t'_i, T_2 + \tau]$ and, from the viewpoint of Opinion i we have that T_3 can take on an arbitrary value in $[t'_i, T_2 + \tau]$.

Case (c) For Opinion i there exists a t'_i such that $x_1(t'_i) \geq 4 \cdot x_i(t'_i)$. From the claim below it follows that w.h.p. $x_1(t) \geq 2x_i(t)$ for all $t \in [t'_i, T_2 + \tau]$ and from the viewpoint of Opinion i we have that T_3 can take on an arbitrary value in $[t'_i, T_2 + \tau]$.

Now Lemma 3.16 follows either immediately from Case (a). Or we can apply Case (b) or Case (c) for each Opinion $i \neq 1$ and then we can choose $T_3 = T_2 + \tau$. It remains to show the following claim.

Claim 3.17. *Let j be an arbitrary subphase and let $t_0 \in [t_j, t_{j+1}]$. Fix an Opinion i and assume $x_i(t_0) \in [20 \cdot \sqrt{n \log n}, x_1(t_0)/4]$. Then $x_1(t) \geq 2 \cdot x_i(t)$ for all $t \in [t_0, T_2 + \tau]$.*

Proof. Recall that we showed for "small" opinions with $x_i < 20 \cdot \sqrt{n \log n}$ that the multiplicative bias is always larger than a constant. Furthermore recall that $\tau = 420 \cdot n^2 \cdot \log n / x_1(T_2)$ and T_2 is the end of Phase 2. Assume w.l.o.g. that we start with the analysis at time $t_0 = 0$. We use Observation 3.14 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$. An interaction is productive w.r.t. to x_1 and x_i (meaning that either x_1 or x_i change) with probability

$$p = \frac{u \cdot x_1 + x_1 \cdot (n - u - x_1) + u \cdot x_i + x_i \cdot (n - x_i)}{n^2} \leq 3 \cdot \frac{x_1(T_2)}{n}$$

for $x_1 \geq 2 \cdot x_i$. It follows from an application of Chernoff bounds that within a sequence of T interactions, the number of 1-productive interactions is at most $2T \cdot p = 6T \cdot (x(T_2)) \cdot \log n/n \leq 2520n \cdot \log n$ with probability at least $1 - n^{-10}$. We define $\tau' = 2520n \cdot \log n$ and consider τ' productive interactions. Let $Z(t) = x_1(t) - 2x_i(t)$. We aim to use the Hoeffding bound (Lemma A.7) to show that this quantity does not decrease significantly throughout τ' productive interactions. Hence, we must calculate the probability that $Z(t)$ increases or decreases. Note that the maximum one step change in $\in [-2, 2]$. Assuming the next interaction is a 1-productive interaction for $x(t)$ we have

$$\begin{aligned}
\Pr\left[Z(t+1)-Z(t)=1 | \mathbf{X}(t)=\mathbf{x}\right] &= \frac{1}{p}\cdot\frac{x_1\cdot u}{n^2} &&= \frac{x_1\cdot u}{x_1\cdot(n-x_1)+x_i\cdot(n-x_i)}\\
\Pr\left[Z(t+1)-Z(t)=-1 | \mathbf{X}(t)=\mathbf{x}\right] &= \frac{1}{p}\cdot\frac{x_1\cdot(n-u-x_1)}{n^2} &&= \frac{x_1\cdot(n-x_1)}{x_1\cdot(n-x_1)+x_i\cdot(n-x_i)}\\
\Pr\left[Z(t+1)-Z(t)=-2 | \mathbf{X}(t)=\mathbf{x}\right] &= \frac{1}{p}\cdot\frac{x_i\cdot u}{n^2} &&= \frac{x_i\cdot(n-x_i)}{x_1\cdot(n-x_1)+x_i\cdot(n-x_i)}\\
\Pr\left[Z(t+1)-Z(t)=2 | \mathbf{X}(t)=\mathbf{x}\right] &= \frac{1}{p}\cdot\frac{x_i\cdot(n-u-x_i)}{n^2} &&= \frac{x_1\cdot u}{x_1\cdot(n-x_1)+x_i\cdot(n-x_i)}
\end{aligned}$$

Therefore

$$\begin{aligned}
\mathbb{E}\left[Z(t+1)-Z(t) | \mathbf{X}(t)=\mathbf{x}\right] &= \frac{x_1\cdot u - 2x_i\cdot u - x_1(n-u-x_1) + 2x_i(n-u-x_i)}{x_1\cdot(n-x_1)+x_i\cdot(n-x_i)}\\
&\geq \frac{(x_1-2x_i)\cdot(2u-n+x_1)+x_1\cdot x_i}{n-x_1+x_i\cdot(n-x_i)}
\end{aligned}$$

where we used that $x_1 \geq 2\cdot x_i$. Since $x_i \geq 20\cdot\sqrt{n\log n}$ and $u \geq n/2 - x_1/2 - 8\cdot\sqrt{n\log n}$ (Lemma 3.6), we get

$$\begin{aligned}
\mathbb{E}\left[Z(t+1)-Z(t) | \mathbf{X}(t)=\mathbf{x}\right] &\geq \frac{(x_1-2x_i)\cdot(2u-n+x_1)+x_1\cdot x_i}{n-x_1+x_i\cdot(n-x_i)}\\
&\geq \frac{(x_1-2x_i)\cdot(-16\sqrt{n\log n})+4x_1\cdot x_i/5+x_1\cdot x_i/5}{n-x_1+x_i\cdot(n-x_i)}\\
&\geq \frac{2x_i\cdot 16\sqrt{n\log n}+x_1\cdot x_i/5}{n-x_1+x_i\cdot(n-x_i)} > 0
\end{aligned}$$

Thus, we have $\mathbb{E}\left[Z(t+1)-Z(t) | \mathbf{X}(t)=\mathbf{x}\right] \geq 0$ if $x_1 \geq 2\cdot x_i$ and $x_i \geq 20\cdot\sqrt{n\log n}$.

We are ready to apply the Hoeffding bound from Lemma A.7. Observe that $|Z(t+1)-Z(t)| \leq 2$ for all $t \in [0, \tau'-1]$ and

$$S = \sum_{t=0}^{\tau'-1} Z(t+1)-Z(t) = Z(\tau') - Z(0)$$

Then it follows from Hoeffding bound (Lemma A.7) with $\lambda = Z(0) \geq x_1(0)/2$ that

$$\Pr\left[S < Z(0) - c_1\cdot Z(0)\right] \leq \Pr\left[S - \mathbb{E}\left[S\right] < -\lambda\right] \leq \exp\left(-\frac{2\lambda^2}{16\tau'}\right) \leq n^{-c\cdot\log^2(n)}$$

for some constant c. Thus, we have that w.h.p. $Z(\tau') \geq Z(0)$. Then,

$$\frac{x_1(\tau')}{x_i(\tau')} = \frac{x_1(\tau')-2x_i(\tau')}{x_i(\tau')} + \frac{2x_i(\tau')}{x_i(\tau')} = \frac{x_1(\tau')-2x_i(\tau')}{x_i(\tau')} + 2 = \frac{Z(\tau')}{x_i(\tau')} + 2 \geq 2.$$

Thus, w.h.p. $x_1(\tau') \geq x_i(\tau')$. The claim follows from the union bound over all $\tau' < n^3$ interactions. $\square$

$\square$

3.4. From Multiplicative Bias To Absolute Majority (Phase 4)

Recall that T_3 is the end of Phase 3, and $\mathbf{X}(T_3)$ is a configuration with multiplicative bias. In the remainder we assume that the bias is at least two. [2] In the following we assume w.l.o.g. that $x_1(T_3) > x_2(T_3) \geq \ldots \geq x_k(T_3)$. The main result for this phase is Lemma 3.21, where we show that the multiplicative bias is grown into a unique majority opinion with support at least $2n/3$ within $\mathrm{O}\big(n \log n + n^2/x_1(T_3)\big)$ interactions, w.h.p. To do so, we first need an improved bound on the number of undecided agents that we reach at the time $T_3 + \mathrm{O}\big(n \log n\big)$. Additionally, we have to show that in the meantime, both x_1 and the multiplicative bias decrease only by a small constant fraction (Lemma 3.18 and Lemma 3.19). The proofs of both lemmas are similar to the proofs of Lemma 3.7 and Claim 3.17, respectively.

Lemma 3.18. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \log n / x_1(T_3)$. Then*

$$\Pr\left[\textit{for all } t \in [T_3, T_3 + T]\colon x_1(t) \geq x_1(T_3)/2\right] \geq 1 - n^{-5}.$$

Proof. Let

$$\hat{T} = \inf\left\{\, t \geq T_3 \;\mid\; u(t) \notin [(n - x_{\max}(t'))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u] \,\right\}$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. From Lemma 3.4 and Lemma 3.6 it follows $\hat{T} - t = \Omega\big(n^2/x_{\max}(t) \cdot \log n\big)$ w.h.p. Thus, $(\mathbf{X}(t))_t$ and $(\hat{\mathbf{X}}(t))_t$ behave the same between time t and $t + \mathrm{O}\big(n^2/x_{\max}(T_3) \cdot \log n\big)$. As long as $x_1(t') \leq 2 \cdot x_1(T_3)$ an interaction is 1-productive with probability

$$\frac{u \cdot x_1 + x_1 \cdot (n - u - x_1)}{n^2} = \frac{x_1 \cdot (n - x_1)}{n^2} \leq 2 \cdot \frac{x(T_3)}{n}$$

It follows from an application of Chernoff bounds that within a sequence of $c \cdot n^2 \cdot \log n / x_1(T_3)$ interactions, the number of 1-productive interactions is at most $4 \cdot c \cdot n \log n$ with probability at least $1 - n^{-10}$. Now consider $\tau = 4 \cdot c \cdot n \log n$ such productive interactions and let Z_t denote the change w.r.t. $x_1(t)$, i.e., the support of the largest opinion increase or decrease by one, respectively. That is, assuming the next interaction is a 1-productive interaction for $x(t)$ we have

$$\begin{aligned}
\Pr\left[Z_t = 1\right] &= \frac{u \cdot x_1}{u \cdot x_1 + x_1 \cdot (n - u - x_1)} = \frac{u}{(n - x_1)} \\
\Pr\left[Z_t = -1\right] &= 1 - \Pr\left[Z_t = 1\right]
\end{aligned}$$

[2] If we start with an initial multiplicative bias $(1 + \epsilon) < 2$, we skip Phase 2 and Phase 3. Technically, we follow the proof of Lemma 3.16 to amplify the bias from $(1 + \epsilon)$ up to factor 2 where we essentially apply Lemma 3.15 constantly many times.

Therefore

$$\mathbb{E}[Z_t] = \frac{u - (n - u - x_1)}{n - x_1} = \frac{2 \cdot u + x_1 - n}{n - x_1} \geq -48 \cdot \frac{\sqrt{n \ln n}}{n}$$

Let Z be the sum of Z_t for all $t \in [1, \tau]$. Then it follows from Hoeffding bound with $\lambda = x_1(T_2)/2 - 200 \cdot \sqrt{n} \ln^{3/2} n$

$$\Pr\left[Z < -\frac{1}{2} \cdot x_1(T_1)\right] \leq \Pr[Z < \mathbb{E}[Z] - \lambda] \leq e^{-\frac{2\lambda^2}{4\tau}} \leq n^{-10}$$

Note that if (ever) $x_1(t') > 2 \cdot x_1(T_3)$ for some $t' \in [T_3, T_3 + T]$ the statement hold by the union bound and the previous part. Thus, starting with $x_1(T_3)$ throughout the next $c \cdot n^2/x_{\max}(T_2) \cdot \log n$ interactions $x_{\max}(t) \geq x_{\max}(T_3)/2$ with probability at least $1 - n^{-5}$. □

Lemma 3.19. *Assume that* $\mathbf{x}(T_3)$ *is a configuration with* $x_1(T_3) \geq 2 \cdot x_i(T_3)$ *for all* $i \neq 1$. *Then*

$$\textit{for all } \; i \neq 1 \quad \Pr\left[\textit{for all } t \in [T_3, 111 \cdot n^2/x_1(T_3)] \colon x_1(t) \geq 7/4 \cdot x_i(t)\right] \geq 1 - 2n^{-3}.$$

Proof. The proof is similar to the proof of Claim 3.17 using $Z(t) = x_1(t) - 7x_i(t)/4$ instead of $Z(t) = (t) - 2x_i(t)$. We have $Z(0) = x_1(0) - 7x_i(0)/4 \geq x_1(0)/8$ and $\mathbb{E}[Z(t+1) - Z(t)] \geq 0$. □

Next, we improve the lower bound on the number of undecided agents from Lemma 3.4. Recall that T_4 is the end of Phase 4, defined as $T_4 = \inf\{\, t \geq T_3 \mid x_1(t) \geq 2n/3 \,\}$.

Lemma 3.20. *Let* $T_u = \inf\{\, t \geq T_3 \mid u(t) \geq n/2 - 7/8 \cdot x_1(t) \,\}$. *Then*

$$\Pr[\min(T_4, T_u) - T_3 \leq \lceil 7n \ln n \rceil] \geq 1 - 4n^{-3}.$$

Proof. To bound $T_u - T_3$, we follow the proof of Lemma 3.2. Let $\alpha = 7/8$ and let

$Z(t) = n - 2u(t) - \alpha \cdot x_1(t)$ and let $r^2 = \sum_{i\in[k]} x_i^2$. Then

$$
\begin{aligned}
&\mathbb{E}[Z(t)-Z(t+1)|\mathbf{X}(t)=\mathbf{x}]\\
&\quad=-\frac{x_1\cdot u}{n^2}\cdot(2-\alpha)-\sum_{i=2}^{k}\frac{x_i\cdot u}{n^2}\cdot 2-\frac{x_1(n-u-x_1)}{n^2}\cdot(-2+\alpha)-\sum_{i=2}^{k}\frac{x_i(n-u-x_i)}{n^2}\cdot(-2)\\
&\quad=(2-\alpha)\cdot\frac{-(x_1\cdot u)+x_1\cdot(n-u-x_1)}{n^2}+2\cdot\sum_{i=2}^{k}\frac{-(x_i\cdot u)+x_i\cdot(n-u-x_i)}{n^2}\\
&\quad=(2-\alpha)\cdot\frac{x_1\cdot(n-2u-x_1)}{n^2}+2\cdot\sum_{i=2}^{k}\frac{x_i\cdot(n-2u)-x_i^2}{n^2}\\
&\quad=2\cdot\frac{x_1\cdot(n-2u)-x_1^2}{n^2}-\alpha\cdot\frac{x_1\cdot(n-2u-x_1)}{n^2}+2\cdot\frac{(n-u-x_1)(n-2u)}{n^2}-2\cdot\frac{r^2-x_1^2}{n^2}\\
&\quad=2\cdot\frac{x_1\cdot(n-2u)+(n-u-x_1)(n-2u)-r^2}{n^2}-\frac{\alpha\cdot x_1\cdot(n-2u-x_1)}{n^2}\\
&\quad=2\cdot\frac{(n-2u)\cdot(n-u)-r^2}{n^2}-\frac{\alpha\cdot x_1\cdot(n-2u-x_1)}{n^2}\\
&\quad=2\cdot\frac{(n-2u-\alpha\cdot x_1)\cdot(n-u)+\alpha\cdot x_1\cdot(n-u)-r^2}{n^2}-\frac{2\alpha\cdot x_1\cdot(n-u)-\alpha\cdot x_1\cdot(n+x_1)}{n^2}\\
&\quad=2\cdot\frac{(n-2u-\alpha\cdot x_1)\cdot(n-u)}{n^2}+\frac{2\alpha\cdot x_1\cdot(n-u)-2r^2}{n^2}-\frac{2\alpha\cdot x_1\cdot(n-u)-\alpha\cdot x_1\cdot(n+x_1)}{n^2}\\
&\quad=2\cdot\frac{Z(t)\cdot(n-u)}{n^2}+\frac{\alpha\cdot x_1\cdot(n+x_1)-2r^2}{n^2}\\
&\quad=\frac{Z(t)}{2n}+\frac{1}{n^2}\cdot\left(3n^2/2-2n\cdot u+\alpha\cdot x_1\cdot n+\alpha\cdot x_1^2-2r^2\right)
\end{aligned}
$$

Note that $r^2 = \sum_{i=1}^{k} x_i^2 \leq x_1^2 + (4/7)\cdot x_1 \cdot \sum_{i=2}^{k} x_i = x_1^2 + (4/7)\cdot x_1 \cdot (n-u-x_1)$. Furthermore, by Lemma 3.6 and Lemma 3.4 and using $x_1 \leq 2n/3$, we have w.h.p. $u < n/2 + \Delta_u$ and $u \geq n/2 - x_1/2 - o(x_1) \geq n/8$ for sufficiently large n. For the last expression in parentheses, we calculate

$$
\begin{aligned}
&3n^2/2 - 2n\cdot u + \alpha\cdot x_1\cdot n + \alpha\cdot x_1^2 - 2r^2\\
&\qquad\geq 3n^2/2 - 2n\cdot u + \alpha\cdot x_1\cdot n + \alpha\cdot x_1^2 - 2(x_1^2 + (4/7)\cdot x_1\cdot(n-u-x_1))\\
&\qquad\geq n^2/2 - 2n\cdot \Delta_u + \alpha\cdot x_1\cdot n + \alpha\cdot x_1^2 - 2(x_1^2 + (4/7)\cdot x_1\cdot(n-u-x_1))\\
&\qquad\geq n^2/2 - 2n\cdot \Delta_u + \alpha\cdot x_1\cdot n + \alpha\cdot x_1^2 - 2(x_1^2 + (4/7)\cdot x_1\cdot((7/8)\cdot n-x_1))\\
&\quad\geq 0
\end{aligned}
$$

for $\alpha = 7/8$.

The remainder of the proof is identical to that of Lemma 3.2 except that we note that either $\mathbb{E}[Z(t)-Z(t+1)|\mathbf{X}(t)=\mathbf{x}] \geq Z/(n)$ or at some time $t \in [T_3, n^3] : x_1(t) < 7/4 \cdot x_i(t)$ for some $i > 1$. The latter event is ruled out w.h.p. by Lemma 3.19.

We now apply Theorem A.9 with $r = 3\ln n$, $s_0 = n - 2u(0) - 7/8 \cdot x_1(0) \le n$, $s_{min} = 1$, $\delta = 1/(2n)$ and get with $T = \inf\{\, t \ge T_3 \mid Z(t) \le 0 \,\}$

$$\begin{aligned}\Pr\left[T - T_3 > \lceil 7n\ln n\rceil\right] &\le \Pr\left[T - T_3 > \left\lceil \frac{6\cdot \ln n + \ln(n - 2u(T_3) - 7/8\cdot x_1(T_3))}{1/(2n)} \right\rceil\right] \\ &\le e^{-3\cdot\ln(n)} = n^{-3}\,.\end{aligned}$$

Note that if ever $x_1(t) \ge 2n/3$ for $t < \lceil 7n\ln n\rceil$, we have $T_4 \le \lceil 7n\ln n\rceil$. Otherwise, we have shown that $T_u \le \lceil 7n\ln n\rceil$. Hence, overall we get $\min\{T_u, T_4\} - T_3 \le \lceil 7n\ln n\rceil$. □

We are ready to analyze the running time of Phase 4.

Lemma 3.21. *Assume that $\mathbf{x}(T_3)$ is a configuration with $x_1(T_3) \ge 2\cdot x_i(T_3)$ for all $i \ne 1$. Then, there exists a constant c such that Then*

$$\Pr\left[T_4 - T_3 \le 7n\ln n + 444\cdot n^2/x_1(T_3)\right] \ge 1 - 2n^{-2}.$$

Proof. To show the statement, we require the following two auxiliary results. First, we establish in Claim 3.22 that the improved bound on the undecided agents from Lemma 3.20 holds throughout the remainder of the phase. As before, we define $T_u = \inf\{\, t \ge T_3 \mid u(t) \ge n/2 - 7/8\cdot x_1(t)/2 \,\}$ and recall that T_4 denotes the end of the phase. The proof follows along the lines of the proof of Lemma 3.6 with the new $Z(t)$.

Claim 3.22. $\Pr\left[\textit{for all } t \in [T_u, \min\{n^3, T_4\}]\colon u(t) \ge n/2 - 7/16\cdot x_1(t) - 8\cdot\sqrt{n\ln n}\right] \ge 1 - 4n^{-3}$.

Proof. We follow the proof idea of Theorem 6 in [64]. We define a new set of random variables with $Y(t) = \exp(\eta\cdot Z(t))$ for $t \ge T$ and $\eta = \sqrt{\ln n/n}$ and let $z_0 = 4\eta\cdot n$.

Fix an arbitrary $i \ge 0$. We first give a bound for $\mathbb{E}\left[Y(i+1) - Y(i) \mid Z(i) = z\right]$. Note that $Z(i+1) - Z(i) \in [-2, 2]$. We get

$$\begin{aligned}&\mathbb{E}\left[Y(i+1) - Y(i) \mid Z(i) = z\right] \\ &\quad = \mathbb{E}\left[e^{\eta\cdot Z(i+1)} - e^{\eta\cdot Z(t)} \mid Z(i) = z\right] \\ &\quad = e^{\eta\cdot z}\cdot\mathbb{E}\left[e^{\eta\cdot(Z(i+1)-z)} - 1 \mid Z(i) = z\right] \\ &\quad = e^{\eta\cdot z}\cdot\sum_{j\in\{\,-2,-9/8,0,9/8,2\,\}}(e^{\eta\cdot j} - 1)\cdot\Pr\left[Z(i+1) - z = j | Z(i) = z\right]\end{aligned}$$

We derive the following bound for $\exp(\eta\cdot j) - 1$. Since $\exp(x) \le 1 + x + x^2$ for $x \le 1$ and $\eta \to 0$ for large n, we have $\exp(2\eta) \le 1 + 2\eta + (2\eta)^2 = 1 + 2\eta + \eta\cdot z_0/n$. For $j \in [-2, 2]$, we thus have $\exp(\eta j) - 1 \le \eta j + \eta\cdot z_0/n$. We know that $\mathbb{E}\left[Z(i+1) - Z(i) | Z(i) = z\right] \le -\frac{z}{n}$ w.h.p. from Part 1.

Thus, for all $z \geq z_0$ we have

$$\begin{aligned}
&\mathbb{E}\left[Y(i+1) - Y(i) \mid Z(i) = z\right] \\
&\quad \leq e^{\eta \cdot z} \cdot \sum_{j \in \{-2, -9/8, 0, 9/8, 2\}} (\eta \cdot j + \eta \cdot z_0/n) \cdot \Pr\left[Z(i+1) - z = j \mid Z(i) = z\right] \\
&\quad = e^{\eta \cdot z} \cdot \eta \cdot \left(\mathbb{E}\left[Z(i+1) - Z(i) \mid Z(i) = z\right] + z_0/n\right) \leq 0.
\end{aligned}$$

In total, we get

$$\mathbb{E}\left[Y(t)\right] = \mathbb{E}\left[Y(0)\right] + \sum_{i=0}^{t-1} \mathbb{E}\left[Y(i+1) - Y(i)\right] \leq 1.$$

We can apply Markov's inequality since $\forall t \geq 0 : Y(t) \geq 0$. Thus,

$$\Pr\left[Z(t) \geq 2z_0\right] = \Pr\left[Y(t) \geq \exp(2\eta z_0)\right] \leq \frac{\mathbb{E}\left[Y(t)\right]}{n^8} \leq n^{-8}.$$

Finally, we apply the union bound over $n^3 - T \leq n^3$ interactions. $\square$

Observation 3.23. *For a given time T_i and configuration $\mathbf{x}(T_i)$ let*

$$\hat{T} = \inf\left\{ t \geq T_i \mid u(t) \notin \left[(n - 7/8 \cdot x_1(t))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u\right] \text{ or } x_{\max}(t) < x_{\max}(\hat{t})/2 \right\}$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. Then $(\hat{X}(t))_t$ and $(X(t))_t$ behave the same for every $t \in [T_i, T_i + \mathrm{O}\big(n^2/x_{\max}(T_i) \cdot \log n\big)]$ w.h.p.

Next, in Claim 3.24, we bound the number of interactions until the support of Opinion 1 has doubled. Similarly to Lemma 3.15, the proof uses the classical gambler's ruin problem to show that in a sequence of $c \cdot n^2/x_1(t)$ interactions, the support of Opinion 1 doubles w.h.p. before it halves.

Claim 3.24. *Let $\mathbf{x}(t)$ be a configuration with $u(t) \geq n/2 - 7/16 \cdot x_1(t) - 8 \cdot \sqrt{n \ln n}$ and $x_1(t) < 2n/3$. We define $t' = c \cdot n^2/x_1(t)$ for a suitable chosen constant c. Then*

$$\Pr\left[\exists t' \in [t, t+t'] \colon x_1(t') \geq 2 \cdot x_1(t) \quad \text{or} \quad x_1(t) \geq 2n/3\right] \geq 1 - n^{-3}.$$

Proof. The proof is similar to the proof of Lemma 3.15, but instead of analyzing the quantity $x_1(t) - x_i(t)$, we only analyze the growth of $x_1(t)$ directly. We use Observation 3.23 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$. First we bound the number of 1-productive interactions in the interval $[t_0, t_0 + \tau]$ for $\tau = 111 \cdot n^2/x_1(t_0)$. Assume for the remainder of the proof that $x_1(t) < 2n/3$ for all $t \in [t_0, t_0 + \tau]$ (otherwise the statement follows immediately). Recall that only 1-productive interactions change the quantity $x_1(t)$, but other interactions may change the remainder of the configuration, e.g., an additional undecided agent is created.

An interaction is 1-productive with probability

$$\frac{u(t)\cdot x_1(t)+x_1(t)\cdot(n-u(t)-x_1(t))}{n^2} \geq \frac{x_1(t)\cdot(n-x_1(t))}{n^2} \geq \frac{x_1(t)}{3n} \geq \frac{x_1(t_0)}{10n}$$

where we use $x_1(t_0)/2 \leq x_1(t) < 2n/3$.

Thus, an application of Chernoff bounds provides for $c_1 = 110$ at least $c_1 \cdot n$ many 1-productive interactions in $[t_0, t_0+\tau]$ w.h.p. For $1 \leq i \leq c_1 \cdot n$ we define t_i as the ith 1-productive interaction in $[t_0, t_0+\tau]$. Then for an arbitrary $i \in [1, c_1 \cdot n]$ we have

$$\begin{aligned}
&\Pr\left[X_1(t_i+1) = x_1+1 \mid \mathbf{X}(t_i) = \mathbf{x}\right] \\
&\quad= \frac{1}{2} + \frac{x_1\cdot u - x_1(n-u-x_1)}{2\cdot(x_1\cdot u + x_1(n-u-x_1))} = \frac{1}{2} + \frac{2u-n+x_1}{2(n-x_1)} \geq \frac{1}{2} + \frac{x_1(t_0)}{110n}
\end{aligned}$$

Note that the last inequality holds as long as $x_1(t) \geq x_1(t_0)/2$. Thus, the quantity $x_1(t_i)$ increases by 1 with probability at least $p = 1/2 + x_1(t_0)/(110n)$ and decreases by 1, otherwise. Observe that starting at time t_0 with $\Delta = x_1(t_0)$ as long as $x_1(t_i) \geq \Delta/2$ for the first $i \leq c_1 \cdot n$ many 1-productive interactions in $[t_0, t_0+\tau]$ the evolution of $x_1(t_i)$ can be viewed as a biased random walk on the line starting at Δ where a "right step" happens with probability p and "left step" with probability $1-p$, otherwise. The correctness follows from a standard coupling argument between two biased coins. Formally let $T_{min} = \inf\{t' \geq t_0 \mid x_1(t') = x_1(t_0)/2\}$ and $T_{max} = \inf\{t' \geq t_0 \mid x_1(t') = 2x_1(t_0)\}$.

First we bound $\Pr[T_{max} > T_{min}]$. It follows from Lemma A.18 the probability of ever having an excess of $\Delta/4$ "left steps" to "right steps" is at most

$$\left(\frac{1-p}{p}\right)^{\Delta/4} = \left(\frac{55n-\Delta}{55n+\Delta}\right)^{\Delta/4} = \left(1 - \frac{2\Delta}{55n+\Delta}\right)^{\Delta/4} \leq e^{-\frac{\Delta^2}{2(55n+\Delta)}} \leq n^{-5}$$

where we use $\Delta \geq x_1(t_0)/2$.

Next we bound $\Pr[T_1 > c_1 \cdot n]$. Now consider $\tau = c_1 \cdot n$ independent Poisson trials ($S_i \in \{-1, 1\}$ for all $i \leq c_1 \cdot n$) each with success probability $p = 1/2 + x_1(t_0)/(36n)$. Let $S = \sum_{i=1}^{c_1 \cdot n} S_i$. Using the Hoeffding bound (Theorem A.6) for $\lambda = \Delta$ we get

$$\begin{aligned}
\Pr[T_1 > c_1\cdot n] &\leq \Pr[S < \Delta] \\
&= \Pr[S - \mathbb{E}[S] < \Delta - \mathbb{E}[S]] \\
&\leq \Pr[|S - \mathbb{E}[S]| > \mathbb{E}[S] - \Delta] \\
&\leq 2\cdot e^{-\frac{2\Delta^2}{4\cdot c_1\cdot n}} \\
&\leq 2\cdot e^{-\frac{\Delta^2}{2\cdot c_1\cdot n}} \\
&\leq n^{-5}
\end{aligned}$$

Hence, the statement follows by the union bound over the high probability events from above. $\square$

We are now ready to show the lemma with these two auxiliary claims. We start with a brief overview of the proof. The proof is similar to the proof of Lemma 3.16, but we only have to consider the analog to Case (a). We repeatedly apply Claim 3.24 to Opinion 1. Then the support of the largest opinion, $x_1(t)$ doubles every $\mathrm{O}\big(n^2/x_1(t)\big)$ interactions until its support becomes larger than $2n/3$. After doubling at most $\log n$ times, we reach a configuration where $x_1(t) \geq 2n/3$. This will be our time T_4.

To show that there exists a t with $x_1(t) \geq 2n/3$, we use Observation 3.23 to utilize the stopped process $(\hat{\mathbf{X}}(t))_t$.

To track the progress of Opinion 1, we divide the interactions from $[T_3 + t_0, T_3 + t_0 + c \cdot n^2/x_1(T_3)]$ into subphases of varying length. Let $T_{(0)} = T_3 + t_0$ and define for $1 \leq \ell \leq \log n$

$$T_{(\ell)} = \inf \{ \, t \geq T_{(0)} \mid x_1(t) \geq 2^\ell \cdot x_1(T_{(0)}) \text{ or } x_1(t) \geq 2n/3 \, \} \, .$$

We call the interactions in the interval $\big[T_{(\ell-1)}, T_{(\ell)}\big)$ *subphase* ℓ. Note that by definition of $T_{(\ell)}$, the support of x_1 doubles in every subphase (or $x_1 \geq 2/3n$ and Phase 4 ends). In more detail, for a fixed but arbitrary subphase ℓ it follows from Claim 3.24 that the length of subphase ℓ is at most $c \cdot n^2/x(T_{(\ell-1)}) \leq c \cdot n^2/(2^{\ell-1} \cdot x_1(T_{(0)}))$, w.h.p. Hence, it follows that there exists a time $t' \in [T_{(\ell-1)}, T_{(\ell-1)} + c \cdot n^2/x(T_{(\ell-1)})]$ such that $x_1(t') \geq 2^\ell \cdot x_1(T_{(0)})$ or $x_1(t') \geq 2/3 \cdot n$, w.h.p. From the union bound over all subphases, we get that after at most $\log n$ subphases, there exists w.h.p. a time $t' \in [T_{(0)}, T_{(\log n)}]$ such that $x_1(t') \geq 2n/3$. This holds since otherwise $x_1(t') \geq 2^{\log n} \cdot x_1(T_{(0)}) \geq n \cdot \sqrt{n} \log^2 n > n.$, a contradiction.

Summing up the length of all subphases for $c = 111$ gives us

$$\sum_{i=1}^{\log n} \frac{c \cdot n^2}{2^{i-1} \cdot x_1(T_{(0)})} = \frac{c \cdot n^2}{x_1(T_{(0)})} \cdot \sum_{i=1}^{\log n} \frac{1}{2^{i-1}} \leq 2 \cdot c \cdot \frac{n^2}{x_1(T_{(0)})}$$

and hence, $T_4 - T_3 \leq 7n \ln n + 4 \cdot c \cdot \frac{n^2}{x_1(T_3)}$ as claimed. □

3.5. From Absolute Majority to Consensus (Phase 5)

Recall that T_4 is the end of Phase 4, and $\mathbf{X}(T_4)$ is a configuration where the support of the largest opinion, $x_{\max}(T_4)$, is at least $2n/3$. The fifth phase ends when all agents agree on the Opinion $\max(T_4)$. In the following we assume w.l.o.g. that $x_1(T_4) \geq x_2(T_4) \geq \ldots \geq x_k(T_4)$.

Lemma 3.25. *Assume that* $\mathbf{x}(T_4)$ *is a configuration with* $x_1(T_3) \geq (2/3) \cdot n$. *Let* $T_5 = \inf \{ \, t \geq T_4 \mid x_1(t) = n \, \}$. *Then*

$$\Pr [T_5 - T_4 \leq c \cdot n \log n] \geq 1 - n^{-3}.$$

Proof. The idea is to show that all but Opinion 1 lose their support, and then the remaining undecided agents eventually adopt the only remaining opinion. We start with the first part by defining the potential function $\Psi_(t) = x_i(t)/x_1(t)$ for a fixed opinion $i \neq 1$. We show that $\Psi(t)$ decreases in expectation and apply a known multiplicative drift theorem. We calculate the expected change of Ψ by considering all possible interactions.

$$\begin{aligned}
&\mathbb{E}\left[\Psi(t+1) - \Psi(t) \mid |\mathbf{X}(t) = \mathbf{x}\right] \\
&= \frac{x_1(n-u-x_1)}{n^2} \cdot \left(\frac{x_i}{x_1 - 1}\right) + \frac{x_i(n-u-x_i)}{n^2} \cdot \left(\frac{x_i - 1}{x_1}\right) \\
&\quad + \sum_{j \neq 1,i} \frac{x_j(n-u-x_j)}{n^2} \cdot \left(\frac{x_i}{x_1}\right) + \frac{ux_1}{n^2} \cdot \left(\frac{x_i}{x_1+1}\right) + \frac{ux_i}{n^2} \cdot \left(\frac{x_i+1}{x_1}\right) \\
&\quad + \sum_{j \neq 1,i} \frac{ux_j}{n^2} \cdot \left(\frac{x_i}{x_1}\right) + \frac{\sum x_j^2 + \sum x_j u + u^2}{n^2} \cdot \left(\frac{x_i}{x_1}\right) - \Psi(t) \\
&= \frac{\Psi(t)}{n^2} \cdot \left(\frac{x_1(n-u-x_1)}{x_1 - 1} + \frac{u}{x_1+1} - (n-u-x_i)\right) \\
&= -\frac{\Psi(t)}{n^2} \cdot \left((x_1 - x_i) - \frac{n-u-x_1}{x_1-1} - \frac{u}{x_1+1}\right) \\
&\leq -\frac{\Psi(t)(x_1 - x_i)}{2n^2} \leq -\frac{\Psi}{20n}
\end{aligned}$$

where we use $x_1 \geq (6/10) \cdot n$ (follows by Lemma 3.26). We now apply the multiplicative drift result (Theorem A.9) with $r = 3\ln n$, $s_0 = x_i(T_4)/x_1(T_4)$, $s_{\min} = (n-1)^{-1}$ and $\delta = n^{-1} \cdot c^{-1}$. Then, we get for $T_i = \inf\{\, t \geq T_4 \mid x_i(t) = 0 \,\}$

$$\Pr\left[T - T_4 > c' n \log n\right] \leq \Pr\left[T - T_4 > \left\lceil \frac{r + \ln(s_0/s_{\min})}{\delta} \right\rceil\right] \leq e^{-r} = n^{-3}$$

From the union bound, it follows w.h.p. that all but Opinion 1 vanishes. It remains to show that all undecided agents vanish as well by adopting the only remaining opinion in the population. This follows by a simple argument about the one-way epidemics in $O(n \log n)$ interactions w.h.p. $\square$

Lemma 3.26. *Let $c > 0$ be an arbitrary constant and define $T = c \cdot n^2 \log n / x_1(T_4)$. Then*

$$\Pr\left[\text{for all } t \in [T_4, T_4 + T] \colon x_1(t) \geq (6/10) \cdot n\right] \geq 1 - n^{-5}.$$

Proof. Let

$$\hat{T} = \inf\{\, t \geq T_4 \mid u(t) \notin [(n - x_{\max}(t'))/2 - 8 \cdot \sqrt{n \ln n}, n/2 + \Delta_u] \,\}$$

be a stopping time and let $(\hat{X}(t))_t$ denote the process with $\hat{X}(t) = X(t)$ for all $t \leq \hat{T}$ and $\hat{X}(t) = X(\hat{T})$ for $t > \hat{T}$. From Lemma 3.4 and Lemma 3.6 it follows $\hat{T} - t = \Omega\big(n^2/x_{\max}(t) \cdot \log n\big)$ w.h.p. Thus, $(\mathbf{X}(t))_t$ and $(\hat{\mathbf{X}}(t))_t$ behave the same between time t and $t + \mathrm{O}\big(n^2/x_{\max}(T_4) \cdot \log n\big)$. An interaction is 1-productive with probability

$$\frac{u \cdot x_1 + x_1 \cdot (n - u - x_1)}{n^2} = \frac{x_1 \cdot (n - x_1)}{n^2} \leq \frac{1}{4}$$

It follows from an application of Chernoff bounds that within a sequence of $c \cdot n \cdot \log n$ interactions, the number of 1-productive interactions is at most $(c/4) \cdot n \log n$ with probability at least $1 - n^{-10}$. Now consider $\tau = (c/4) \cdot n \log n$ such productive interactions and let Z_t denote the change w.r.t. $x_1(t)$, i.e., the support of the largest opinion increase or decrease by one, respectively. That is, assuming the next interaction is a 1-productive interaction for $x(t)$ we have

$$\begin{aligned}\Pr[Z_t = 1] &= \frac{u \cdot x_1}{u \cdot x_1 + x_1 \cdot (n - u - x_1)} = \frac{u}{(n - x_1)} \\ \Pr[Z_t = -1] &= 1 - \Pr[Z_t = 1]\end{aligned}$$

Therefore

$$\mathbb{E}[Z_t] = \frac{u - (n - u - x_1)}{n - x_1} = \frac{2 \cdot u + x_1 - n}{n - x_1} \geq -48 \cdot \frac{\sqrt{n \ln n}}{n}$$

Let Z be the sum of Z_t for all $t \in [1, \tau]$. Then it follows from Hoeffding bound with $\lambda = n/15 - 200 \cdot \sqrt{n} \ln^{3/2} n$

$$\Pr\left[Z < -\frac{n}{15}\right] \leq \Pr[Z < \mathbb{E}[Z] - \lambda] \leq e^{-\frac{2\lambda^2}{4\tau}} \leq n^{-10}$$

The statement follows by the union bound. $\square$

Chapter 4.

Exact Plurality Consensus

In this part, we propose and analyze protocols for the exact plurality consensus problem in the population protocol model. Recall that *plurality opinion* refers to the opinion with the initially largest support (assuming it is unique), and *bias* denotes the difference between that opinion's initial support and that of the second largest opinion. A major part of research seeks to identify this plurality opinion for *any* initial bias, even if it is only 1.

We present new population protocols for plurality consensus with a primary focus on space complexity. Commonly, exact consensus problems are preferably solved by stable protocols. The exact majority problem for $k = 2$ opinions has been solved stably time- and space-optimal by Doty et al. [47]. Unfortunately, there is a known $\Omega(k^2)$ lower bound on the state space by Natale and Ramezani [72]. To beat this quadratic lower bound, we allow our protocols to fail with negligible probability as an essential aspect.

Results and Methodology We design protocols that, w.h.p., identify the plurality opinion quickly and have an almost optimal space complexity, even if the initial bias is only 1 (hence we solve *exact* plurality consensus). With this goal, allowing a negligible failure probability is essential, as otherwise – independently of the runtime – any protocol requires $\Omega(k^2)$ states [72].

Our first protocol uses $O(k + \log n)$ states. It consists of $k - 1$ tournaments, during which a *defender* and *challenger* opinion compete. To be more precise, we utilize the optimal majority protocol of Doty et al. [47] in each tournament and update the champion and challenger opinions afterward. W.h.p., the plurality opinion emerges victorious from all tournaments in time $O(k \cdot \log n)$. This protocol relies on ordering the opinions to determine the next challenger opinion.

Our second protocol avoids the requirement of such an order by using instead a *leader election* subprotocol to determine the next challenger opinion. Using the leader election protocol of Gasieniec and Stachowiak [51] for this,[1] our protocol for unordered opinions still uses $O(k+\log n)$ states but has a slightly increased runtime of $O(k\cdot n\cdot\log n+n\cdot\log^2 n)$ interactions. Avoiding such an ordering might seem like an esoteric challenge by itself. Still, this approach plays a crucial role in our third protocol (see below), where it is used

[1] W.h.p., that protocol finishes in $O(\log^2 n)$ time, leading to the corresponding term in our increased runtime. While there is a $O(\log n)$ time leader election protocol [32], that runtime holds only in expectation, which is too weak for our purpose.

to perform tournaments only for a subset of a priori unknown (significant) opinions that remain after an initial pruning phase. The following theorem states the properties of our first two protocols.

Theorem 4.1. *Assume we have a population of size n with $k \leq n/40$ initial opinions.*

1. *If the opinions are numbered $1, \ldots, k$ then SimpleAlgorithmconverges w.h.p. to the plurality opinion in $\mathrm{O}\big(k \cdot n \cdot \log n\big)$ interactions using $\mathrm{O}\big(k + \log n\big)$ states.*

2. *If there is no order among the opinions, SimpleAlgorithmcan be modified to converge w.h.p. to the initial plurality opinion in $\mathrm{O}\big(k \cdot n \cdot \log n + n \cdot \log^2 n\big)$ interactions using $O(k + \log n)$ states.*

For constant values of k, the unmodified SimpleAlgorithmconverges w.h.p. in optimal $\mathrm{O}\big(\log n\big)$ parallel time and requires only $\mathrm{O}\big(\log n\big)$ states. This matches the state and time complexities of the state-of-the-art exact majority protocol [47]. Note that the protocol from [47] is stable but ours gives w.h.p. guarantees only.

Our main contribution is the third protocol, which uses a pruning process to remove *insignificant* opinions before the tournaments start, reducing their number from $k - 1$ to $n/x_{\max}$ (remember that $x_{\max}$ denotes the initial size of the plurality opinion).

The following theorem states the results for our final population protocol formally.

Theorem 4.2. *Assume we have a population of size n with k initial opinions where $x_{\max} > n^{1/2+\varepsilon}$ for some small constant $1/2 > \varepsilon > 0$. ImprovedAlgorithmconverges w.h.p. to the plurality opinion in $\mathrm{O}\big(n^2/x_{\max} \cdot \log n + n \cdot \log^2 n\big)$ interactions using $\mathrm{O}\big(k \cdot \log\log n + \log n\big)$ states.*

The idea of the pruning process is to have each *subpopulation of opinions* run through a few *preprocessing phases* controlled by their own dedicated *phase clock*. Phase clocks [1, 10, 28, 51] are a common tool in population protocols to synchronize agents into phases. We will show that larger subpopulations finish their preprocessing phase faster than smaller subpopulations.

This way, we can filter out smaller opinions during the preprocessing before starting the tournaments. Assuming the initial largest opinion is of order $n^{1/2+\Omega(1)}$, we solve exact plurality in $\mathrm{O}\big(n^2/x_{\max} \log n + \log^2 n\big)$ interactions using $\mathrm{O}\big(k \log\log n + \log n\big)$ states. Note that if $k < n^{1/2-\varepsilon}$, the requirement $x_{\max} > n^{1/2+\varepsilon}$ is always fulfilled (this follows from $x_{\max} \geq n/k$). The assumption on $x_{\max}$ is necessary to wrap up the preprocessing and initiate the tournament part. In addition to the slightly increased state space due to using phase clocks, we cover a broader range of initial configurations. Furthermore, the runtime is improved in many cases since it does not only depend on k anymore. For example, in Fig. 4.1, we observe the power of the pruning process in a skewed configuration. After the pruning, all but the two largest opinions (below the threshold) are eliminated and do not proceed towards the tournament phase.

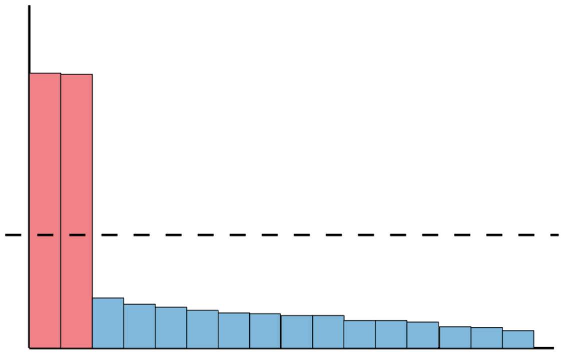

Figure 4.1.: Skewed configuration: Pruning Process

4.1. The Simple Algorithm

In this section, we present our first algorithm called SimpleAlgorithm, where each agent has one of k possible opinions numbered from 1 to k. The main idea of the protocol is as follows. It performs a sequence of *tournaments* of length $O(n \log n)$ synchronized by a *phase clock* [1]. In each tournament, two fixed opinions are chosen, and an exact majority protocol [47] is used to determine the majority opinion among the two of them. In the first tournament, opinions 1 and 2 compete. In tournament $i > 1$ the winner of tournament $i-1$ (called *defender*) competes against opinion $i+1$ (called *challenger*). The winner of tournament i has the largest support among the first $i+1$ opinions, and the last tournament winner is the plurality opinion.

To reach our state bound of $O(k + \log n)$, our protocol has to be very economical with the states. For example, an agent can't store two different opinions, which would already require $\Omega(k^2)$ states. Our protocol starts with an initialization phase, which splits the agents into four *roles*: `collector`, `player`, `clock`, and `tracker`. Every agent u has a variable $\mathsf{role}[u]$ to store its role in the protocol. The protocol consists of an initialization part (see Algorithm 4.3) and three different subprotocols specific to the corresponding roles.

We already argued that no agent can store two different opinions. Hence, the initialization phase is used to "collect" opinions: Initially, each agent is a `collector`-agent for its initial opinion. Each agent has a variable tokens, which can take on values between 1 and 10. The total number of tokens equals the number of agents initially supporting that opinion. When a `collector`-agent meets another agent with the same opinion, it increases the token counter accordingly. This frees up the other agent, which takes on a role in $\{$`clock`, `tracker`, `player`$\}$. During the tournament, the `collector`-agents are responsible for initiating the majority protocols between the actual challenger and defender. To this end, they have two Boolean variables $\mathsf{defender}$ and $\mathsf{challenger}$, which indicate that their opinion participates in the match as defender or challenger, respectively. Additionally, all `collector`-agents have a bit winner, which indicates the majority opinion of the last tournament. This bit is used to broadcast the final majority opinion.

$$\mathcal{S} = \underbrace{\begin{Bmatrix}\texttt{collector,}\\ \texttt{player,}\\ \texttt{clock,}\\ \texttt{tracker}\end{Bmatrix}}_{\texttt{role}} \times \underbrace{[-1;9]}_{\texttt{phase}} \times \overbrace{\underbrace{[\Theta(\log n)]}_{\texttt{count}}}^{\texttt{clock}} \times \overbrace{\underbrace{[k]}_{\texttt{tcnt}}}^{\texttt{tracker}} \times \overbrace{\underbrace{[k]}_{\texttt{opinion}} \times \underbrace{[10]}_{\texttt{tokens}} \times \underbrace{\{0,1\}^3}_{\texttt{defender, challenger, winner}} \times \underbrace{[-10;10]}_{\ell}}^{\texttt{collector}} \times \overbrace{\underbrace{\{A,B,U\}}_{\texttt{playeropinion}} \times \mathcal{S}_{\texttt{maj}}}^{\texttt{player}}$$

Figure 4.2.: State Space $\mathcal{S}$. Note that $[i] = \{1,\ldots,i\}$ and $[-i;j] = \{-i,\ldots,j\}$.

Finally, a value $\ell \in [-10, 10]$ cancels opposing opinions before a match.

Internally, the `clock` agents run the leaderless phase clock from [1] on a local counter `count` (see Section 4.2). Whenever the local counter passes through zero, the agent increases a variable `phase` modulo 10. The new value is disseminated to all other agents via one-way epidemics. The role of the `tracker`-agents is to store the number of the current challenger in a variable `tcnt` (short for *tournament counter*). Whenever one of the tournaments is over, this variable is increased by one. The *Collector*-agents use this to set the challenger bit at the beginning of a new tournament. The `player`-agents are the ones performing the $k-1$ tournaments. At the beginning of a tournament, these agents adopt the opinions from `collector`-agents, which have either the *defender* or *challenger* bit set and set their `playeropinion` to A or B, respectively.

Overview of the State Space We use $\mathcal{S}_{\texttt{maj}}$ to denote the set of states used by the exact majority protocol. Figure 4.2 gives an overview of the variables used by our protocol and how some can be attributed to the different roles. Note that $\mathcal{S}$ is *not* the actual state space used by our protocol. Our actual state space is much smaller since the role-specific variables must only be tracked by the corresponding roles. We describe this more thoroughly in the corresponding proof of the state complexity in Section 4.5.

Simplifications for the Pseudocode In our formal algorithms, we define how both involved agents (u, v) update their states in an interaction: u is the initiator, and v is the responder of that interaction. To simplify the exposition of our protocols, we allow the use of a "do once" statement in the pseudocode for state transitions that are to be executed *only once* in a given phase. For example, consider the scenario where the challenger wins the match. In the subsequent conclusion phase, all defender agents are removed, and all challenger agents set the `defender` bit. This must be done exactly once since. Otherwise, all bits are lost. See Line 27 to Line 29 in Algorithm 4.4 for the corresponding pseudocode using a “do once” statement. Similarly to the “do once” statements, we assume that agents can determine whether they interact for the “first time” in a phase. Note that these statements can be implemented using constantly many bits, such that the overall state space size increases only by a constant factor.

4.2. Clock and Tracker Agents

The `clock`-agents have two different tasks (see Algorithm 4.1). First, they decide when the initialization phase is over. They use their local counter count (initialized to zero) for that. Whenever they interact with a non-`collector`-agent, they increase count by one. If they interact with a `collector`-agent count is decreased by one as long as it is larger than zero. As soon as count reaches $5 \log n$, the agent decides that the initialization phase is over (constant fraction of non-`collector`–agents is reached) and sets phase $= 0$, which is then spread via broadcast (phase is initialized at the beginning of the whole protocol to -1). From there on, the `clock`-agents use count to run the leaderless phase clock from [1] for the synchronization, which works as follows. The counter count is used modulo $\Psi = \Theta(\log n)$. Whenever two `clock`-agents interact, the one with the lower counter value (w.r.t. the circular order modulo Ψ) increments its count. If both `clock`-agents have the same count value, ties are broken arbitrarily. When count $= 0$, the variable phase is increased by one (modulo 10). Alternatively to this simple clock, any phase clock that requires $\mathrm{O}\big(\log n\big)$ states can be used.

Algorithm 4.1: Clock Synchronization

Clock Synchronization. We assume that u is a `clock` agent.

if $\mathsf{phase}[u] = -1$ **then**
  $\mathsf{count}[u] \leftarrow$
    $$\begin{cases} \mathsf{count}[u] + 1 & \text{if } \mathsf{role}[v] \neq \texttt{collector} \\ \mathsf{count}[u] - 1 & \text{if } \mathsf{role}[v] = \texttt{collector} \wedge \mathsf{count}[u] > 0 \end{cases}$$
  if $\mathsf{count}[u] = 5 \cdot \log n$ **then**
    $\mathsf{phase}[u] \leftarrow 0$

if $\mathsf{phase}[u] \neq -1$ **and** $\mathsf{phase}[v] \neq -1$ **then**
  ▷ leaderless phase clock from [1]
  leaderless_phase_clock ($\mathsf{count}[u], \mathsf{count}[v]$)

if $\mathsf{count}[u]$ passes through zero **then**
  $\mathsf{phase}[u] \leftarrow \mathsf{phase}[u] + 1 \mod 10$

The `tracker`-agents determine which opinion has to take over the role as a challenger (see Algorithm 4.2). The state variable tcnt is initialized (see initialization phase) with 1 and incremented by one (modulo k) whenever phase switches over to zero. Note that tcnt $= 2$ during the first tournament. This holds due to the initialization of tcnt with one and the fact that it is incremented as soon as phase is incremented from -1 to 0 when the initialization phase ends.

Algorithm 4.2: Synchronization-Subroutine

We assume that u is a `tracker`–agent.

if $\mathsf{phase}[u] = 0$ **and** u interacts for the first time in this phase **then**
  $\mathsf{tcnt}[u] \leftarrow \mathsf{tcnt}[u] + 1$.

4.3. Initialization

The objective of this phase is to partition the population into four different roles: `collector`, `player`, `tracker`, and `clock`. Initially, every agent has the `collector` role, storing one token of its initial opinion. Whenever two `collector`-agents with the same opinion and at most 10 tokens in total interact, the responder sets its tokens variable to the sum of the tokens of both agents, and the initiator switches to a role in $\{\,\mathtt{clock}, \mathtt{tracker}, \mathtt{player}\,\}$ uniformly at random. Agents with opinion 1 set $\mathsf{defender} = 1$ during their first interactions. As soon as agent u becomes `clock`-agent, it uses the state variable count to determine when the initialization is over by setting $\mathsf{phase}[u]$ equals 0, which is then spread via broadcast. At this point, the first tournament starts with the `setup` phase.

Algorithm 4.3: Initialization Phase

We assume that u **and** v are initially in phase$[u]$ = phase$[v] = -1$.

if u is the initiator for the first time **and** $\mathsf{opinion}[u] = 1$ **then**
  $\mathsf{defender}[u] = \textsc{True}$

if $\mathsf{role}[u] = \mathsf{role}[v] = \mathtt{collector}$ **and** $\mathsf{opinion}[u] = \mathsf{opinion}[v]$ **and** $\mathsf{tokens}[u] + \mathsf{tokens}[v] \leq 10$ **then**
  $(\mathsf{tokens}[u], \mathsf{tokens}[v]) \leftarrow (0, \mathsf{tokens}[u] + \mathsf{tokens}[v])$

$$\text{with probability } 1/3\text{:} \begin{cases} (\mathsf{role}[u], \mathsf{count}[u]) & \leftarrow (\mathtt{clock}, 0) \\ (\mathsf{role}[u], \mathsf{tcnt}[u]) & \leftarrow (\mathtt{tracker}, 1) \\ (\mathsf{role}[u], \mathsf{playeropinion}[u]) & \leftarrow (\mathtt{player}, U) \end{cases}$$

if $\mathsf{phase}[v] = 0$ **then**
  $\mathsf{phase}[u] \leftarrow 0$

Lemma 4.3. *Let $\hat{t}$ denote the interaction, in which the first agent sets $\mathsf{phase} = 0$. Then, the following statements hold w.h.p.:*

1. $\hat{t} = \mathrm{O}\big(n \cdot (k + \log n)\big)$.
2. *After interaction $\hat{t}$ each of the roles `collector`, `clock`, `tracker`, and `player` are held by at least $n/10$ agents.*
3. *After interaction $\hat{t}$ all `collector`-agents of opinion 1 have their defender bit set.*

Proof. We consider a modified process, which mimics the original process. The only difference is that in this process, we prevent `clock`-agents from setting their phase to 0 by removing Line 6 of Algorithm 4.1. This causes all agents to remain in the `init` phase, i.e., they have phase set to -1 indefinitely. All agents perform according to Algorithm 4.3 in this setting. This simplifies the analysis as we do not have to deal with some agents that have already started the tournament. In the following, we assume that this modified process runs alongside the original process and that the same random choices are made in both processes. Let now $\tau_m(x)$ denote the first interaction in which at most $x \cdot n$ `collector`-agents remain in the modified process. Similarly, let $\hat{t}_m$ denote the first interaction in which some `clock`-agent counts to $5 \log n$. Additionally, we define the same notation with subscript o with respect to the original process. Observe that $\hat{t}_m = \hat{t}_o$ as until this interaction occurs, both processes are identical.

We start by establishing that, in the modified process, $\tau_m(1/3)$ is reached quickly. The number of remaining `collector`-agents decreases fast as long as all nodes follow Algorithm 4.3.

Claim 4.4. *It holds that* $\tau_m(1/3) = \mathrm{O}\big(n \cdot k\big)$ *w.h.p.*

Proof. To reach interaction $\tau_m(1/3)$ exactly $\lceil 2n/3 \rceil$ agents need to leave their `collector` role due to the token transfer in Line 8 of Algorithm 4.3. In the following, we say that an interaction is *good* if two agents interact that both are `collector`-agents, have the same opinion, and have at most 10 tokens. Such an interaction decreases the number of `collector`-agents by one. Let $z_i(t)$ denote the number `collector`-agents of opinion i, which have at most 5 tokens before interaction t is executed. If two such agents of the same opinion interact, the interaction is guaranteed good. Fix now some interaction $t < \tau(1/3)$, i.e., an interaction before which more than $n/3$ `collector`-agents are still present. Then, the probability for interaction t to be *good* is

$$\sum_{i=1}^{k} \frac{z_i(t)}{n} \cdot \frac{z_i(t)-1}{n-1} \geq \frac{1}{n^2} \sum_{i=1}^{k} z_i(t)^2 - \frac{1}{n^2} \sum_{i=1}^{k} z_i(t)$$
$$\stackrel{(a)}{\geq} \frac{1}{n^2} \frac{\left(\sum_{i=1}^{k} z_i(t)\right)^2}{k} - \frac{1}{n} \stackrel{(b)}{\geq} \frac{1}{n^2} \frac{n^2}{36 \cdot k} - \frac{1}{n} \stackrel{(c)}{\geq} \frac{1}{500k}$$

For the third inequality (b), we apply the following counting argument to bound $\sum_{i=1}^{k} z_i(t)$: only $n/6$ agents may have at least 6 tokens as the number of tokens sums to n at all times. We assume that at time t, there are still $n/3$ total `collector`-agents remaining. Hence, $\sum_{i=1}^{k} z_i(t) \geq n/3 - n/6 = n/6$. For the last inequality (c), we use that $k \leq n/40$ as assumed in Theorem 4.1. As each interaction is good with a probability of at least $1/500k$, independently, we consider a sequence of $500nk$ interactions and apply Chernoff's bounds. This yields that, w.h.p., there will be at least $\lceil 2n/3 \rceil$ good interaction in this sequence, reducing the number of `collector`-agents below $n/3$. In other words: $\tau(1/3) < 500nk$ w.h.p. ☐

In the following claim, we bound the time for the first `clock`-agent to count until $5 \log n$ in the modified process.

Claim 4.5. *It holds that* $\tau_m(2/3) < \hat{t}_m$ *and* $\hat{t}_m = \mathrm{O}\big(n \cdot (k + \log n)\big)$ *w.h.p.*

Proof. We consider the modified process and couple the counting procedures of any fixed `clock`-agent with a biased random walk on the non-negative line. The current value of the counter variable corresponds to the position of the walk on the line. Each time the `clock`-agent interacts as an initiator with a non-`collector`-agent, the random walk process moves to the right. Similarly, when interacting with a `collector`-agent the random walk moves to the left (or remains at 0 if its current position is 0). We are interested in the interactions required for the random walk to hit the value $5\log n$ as this corresponds to the `clock`-agent counting until $5\log n$. Until $\tau_m(2/3)$ is reached, this hitting time may be minorized with the hitting time of a random walk that has probability exactly $q = 2/3$ to move to the left and probability $p = 1/3$ to move to the right. Due to the strong drift towards 0, it is known that such a random walk takes poly(n) steps w.h.p. to hit $5\log n$. We utilize a variant of a known random walk result (Lemma A.17) to determine the constant hidden in poly(n). It implies that this hitting time is at least $n^{2.5}$ with probability at least $1-n^{-2.5}$. Therefore, w.h.p., the `clock`-agent will *not* reach a counter value of $5\log n$ before, either, $\tau_m(2/3)$ is reached or $n^{2.5}$ interactions have passed. Now, observe that $\tau_m(2/3) < \tau_m(1/3)$ as the number of `collector`-agents can only decrease over time. This implies by Claim 4.4 that $\tau_m(2/3) < n^{2.5}$ w.h.p. for large enough n. Hence, $\hat{t}_m(2/3)$ precedes $n^{2.5}$ w.h.p. and $\hat{t}_m > \tau_m(2/3)$ follows.

To show the upper bound on $\hat{t}_m$, we first argue that soon after $\tau_m(1/3)$ some `clock`-agents increases its counter to $5\log n$. We follow a similar approach and fix the modified process at some interaction $t \geq \tau_m(1/3)$ together with a `clock`-agent and its corresponding random walk. This time, we majorize the time for the counter to reach $5\log n$ with the hitting time of a random walk with $p = 2/3$ and $q = 1/3$. Such random walk is known (e.g., Theorem 18.2 of [65]) to have a hitting time of $\mathrm{O}\big(\log n\big)$ w.h.p. For convenience, we included a similar statement in Lemma A.17. Each movement of the random walk corresponds to one interaction as the initiator of the `clock`-agent. As the agent is selected as an initiator with probability $1/n$ in each interaction, it follows from a Chernoff bound that $\mathrm{O}\big(n\log n\big)$ interactions guarantee sufficient movements of the random walk w.h.p. Therefore, some `clock`-agents hits $5\log n$ before time $\tau_m(1/3) + \mathrm{O}\big(n\log n\big)$ w.h.p. From Claim 4.4 we know that $\tau_m(1/3) = \mathrm{O}\big(n\cdot k\big)$ w.h.p., allowing us to simplify this upper bound to $\mathrm{O}\big(n\cdot(k+\log n)\big)$. □

To show the first two statements of the lemma, we need the guarantees of Claim 4.5 in terms of the original process. Initially, we established that $\hat{t}_o = \hat{t}_m$ and that until this interaction, both processes act identically per definition. Additionally, note that $\Pr[\tau_o(2/3) = \tau_m(2/3)] \geq \Pr[\tau_m(2/3) \leq \hat{t}_m]$. This inequality holds because, if the event $\tau_m(2/3) \leq \hat{t}_m$ occurs, then both processes acted identically until interaction $\tau_m(2/3)$. Therefore, the amount of `collector`-agents is the same in both processes until this interaction, implying that $\tau_m(2/3) = \tau_o(2/3)$. By Claim 4.5 we have that $\tau_m(2/3) \leq \hat{t}_m$ w.h.p. and therefore $\tau_m(2/3) = \tau_o(2/3)$ is also a high probability event. Hence, w.h.p., Claim 4.5 also holds when exchanging $\hat{t}_m$ by $\hat{t}_o$ and $\tau_m(2/3)$ by $\tau_o(2/3)$, leading to the statement: $\tau_o(2/3) < \hat{t}_o = \mathrm{O}\big(n\cdot(k+\log n)\big)$ w.h.p. This inequality immediately yields the first statement of the lemma. We also use this inequality to show the second statement of the lemma as it implies that at $\hat{t}_o$ at most $2n/3$ `collector`-agents remain w.h.p. Therefore, at time $\hat{t}_o$, at least $n/3$ `collector`-agents must have left their role. Every agent that switches its role selects a new role uniformly and independently at random. Hence, it follows from Chernoff bounds that each non-`collector` role consists of at least

$(n/3) \cdot (1/3)(1 - o(1)) > n/10$ agents. Additionally, note that there must be at least $n/10$ collector-agents at all times. This follows since there are n tokens in total, and only collector-agents can hold up to 10 tokens each.

The proof for the final statement of the lemma is straightforward. It suffices to show that every agent interacts at least once before the first clock-agent sets phase to 0. Even if a clock-agent interacts with a non-collector-agent each time it is selected as initiator, it takes at least $5 \log n$ such interactions for it to set phase to 0. From Chernoff's bounds, it follows w.h.p. that it requires more than $2n \log n$ overall interactions for any clock-agent to be selected as initiator sufficiently many times. However, any fixed agent acts as an initiator at least once within $2n \log n$ interactions w.h.p. As each node is selected with probability $1/n$ as an initiator, the probability that an arbitrary but fixed agent is not selected is at most $(1 - 1/n)^{2n \log n} \leq \exp(-2 \log n) \leq n^{-2}$. A union bound over all agents shows that this is enough time for every agent to act as an initiator at least once w.h.p. □

Algorithm 4.4: Tournament Algorithm

if $\mathsf{phase}[u] = \mathsf{phase}[v] = 0$ **then** ▷ *Setup Phase*

  if $\mathsf{role}[u] = \texttt{collector}$ **and** $\mathsf{role}[v] = \texttt{tracker}$ **and** $\mathsf{opinion}[u] = \mathsf{tcnt}[v]$ **then**

    $\mathsf{challenger}[u] \leftarrow \textsc{True}$

  if $\mathsf{role}[u] = \texttt{collector}$ **then**

$$\ell[u] \leftarrow \begin{cases} \mathsf{tokens}[u] & \text{if } \mathsf{defender}[u] \\ -\mathsf{tokens}[u] & \text{if } \mathsf{challenger}[u] \\ 0 & \text{otherwise.} \end{cases}$$

if $\mathsf{phase}[u] = \mathsf{phase}[v] = 2$ **then** ▷ *Cancellation Phase*

  if $\mathsf{role}[u] = \mathsf{role}[v] = \texttt{collector}$ **then**

$$(\ell[u], \ell[v]) \leftarrow \left(\left\lfloor \frac{\ell[u] + \ell[v]}{2} \right\rfloor, \left\lceil \frac{\ell[u] + \ell[v]}{2} \right\rceil \right)$$

if $\mathsf{phase}[u] = \mathsf{phase}[v] = 4$ **then** ▷ *Lineup Phase*

  if $\mathsf{role}[u] = \texttt{collector}$ **and** $\mathsf{role}[v] = \texttt{player}$ **and** $\mathsf{playeropinion}[v] = U$ **then**

$$\mathsf{playeropinion}[v] \leftarrow \begin{cases} A & \text{if } \ell[u] > 0 \\ U & \text{if } \ell[u] = 0 \\ B & \text{if } \ell[u] < 0. \end{cases}$$

$$\ell[u] \leftarrow \operatorname{sign}(\ell[u]) \cdot (|\ell[u]| - 1)$$

if $\mathsf{phase}[u] = \mathsf{phase}[v] = 6$ **then** ▷ *Match Phase*

if $\mathsf{role}[u] = \mathsf{role}[v] = \texttt{player}$ **then**
execute $\texttt{majority}(\mathcal{S}_{\texttt{maj}})$ ▷ *execute the exact majority protocol from [47]*

if $\mathsf{phase}[u] = \mathsf{phase}[v] = 8$ **then** ▷ *Conclusion Phase*

if $\mathsf{role}[u] = \texttt{collector}$ **and** $\mathsf{role}[v] = \texttt{player}$ **and** $\mathsf{playeropinion}[v] = B$ **do** once
$\mathsf{defender}[u] \leftarrow \mathsf{challenger}[u]$
$\mathsf{challenger}[u] \leftarrow$ False

if $\mathsf{role}[u] = \texttt{collector}$ **and** $\mathsf{role}[v] = \texttt{player}$ **and** $\mathsf{playeropinion}[v] \in \{A, U\}$ **do** once
$\mathsf{challenger}[u] \leftarrow$ False

if $\mathsf{phase}[v] >_{(10)} \mathsf{phase}[u]$ **then**
$\mathsf{phase}[u] \leftarrow \mathsf{phase}[v]$

4.4. Player and Collector Agents

The tournaments are performed by both `player`- and `collector`-agents. Each tournament has five phases: `setup`, `cancellation`, `lineup`, `match`, and `conclusion`. To synchronize the beginning of the phases, we assume that there are phases (numbered with odd numbers) in which none of the `player`- and `collector`-agents are activated.

In the `setup` phase, `collector`-agents determine if their opinion is the challenger (in the i-th tournament, Opinion $i+1$ is the challenger opinion and $\mathsf{tcnt} = i$). Furthermore, all challenger and defender agents initialize a variable $\ell[u]$ with the (positive or negative) amount of tokens they store. In the `cancellation` phase, the agents use the load balancing protocol from [30, 69]. At the end of the protocol each agent u will have $\ell[u] \in \{\, \overline{\boldsymbol{\ell}} - 1, \overline{\boldsymbol{\ell}}, \overline{\boldsymbol{\ell}} + 1 \,\}$ where $\overline{\boldsymbol{\ell}}$ is the average of all the $\ell[u]$ values from challengers and defender agents rounded to the nearest integer. This phase reduces the number of tokens such that each token can be assigned to a different `player`-agent. This will be done in the `lineup` phase. The load balancing protocol can be used (see [47]) to calculate the majority opinion for the case of $k = 2$ and large bias. In that case, the majority opinion is the opinion for which a `collector`-agent exists with $\ell[u] \le -2$ or $\ell[u] \ge 2$. For the ease of presentation of our protocol, we do not distinguish between cases where the majority is already determined after this phase.

In the `match` phase, the `player`-agents, now having opinions A (defender opinion), B (challenger opinion), or U (undecided, held by `player`-agents which do not receive any opinion) determine the majority opinion using the majority protocol of [47]. We assume that the protocol returns the result in the state $\mathsf{playeropinion}$, which takes the values of the majority opinion.

Every `player`-agent u has a variable called $\mathsf{output}[u]$, which finally stores (unless there is a tie) the majority opinion. The initial opinion is stored in $\mathsf{input}[u]$. The protocol uses a variable $\mathsf{bias}[u]$ and sets $\mathsf{bias}[u] = +1$ if $\mathsf{input}[u] = A$ and $\mathsf{bias}[u] = -1$ if $\mathsf{input}[u] = B$.

In our protocol, we execute the exact majority protocol among the `player`-agents only. Hence, each `player`-agent needs the same set of states (additionally to the ones given in Section 4.1) as the exact majority protocol from [47]. SimpleAlgorithmnow is initialized as follows. A `player`-agent u with $\mathsf{playeropinion}[u] \neq U$ sets $\mathsf{input}[u] = \mathsf{playeropinion}[u]$. A `player`-agent u with $\mathsf{playeropinion}[u] = U$ sets $\mathsf{bias}[u] = 0$. With this initialization, the protocol determines the majority in time $\mathrm{O}(n \log n)$ since the number of `player`-agents is at least $\Omega(n)$. Note that in contrast to [47] we do not need the slow and always correct algorithm used since we are only interested in results that hold with high probability. We assume for every `player`-agent u that $\mathsf{playeropinion}[u]$ stores the output of the protocol.

In the `conclusion` phase, `collector`-agents holding the majority opinion set their defender bit. They have to participate in the next tournament. In Lines 34 and 35 agents broadcast phase to remain synchronized.

The following lemma provides an invariant about the tournament approach. For $1 \leq j < k, 0 \leq i \leq 9$, let $t_i(j)$ be the interaction in which the first agent enters phase i for the j-th time. Let ℓ_j be the plurality opinion ℓ among $1, \ldots j$.

Lemma 4.6. *Fix a j with $1 \leq j < k$ and assume that the first $j-1$ tournaments worked correctly. Then we have w.h.p.*

1. *At time $t_2(j)$ all `collector`-agents u with opinion $j+1$ have $\mathsf{challenger}[u] = \textsc{True}$. All other `collector`-agents have $\mathsf{challenger}[u] = \textsc{False}$. Furthermore, all `collector`-agents v not having opinion ℓ_j have $\mathsf{defender}[v] = \textsc{False}$.*
2. *Let $\mathcal{A}$ be the set of agents u with $\mathsf{playeropinion}[u] = A$ and let $\mathcal{B}$ be the set of agents with $\mathsf{playeropinion}[u] = B$. At time $t_6(j)$ we have $|\mathcal{A}| \geq |\mathcal{B}|$ iff $x_{\ell_j}(0) \geq x_{j+1}(0)$.*
3. *If $|\mathcal{A}| \geq |\mathcal{B}|$ ($|\mathcal{A}| < |\mathcal{B}|$) at time $t_8(j)$ we have $\mathsf{playeropinion}[u] \in \{A, U\}$ ($\mathsf{playeropinion}[u] = B$) for all `player`-agents u.*
4. *At time $t_0(j+1)$ all `collector`-agents u with opinion ℓ_{j+1} have $\mathsf{defender}[u] = \textsc{True}$.*

Proof. From Lemma 4.3, it follows that each role `collector`, `player`, `clock`, and `tracker` is held by at least $n/10$ agents, w.h.p. We denote the set of `player`-agents by P and the set of `collector`-agents by C. In the following, we prove the statements one after the other.

Statement (1) First we show that in Phase 0 of tournament j each agent interacts at least twice with a `tracker`-agent. w.h.p. we have at least $n/10$ `tracker`-agents; hence, the probability of interacting in a fixed step with a `tracker`-agent is at least $1/10$. The claim now follows from Chernoff's bounds.

Fix an agent u with opinion $j+1$. Since agent u interacts at least once with a `tracker`-agent in Phase 0, u sets $\mathsf{challenger}[u] = \textsc{True}$ in Line 3 of Algorithm 4.4. $\mathsf{defender}[u] = \textsc{False}$ follows from the initialization phase (see Algorithm 4.3). Now consider an agent u with opinion $\ell \notin \{\ell_j, j+1\}$. If $\ell > j+1$ $\mathsf{challenger}[u] = \mathsf{defender}[u] = \textsc{False}$ follows from the initialization phase (see Algorithm 4.3). Now assume that $\ell < j+1$. If opinion $\ell > 1$ the opinion was $\mathsf{challenger}$ opinion in tournament $\ell - 1$. If $\ell = 1$, the opinion was the defender in the first tournament. In either case, $\mathsf{challenger}[u]$ and $\mathsf{defender}[u]$ are set to $=$ FALSE in Line $17-21$ of Algorithm 4.4 or in Line 5 of Algorithm 4.3.

Statement (2) In this proof we assume w.l.o.g. that $x_{\ell_j}(0) \geq x_{j+1}(0)$. Fix a `collector`-agent u. From Statement (1) and Statement (4) of the previous tournament, it follows that in Line 5 of Algorithm 4.4 $\ell[u]$ is set to $\mathsf{tokens}[u]$ if u is a defender agent and to $-\mathsf{tokens}[u]$ if u is a challenger agent. In Line 8 of Algorithm 4.4, the defender and challenger agents perform a load balancing protocol for the rest of Phase 2. The protocol is analyzed in [30, 69]). We define $L = \sum_{u:\texttt{collector}} \ell[u]$ as the total load at the end of Phase 2 and $\hat{L} = \sum_{u:\texttt{collector}} |\ell[u]|$ as the total remaining load . From [30, 69] it follows that at the end of Phase 2 we have w.h.p. (a) $L = x_{\ell_j}(0) - x_{j+1}(0)$ and (b) for every `collector`-agent u it holds either $\ell[u] \in \{0, 1, 2\}$ if $x_{\ell_j}(0) - x_{j+1}(0) \geq |C|/2$ or $\ell[u] \in \{-1, 0, 1\}$, otherwise.

In Phase 2 of Algorithm 4.4 every `collector`-agent u recruits $|\ell[u]|$ many undecided `player`-agents v. If $\ell[u] > 0$ it sets $\mathsf{playeropinion}[v] = A$ and $\ell[u] = \ell[u] - 1$. If $\ell[u] < 0$ it sets $\mathsf{playeropinion}[v] = B$ and $\ell[u] = \ell[u] + 1$. For rest of the `player`-agents it remains $\mathsf{playeropinion}[v] = U$. This is done in Lines $10-12$ of Algorithm 4.4. It remains to show that each of these agents can recruit the sufficient amount of `player`-agents. We will show the following claim.

Claim 4.7. *Assume $|P|$ is the number of `player`-agents. Fix the configuration at the time of the end of Phase 2. W.h.p. we have either*

(i) $\hat{L} \leq |P|/2$ *, or*

(ii) for every `collector`-agent u we have $\ell[u] \in \{0, 1, 2\}$ and then there exists a `collector`-agent u with $\ell[u] > 0$.

Proof. Statement (i) follows directly for $x_{\ell_j} + x_{j+1} \leq |P|/2$. Hence, for the rest of the proof, we can assume that $\hat{L} > |P|/2$. (Note that (i) would immediately follow if $|P| \geq 2|C|$, which is, unfortunately, quite unlikely). From the analysis in [30], it follows that w.h.p. at least $|C|/4$ agents u have $\ell[u] = 0$ (this holds due to the length of the phase and the fact that in Line 8 of Algorithm 4.4 "$+1$"-s are canceled against "-1"-s). From Chernoff's bounds, it follows that w.h.p. $|P| \geq |C|/2$. Statement (ii) follows directly for $x_{\ell_j}(0) - x_{j+1}(0) \geq |C|/2$ and the fact that $L = x_{\ell_j}(0) - x_{j+1}(0)$. The claim follows from a union bound over both statements. □

At last it remains to show that Statement (2) follows by the claim and the fact that $L = x_{\ell_j}(0) - x_{j+1}(0)$. Assume Statement (i) holds. Chernoff bounds show that every `collector`-agent u is able to recruit $|\ell[u]|$ many `player`-agents in $\mathrm{O}(n \log n)$ interactions w.h.p.

Now assume Statement (ii) holds instead. That is, no `player`-agent u is able to sets $\mathsf{playeropinion}[u] = B$ in Line 11 of Algorithm 4.4 and hence, it is sufficient that at least a `collector`-agent u with $\ell[u] > 0$ is able to recruit a `player`-agent. Again, this follows from Chernoff bounds for $\mathrm{O}(n \log n)$ interactions w.h.p. Then Statement (2) follows from a union bound.

Statement (3) We execute the exact majority protocol from [47] among the `player`-agents in Phase 6. (The detailed explanation can be found at the beginning of Section 4.4.) Since the `player` size is at least $n/10$, Chernoff bounds provide enough meaningful interactions in $\Theta(n \log n)$ interactions w.h.p. Together with Statement (2), this implies the claim.

Statement (4) Similarly to the proof of Statement (1), we can argue that in Phase 8 of tournament j, each agent interacts at least twice with a `player`-agent. From Statement (3), it follows that every `player`-agent v has $\mathsf{playeropinion}[v] = A$ ($\mathsf{playeropinion}[v] = B$, respectively) if the defender (challenger, respectively) opinion won the majority protocol in Phase 6. Note that the j-th tournament competition is between opinion $j + 1$ and ℓ_j.

First, let us assume that for each `player`-agent v, we have $\mathsf{playeropinion}[v] = B$, i.e., the challenger opinion won. Consider `collector`-agent u. In Phase 8 *Algorithm* 4.4 sets (see Line 17-19) $\mathsf{defender}[u] = \mathsf{challenger}[u]$, i.e., every `collector`-agent u with the challenger opinion has $\mathsf{defender}[u] = \text{True}$, and afterwards Algorithm 4.4 sets $\mathsf{challenger}[u] = \text{False}$.

Now we assume that for each `player`-agent v we have $\mathsf{playeropinion}[v] \in \{A, U\}$. In that case, the defender opinion won the competition, and we have $\mathsf{defender}[u] = \text{True}$ for all `collector`-agents with the defender opinion, as before. □

4.5. Aftermath

This subsection briefly describes how our protocol finishes after the last tournament. That is, the agents still need to ensure that the ultimate defender – w.h.p. the initial plurality opinion – is disseminated to all agents.

The `tracker`-agents initiate a final broadcast. Recall that the `tracker`-agents have a variable tcnt that keeps track of the challenger in each tournament. Once this variable reaches $k+1$, all opinions have participated in a tournament, and those `collector`-agents with the defender bit set have w.h.p. the initial plurality opinion. Now when a `tracker`-agent u with $\mathsf{tcnt}[u] = k+1$ interacts with a `collector`-agent v with $\mathsf{defender}[v] = \text{True}$, the defender agent sets its *winner* bit $\mathsf{winner}[v] \leftarrow \text{True}$. This winner bit and the corresponding opinion is disseminated to all agents: any agent w for which $\mathsf{winner}[w] = \text{False}$ sets $(\mathsf{role}[w], \mathsf{opinion}[w], \mathsf{winner}[w])$ to $(\texttt{collector}, \mathsf{opinion}[v], \text{True})$ when it interacts with such a winner agent v (with $\mathsf{winner}[v] = \text{True}$).

4.6. Proof of Theorem 4.1

Next, we prove the runtime and the bound on the state space from the first statement in Theorem 4.1.

Proof: Correctness and Runtime for Statement (1) of Theorem 4.1.
We first apply Lemma 4.3. Then it holds that $t_0(1) = O(n \cdot (k + \log n))$ and the population is partitioned into the roles `player`, `tracker`, `clock` and `collector` where

each role consists of at least $n/10$ agents w.h.p. The `clock`-agents run the phase clock from [1] that provides synchronized phases of length $\Theta(n \log n)$ w.h.p. In particular, the separation between even phases is sufficiently large, i.e., it lasts longer than the time to broadcast a message via one-way epidemic (see [10]). Now, we do an induction over $k-1$ tournaments to show that opinion ℓ_k is the defender at the end of the tournaments. At the beginning of the first tournament, the time $t_0(1)$ Lemma 4.3 implies that opinion 1 is the initial defender w.h.p., i.e., $\ell_1 = 1$. The induction step from tournament j to $j+1$ follows by Lemma 4.6 w.h.p. Thus, the initial plurality opinion is the defender at the end of the last tournament w.h.p. At last, all agents agree on the unique defender opinion, followed by a final broadcast in $\mathrm{O}(n \log n)$ interactions w.h.p. Summing up over the initialization phase and all tournaments, SimpleAlgorithmrequires $\mathrm{O}(n \cdot k \cdot \log n)$ interactions. □

Proof: Space Complexity for Statement (1) of Theorem 4.1.
Figure 4.2 shows a *superset* $\mathcal{S}$ of our protocol's state space. Depending on their role, the agents only use a much smaller portion of $\mathcal{S}$ as described below.

Each agent's state space consists of a set of *shared* variables, which any agent keeps track of, and of *role-specific* variables, which only agents of that role keep track of. We use $\mathcal{S}_{\mathsf{shared}}$ to denote the state set represented by all shared variables and $\mathcal{S}_r$ to denote the variables required for role $r \in \{\, \texttt{clock}, \texttt{tracker}, \texttt{collector}, \texttt{player} \,\}$.

Note that $|\mathcal{S}_{\mathsf{shared}}| = \Theta(1)$. Indeed, the shared variables encompass the constant size role variable, the constant size phase variable, and the constantly many bits required for the do-once statements (see the overview of the state space at the beginning of Section 4.1). The gray box in Fig. 4.2 indicates the role-specific variables. Specifically:

- `clock`-agents use count variable ($\Theta(\log n)$ values).
- `tracker`-agents use the tcnt variable (k values).
- `collector`-agents use the opinion variable (k values), the tokens variable (10 values), the defender, challenger, winner bits, and the load balancing values ℓ (21 values).
- `player`-agents use the playeropinion variable (3 values) and $O(\log n)$ states for the majority protocol from [47].

The maximum number of states required by any agent is then calculated as

$$\begin{aligned} & |\mathcal{S}_{\mathsf{shared}}| \cdot \max\{ \quad \mathcal{S}_{\texttt{clock}}, \; \mathcal{S}_{\texttt{tracker}}, \qquad \mathcal{S}_{\texttt{collector}}, \qquad \mathcal{S}_{\texttt{player}} \,\} \\ = \;\; & \Theta(1) \cdot \max\{\, \Theta(\log n), \qquad k, \; k \cdot 10 \cdot 2^3 \cdot 21, \; 3 \cdot \mathrm{O}(\log n) \,\} \\ = \;\; & \Theta(k + \log n), \end{aligned}$$

Finishing the proof of the first protocol's state complexity. □

4.7. The Improved Algorithm

The goal in this section is to remove *insignificant* opinions before they even participate in the tournament. For now, let us assume that every agent u has a counter $c[u]$, which counts the number of interactions with the same opinion. As soon as the first counter reaches a fixed value $t \in O(\log n)$, the agent sets $\mathsf{phase}[u] = 0$, which triggers the beginning of the tournaments. Only agents with at least a $t/2$ counter will participate in the tournament. *Insignificant* opinions (those of support $x_i < x_{\max}/c_s$ for some constant $c_s > 1$) are effectively out of the race. This reduces the required tournaments to $\mathrm{O}\big(n/x_{\max}\big)$, improving the runtime. To show the correctness of this approach, it remains to show that w.h.p. every agent of the initial plurality opinion is among these remaining agents, while no agents of insignificant opinions participate in the tournament. The rest of the analysis follows along the lines of Statement (2) of Theorem 4.1. Unfortunately, this simple approach requires an additional counter per agent, which exceeds the state space bounds of Theorem 4.2.

Our main idea to save on states is to use phase clocks instead of the counters, one per opinion. In the following, we call interactions *meaningful* if an agent interacts with another agent of the same opinion. We split the agents into *subpopulations*; agents with opinion i belong to subpopulation i. Every subpopulation runs its own phase clock as follows. Every agent u has all states of the *junta-driven* phase clock (see [10, 28, 51]), which requires only $\mathrm{O}\big(\log\log n\big)$ states compared to the $\Theta(\log n)$ used by the simple counter. The clocks work as follows. First, in every subpopulation, so-called *junta* agents are selected in meaningful interactions. Then, the phase clock runs on a counter in meaningful interactions only. Note that phase clocks of large subpopulations run faster than phase clocks of small ones. Whenever a phase clock passes through 0, the agents increment phase, which is initialized to $-c$ (we assume that the value $c \in \mathbb{N}$ is a sufficiently large constant). Once $\mathsf{phase}[u]$ becomes 0 for some agent u, this value is broadcasted to all agents as before. All agents u for which $\mathsf{phase}[u]$ is still stuck at the initial value $\mathsf{phase}[u] = -c$ will not participate in any tournament. Instead, they change their role (with probability 1/3 each) from `collector` to `clock`, `tracker`, or `player`. Note that in contrast to before, an agent u does not immediately adopt a new role when it sets $\mathsf{tokens}[u] = 0$ in an interaction with another `collector`-agent (see Line 11 of Algorithm 4.5). Instead, agent u waits until $\mathsf{phase}[u] = 0$. Then, agent u adopts a new role iff. it either has no tokens ($\mathsf{tokens}[u] = 0$) or its $\mathsf{phase}[u] = -c$ (the latter implies that the clock of agent u did not pass through zero even once).

It is now easy to see that this results in a faster convergence time. Indeed, this follows from how SimpleAlgorithmselects the next challenger opinion if there is no order among the opinions (see description in Section 4.9): In a modified setup phase, a leader selects an opinion as challenger randomly from the `collector` agents who have not yet been defeated in a tournament (using a cascade of one-way epidemic processes on the way). Hence, if no `collector` agents are left for some of the opinions, there will not be a tournament involving that opinion. Thus, the total runtime will be reduced accordingly.

When the first agent reaches $\mathsf{phase}[u] = 0$, all agents proceed with the modified version of SimpleAlgorithm. We remark that it can happen that only $o(n)$ `collector`-agents remain after removing all insignificant opinions. In this case, the `cancellation` phase will not achieve a balanced state. However, all tokens will fit into the `player` agents nonetheless, as we will show in Statement (3) of Lemma 4.12 that there will be a constant fraction of agents for each role in { `clock`, `tracker`, `player` }.

While the overall approach sounds very easy, the crux lies in the analysis. First of all, we have to analyze the *speed* of the clocks running via meaningful interactions only (Lemma 4.9) Then, we have to show that all agents of the plurality opinion pass through 0 at least once, meaning they will participate in the tournament (Lemma 4.12). Finally, we have to show that all agents with insignificant opinions will not participate in any tournament, either because they did not finish the FormJuntaprotocol (Lemma 4.11) or because their phase clock runs too slow (Lemma 4.12).

Junta-Driven Phase Clock We use the phase clock implementation from [28], which starts by electing a junta. We select the junta in the same way but using meaningful interactions only. Each agent is equipped with a level variable, initially 0, and a bit, indicating whether the agent is still active. Agents progress through levels: They are initially active and remain active and increase their level as long as they interact (as initiators) with another agent on the same or higher level. They become inactive if they initiate an interaction with another agent on a lower level. Finally, agents become inactive if they hit the maximum level $\ell_{\max} = \lfloor \log\log n \rfloor - 3$. All agents that reach this maximum level form the junta and start the phase clock protocol.

Every agent is equipped with a phase counter $p[u]$ (initially 0) in the phase clock. Whenever a junta agent u initiates an interaction with an agent v it sets

$$p[u] = \max\{p[u], p[v] + 1\}$$

. If the initiating agent u is not a junta agent, then u sets $p[u] = \max\{p[u], p[v]\}$. For $i > 0$, we say that an agent u *passes through zero* for the i-th time if its phase counter $p[u]$ fulfills $\lfloor p[u]/m \rfloor \geq i$ for the first time ($m \in \mathbb{N}$ is a fitting large enough constant). Note that in [28], the same property is referred to as u reaching hour i for the first time.

Our protocol sets the maximum level to $\ell_{\max} = \lfloor \log\log n \rfloor - 2$. We show in the proof of Lemma 4.9 that this modified maximum level still allows the election of a junta w.h.p. as long as the subpopulation has size at least $\sqrt{n}$.

We denote by $\mathcal{S}_c$ the $\Theta(\log\log n)$ states required to execute the junta election and phase clock protocols. We assume that all agents are initially equipped with sufficiently many additional states to run this clock. As soon as an agent u sets $\mathsf{phase}[u]$ to 0; it may reuse these states. The following lemma states the properties of this phase clock.

Algorithm 4.5: Modified Initialization

```
Modified Initialization. We assume that phase[u] < 0.

  if opinion[u] = opinion[v] and phase[v] < 0 then
      form_junta_protocol (S_c) ▷ execute the junta-election protocol from [28]
```

loglog_phase_clock ($\mathcal{S}_c$) ▷ *execute the phase clock protocol from [28]*

  if phase clock of u passes through zero **then**
    $\mathsf{phase}[u] \leftarrow \mathsf{phase}[u] + 1$

  if $\mathsf{tokens}[u] + \mathsf{tokens}[v] \leq 10$ **then**
    $(\mathsf{tokens}[u], \mathsf{tokens}[v]) \leftarrow (0, \mathsf{tokens}[u] + \mathsf{tokens}[v])$

if $\mathsf{phase}[u] = 0$ **or** $\mathsf{phase}[v] = 0$ **then**
  if $\mathsf{phase}[u] = -c$ **or** $\mathsf{tokens}[u] = 0$
    with probability 1/3: $\begin{cases} (\mathsf{role}[u], \mathsf{count}[u]) & \leftarrow (\mathtt{clock}, 0) \\ (\mathsf{role}[u], \mathsf{tcnt}[u]) & \leftarrow (\mathtt{tracker}, 1) \\ (\mathsf{role}[u], \mathsf{playeropinion}[u]) & \leftarrow (\mathtt{player}, U) \end{cases}$
  $\mathsf{phase}[u] \leftarrow 0$

Lemma 4.8. *Assume that we run the junta-election process and phase clock from [28] on a population of n agents. Let $s(0)$ ($e(0)$, resp.) be the interaction when the first (last, resp.) junta agent is elected, and let $s(i)$ ($e(i)$, resp.) be the interaction when the first (last, resp.) agent passes through zero for the i-th time. Then, for any constant $a > 0$, there exist two properly chosen constants c_1' and c_2', such that we have with probability at least $1 - n^{-a}$,*

1. *The protocol elects a non-empty junta of size at most $n^{0.98}$.*
2. $s(0) \leq c_2' \cdot n \log(n)$.
3. $c_1' \cdot n \log n \leq s(i+1) - s(i) \leq c_2' \cdot n \log n$ *for any $i = O(\text{poly}(n))$,*
4. $s(i+1) > e(i)$ *for any $i = O(\text{poly}(n))$.*

Proof. Follows from Theorem 1 and Lemma 6 in [28]. □

We denote by $s_j(0)$ ($e_j(0)$) the interaction at which the first (last, respectively) junta agent is elected in subpopulation j. For $i > 0$, we denote by $s_j(i)$ ($e_j(i)$) the time when the first (last, respectively) agent of opinion j passes through zero for the i-th time. The following lemma adjusts the results of Lemma 4.8 to subpopulations.

Lemma 4.9. *Fix a subpopulation j and assume that $x_j \geq n^{1/2}$. Consider the phase clock driven by subpopulation j. Then, for any constant $a > 0$, there exist constants $c_1 \leq c_2 \in \mathbb{N}$ such that the following statements hold with probability $1 - {x_j}^{-a}$.*

1. *Subpopulation j elects a non-empty junta with at most $(x_j)^{0.98}$ agents.*
2. $s_j(0) \leq c_2 \cdot \frac{n^2}{x_j} \log(n)$.
3. $c_1 \cdot \frac{n^2}{x_j} \log(n) \leq s_j(i+1) - s_j(i) \leq c_2 \cdot \frac{n^2}{x_j} \log(n)$ *for any $i = O(\text{poly}(n))$.*

4. $s_j(i+1) > e_j(i)$ *for any* $i = O(\text{poly}(n))$.

Proof. The statements of this lemma would directly follow from Lemma 4.8 by replacing n with x_j. However, the junta-election mentioned in Lemma 4.8 assumes that a maximum level $\lfloor \log\log x_j \rfloor - 3$ is set. Our agents do not know the value x_j, so we set this level to $\ell_{\max} = \lfloor \log\log n \rfloor - 2$ instead. With the following claim, we show that this modification still leads to a junta of desired size if $x_j \geq \sqrt{n}$

Claim 4.10. *If* $x_j \geq \sqrt{n}$ *then the FormJuntaprotocol [28] configured with maximum level* $\ell_{\max} = \lfloor \log\log n \rfloor - 2$ *elects a non-empty junta of* $\leq x_j^{0.98}$ *agents within* $\mathrm{O}\big(x_j \log(x_j)\big)$ *meaningful interactions and with probability at least* $1 - x_j^{-a}$ *(for any constant* $a > 0$*).*

Proof. We start by showing the bounds on the junta size. Depending on the size $x_j \geq \sqrt{n}$ of the subpopulation j, we can express $\ell_{\max}$ as either (i) $\lfloor \log\log x_j \rfloor - 3$, or (ii) $\lfloor \log\log x_j \rfloor - 2$. Consider the first case. In this case, $\ell_{\max}$ matches the maximum level in the specification of the FormJunta[28] protocol for populations of size x_j. Therefore, we can apply the corresponding Theorem 1, which states that a non-empty junta of size $\leq x_j^{0.98}$ is formed with probability $1 - x_j^{-a}$ (for any constant $a > 0$).

In the other case $\ell_{\max} = \lfloor \log\log x_j \rfloor - 2$. Throughout the FormJuntaprocess, only active agents may modify their level. That is, if an active agent u initiates a meaningful interaction with a node v, then (i) it becomes inactive if v has a level lower than u, or (ii) it remains active otherwise.[2] Just as in [28], we denote by B_ℓ the number of agents that reach at least level i. Per definition, it must hold that $B_\ell \geq B_{\ell+1}$ for any level $\ell \geq 0$. First, we show that between 1 and $x_j^{0.98}$ agents make it to level $\ell_{\max}$ with probability $1 - x_j^{-a}$ (for any constant $a > 0$). The upper bound on this number follows directly from Lemma 5 of [28]. It states that $B_{\lfloor \log\log x_j \rfloor - 3} < x_j^{0.98}$ with probability $1 - x_j^{-a}$ (again for arbitrary constants $a > 0$). Due to the monotonicity of B_ℓ, it follows that $B_{\ell_{\max}} < x_j^{0.98}$ as well. To show the lower-bound $B_{\ell_{\max}} > 1$, we would like to use Lemma 4 of [28]. Unfortunately, it only yields that $B_{\ell_{\max}-1} > 1$. Fortunately, in the proof of Lemma 4, they show the slightly stronger statement of $B_{\ell_{\max}-1} > x_j^{2/3}$ with probability at least $1 - x_j^{-a}$. We argue that this implies that $B_{\ell_{\max}} > 1$ with probability $1 - x_j^{-a}$. To show this, we rely on the coupling idea described in Footnote 6 on page 100 of [28]. That is, we serialize the points in time $\{t(l)\}_{l=1}^{x_j^{2/3}}$ at which the first $x_j^{2/3}$ agents that entered level $\ell_{\max} - 1$ make their first interaction as an initiator. At time $t(l)$, the l-th such agent decides whether it stays active and progresses to level $\ell_{\max}$ or becomes inactive (according to (i) and (ii) above). Observe that for any such agent that makes its decision after $t(x_j^{2/3}/2)$, the probability of remaining active is at least $x_j^{2/3}/(2x_j)$ (as at this time already $x_j^{2/3}/2$ agents entered level $\ell_{\max} - 1$). Hence, in expectation, at least $x_j^{2/3}/(2x_j) \cdot x_j^{2/3} = x_j^{1/3}/2$ agents progress to level $\ell_{\max}$. From Chernoff bounds it follows that at least $x_j^{1/3}(1 - o(1))/2$ agents reach $\ell_{\max}$ with probability $1 - x_j^{-\omega(1)}$.

[2] Note that the state transitions for agents on the first level 0 are slightly different but not relevant for this proof.

It remains to show that $\mathrm{O}\big(x_j \log(x_j)\big)$ meaningful interactions suffice for the first agent to reach level $\ell_{\max}$. This follows from Lemma 3 of [28]. There, it is shown that even if the maximum level is unbounded, all nodes become inactive within $\mathrm{O}\big(x_j \log(x_j)\big)$ interactions and with probability at least $1 - x_i^{-a}$. ☐

Statement (1) now follows directly from this claim. For Statement (2), we also refer to this claim and note that the junta election is driven by the subpopulation. Hence, the $\mathrm{O}\big(x_j \log(x_j)\big)$ *meaningful interactions* need to be converted into global interactions. To that end, observe that $(n^2/x_j) \cdot (1 + o(1))$ global interactions suffice for x_j meaningful interactions to occur with probability $1 - x_j^{-\omega(1)}$. Because the probability for any fixed interaction to be meaningful is x_j^2/n^2, this immediately follows from Chernoff's bounds. A symmetric approach also yields that at least $(n^2/x_j) \cdot (1 - o(1))$ global interactions are required for x_j many meaningful interactions to occur. This implies that $s_j(0) = \mathrm{O}\big((n^2/x_j) \cdot \log x_j\big)$. Due to the constraint on x_j, it holds that $\log(n) \geq \log(x_j) \geq \log(n)/2$ and Statement (2) follows.

The proof of Statement (3) follows from Statement (3) of Lemma 4.8 and a conversion to global interactions. Additionally, observe that due to the constraint on x_j, we have $\mathrm{poly}(n) = \mathrm{poly}(x_j)$ and note that the constant hidden in the exponent of $\mathrm{poly}(n)$ in Lemma 4.8 can be made arbitrarily large. The proof of Statement (4) again directly follows from Lemma 4.8 together with the above observation of $\mathrm{poly}(n) = \mathrm{poly}(x_j)$. ☐

Lemma 4.11. *Fix a subpopulation j of $x_j \leq \sqrt{n}$ agents. Let $\varepsilon > 0$ be an arbitrary small constant. Then, subpopulation j will not elect a junta agent before interaction $n^{1.5-\varepsilon}$ with probability $1 - n^{-\omega(1)}$.*

Proof. To join the junta, agents must increase their level from 0 to $\ell_{\max} = \lfloor \log\log n \rfloor - 2$. Per the definition of the junta election [28], an agent u may only increase its level if it interacts as an initiator (and some additional conditions hold). Furthermore, this increase is, at most, an increment of 1. Therefore, any fixed agent u of subpopulation j requires at least $\ell_{\max}$ meaningful interactions as an initiator to join the junta. We call such an interaction *bad* in the following. The probability that any fixed interaction is bad is $(1/n) \cdot x_j/n \leq n^{-1.5}$. Let $\varepsilon > 0$ be an arbitrary small constant. We show that in a sequence of $n^{1.5-\varepsilon}$ there will be less than $\ell_{\max}$ bad interactions with probability $1 - n^{-\omega(1)}$. The lemma's statement follows from a union bound over all agents in subpopulation i.

The number of bad interactions of u in this may be majorized by $\mathrm{Bin}(n^{1.5-\varepsilon}, n^{-1.5})$. It holds that

$$\begin{aligned}
&\Pr\left[\mathrm{Bin}(n^{1.5-\varepsilon}, n^{-1.5}) \geq \ell_{\max}\right] \\
&= \sum_{i=0}^{n^{1.5-\varepsilon}-\ell_{\max}} \Pr\left[\mathrm{Bin}(n^{1.5-\varepsilon}, n^{-1.5}) = \ell_{\max} + i\right] \\
&\overset{(a)}{\leq} n^{1.5-\varepsilon} \cdot \Pr\left[\mathrm{Bin}(n^{1.5-\varepsilon}, n^{-1.5}) = \ell_{\max}\right]
\end{aligned}$$

In step (a), we use that $\ell_{\max}$ is much larger than the expected value of this distribution. Hence, the terms in the sum decline with further i. This allows us to upper-bound each term in the sum by $p = \Pr\left[\mathrm{Bin}(n^{1.5-\varepsilon}, n^{-1.5}) = \ell_{\max}\right]$. Using the PDF of the binomial distribution, we can further bound p.

$$\begin{aligned} p &= \binom{n^{1.5-\varepsilon}}{\ell_{\max}} \cdot (n^{-1.5})^{\ell_{\max}} \cdot (1 - n^{-1.5})^{n^{1.5-\varepsilon} - \ell_{\max}} \\ &\leq \left(\frac{e \cdot n^{1.5-\varepsilon}}{\ell_{\max}}\right)^{\ell_{\max}} (n^{-1.5})^{\ell_{\max}} = \left(\frac{e}{\ell_{\max}}\right)^{\ell_{\max}} \cdot \frac{1}{n^{\varepsilon \cdot \ell_{\max}}} \end{aligned}$$

Since $\ell_{\max} = \Theta(\log\log n)$, this implies that $p = n^{-\omega(1)}$. Hence,

$$\Pr\left[\mathrm{Bin}(n^{1.5-\varepsilon}, n^{-1.5}) \geq \ell_{\max}\right] = n^{1.5-\varepsilon} \cdot n^{-\omega(1)} = n^{-\omega(1)}$$

for sufficiently large n and the result follows. □

In the following we define $T_i(t)$ as the total number of tokens for opinion i at interaction t, i.e.,

$$T_i(t) := \sum_{\{\, u \,\mid\, \mathsf{opinion}[u](t) = i \,\}} \mathsf{tokens}[u](t)$$

where $\mathsf{opinion}[u](t)$ and $\mathsf{tokens}[u](t)$ denote the values of the variables $\mathsf{opinion}[u]$ and $\mathsf{tokens}[u]$, respectively, in interaction t. Note that $T_i(0)$ is the initial support of opinion i.

Lemma 4.12. *Assume that $x_{\max} > n^{1/2+\varepsilon}$ for a small constant $\varepsilon > 0$. Let i be the initial plurality opinion and $\hat{t}$ denote the first interaction in which $\mathsf{phase} = 0$ for all agents. Then, w.h.p., $\hat{t} = \Theta((n^2/x_{\max}) \cdot \log n)$ and the following holds after interaction $\hat{t}$ w.h.p.:*

1. *There are at most $\mathrm{O}(n/x_{\max})$ distinct opinions left.*
2. *For the initial plurality opinion i it holds that $T_i(\hat{t}) = T_i(0)$.*
3. *Each of the roles `clock`, `tracker`, and `player` is held by at least $n/10$ agents.*

Proof. We first show the bound on $\hat{t}$. Recall that $s_i(0)$ is defined as the interaction when the first junta agent in subpopulation i is elected, and $s_i(c)$ is defined as the interaction when the clock of the first agent of opinion i ticks for the c-th time. We will prove upper and lower bounds for $\hat{t}$ based on $s_i(c)$.

From Statements (1) and (3) of Lemma 4.9 (with $a = 4$) it follows that $s_i(c) \leq (c+1)c_2 \cdot \frac{n^2}{x_{\max}} \log n$ with probability at least $1 - (1+c) \cdot x_{\max}^{-4} \geq 1 - (1+c) \cdot n^{-2-4\varepsilon}$ (since we assume that $x_i \geq n^{1/2+\varepsilon}$). Once an agent u has reached $\mathsf{phase}[u] = 0$, this phase value is disseminated to all other agents via one-way epidemics. It follows that, w.h.p., $\hat{t} \leq s_i(c) + \tau_{BC}$, where τ_{BC} is the *broadcast time* with $\tau_{BC} \leq c_2 \cdot n^2/x_{\max} \log n$ w.h.p. [10]. Ultimately, $\hat{t} \leq c_2 \cdot (c+2) \cdot n^2/x_{\max} \cdot \log n$ w.h.p.

For the lower bound, we observe that $\hat{t} \geq s_j(c) \geq c \cdot c_1 \cdot n^2/x_j \cdot \log n \geq c \cdot c_1 \cdot n^2/x_{\max} \cdot \log n$ with probability at least $1 - c \cdot n^{-2-4\varepsilon}$. A union bound over all opinions yields $\hat{t} \geq c \cdot c_1 \cdot n^2/x_{\max} \cdot \log n$ w.h.p. Together with the upper bound, the result for $\hat{t}$ follows. Next, we show the three statements individually.

Statement (1) Let $c^* = (c+2) \cdot c_2$ be the constant from the upper bound on $\hat{t}$ and define $c_s = c^*/c_1$. In the following, we show that any insignificant opinion j vanishes. For this, let j be an arbitrary but fixed opinion with $x_j < x_{\max}/c_s$. We distinguish two cases.

Case 1: $x_j \geq \sqrt{n}$ From Lemma 4.9 we get that w.h.p.

$$\hat{t} \leq c^* \cdot \frac{n^2}{x_{\max}} \log n \quad , \qquad s_j(1) \geq c_1 \cdot \frac{n^2}{x_j} \log n > c^* \cdot \frac{n^2}{x_{\max}} \log n,$$

where the last inequality uses the definition of c_s and $x_j < x_{\max}/c_s$. Together with a union bound, this implies that w.h.p. the clocks of all agents of opinion j do not tick even once. Hence, opinion j vanishes at the latest interaction $\hat{t}$ w.h.p.

Case 2: $x_j < \sqrt{n}$ Similarly to before, we get from above bounds on $\hat{t}$ and from Lemma 4.11 that w.h.p.

$$\hat{t} = \mathrm{O}\big(n^{3/2-\varepsilon} \cdot \log n\big) \qquad \text{and} \qquad s_j(1) \geq s_j(0) \geq n^{3/2-\varepsilon'}.$$

With $\varepsilon' < \varepsilon$ and a union bound, this again implies that w.h.p. the clocks of all agents of opinion j do not tick even once. Hence, in this case, opinion j vanishes at the latest in interaction $\hat{t}$ w.h.p.

The two cases show that any opinion j with $x_j < x_{\max}/c_s$ w.h.p. does not compete in the tournaments. Since we have n agents, at most $n \cdot c_s/x_{\max} = \mathrm{O}\big(n/x_{\max}\big)$ opinions remain after $\hat{t}$ interactions w.h.p.

Statement (2) To show the statement, we need to show that the clocks of any agent of the initial plurality opinion i pass through zero at least once before the first agent u hits $\mathsf{phase}[u] = 0$. Recall that $s_j(c)$ is the interaction when the clock of any agent with opinion j passes through zero for the c-th time (this is the first interaction when any agent u sets $\mathsf{phase}[u] = 0$).

We only consider opinions j with $x_j = \Omega(x_i)$ in the following. The statement for smaller opinions follows from the above proof of Statement (1). There, we have shown that the clocks of agents of smaller opinions will not pass through zero even once before interaction $\hat{t}$. For significant opinions j, we observe $s_j(c) \geq c \cdot c_1 \cdot n^2/x_{\max}$ w.h.p. as shown at the beginning of the proof. From the bound on $s_j(c)$ and from Lemma 4.9 we get that w.h.p.

$$s_j(c) \geq c \cdot c_1 \cdot \frac{n^2}{x_{\max}} \log n \qquad \text{and} \qquad e_i(1) \leq s_i(2) \leq 3c_2 \cdot \frac{n^2}{x_{\max}} \log n.$$

By choosing a sufficiently large constant $c > 3c_2/c_1$ in Algorithm 4.5, this yields $s_j(c) > s_i(2)$ w.h.p. In other words, when the first agent's clock has passed through zero for the

c-th time, all agents of opinion i have passed through zero at least once. In particular, $\mathsf{phase}[u] \neq -c$ for any agent u with opinion i in that interaction.

The total number of tokens $T_i(t)$ of opinion i can only change in some interaction t if an agent u of opinion i adopts another role in Line 15 of Algorithm 4.5 while $\mathsf{tokens}[u] > 0$. However, we have just shown that when the first agent v sets $\mathsf{phase}[v] = 0$, any agent u of opinion i has $\mathsf{phase}[u] \neq -c$. Hence, it follows that agent u can adopt a different role in Line 15 of Algorithm 4.5 only if agent u had $\mathsf{tokens}[u] = 0$ in Line 14 of Algorithm 4.5. Therefore, such an interaction does not change the total number of tokens for opinion i; the statement follows.

Statement (3) The proof follows from similar arguments as the proof of Statement (2) of Lemma 4.3. □

4.8. Proof of Theorem 4.2.

We split the proof of Theorem 4.2 into three parts: the proof of the correctness of the result, the proof of the runtime, and the proof of the state space requirements. Essentially, the theorem follows from Lemma 4.12 for the correctness of the modified initialization phase (Algorithm 4.5) and from Statement (2) of Theorem 4.1 for the correctness of SimpleAlgorithm.

Proof: Correctness of Theorem 4.2.
In ImprovedAlgorithm, all agents start with the modified initialization phase defined in Algorithm 4.5. After this initialization, they execute the tournament according to the variant of SimpleAlgorithm, which does not need an order among the opinions (see Section 4.9). By Statement (3) of Lemma 4.12, we get that all roles in $\{\,\mathtt{clock}, \mathtt{tracker}, \mathtt{player}\,\}$ are held by at least a constant fraction of agents at time $\hat{t}$. The number of agents with role `collector` may be asymptotically much smaller. However, their number does not affect the outcome of SimpleAlgorithm. Statement (2) of Lemma 4.12 guarantees that the initial plurality still has all of its initial tokens at the beginning of the tournaments. It follows along the lines of the proof of Statement (1) of Theorem 4.1 that this opinion will be the defender at the end of the tournament, and all agents will output this opinion after the final broadcast as described in Section 4.5. □

Proof: Runtime of Theorem 4.2.
From Lemma 4.12 we get that, w.h.p., after $\hat{t} = \mathrm{O}\big(n^2/x_{\max} \cdot \log n\big)$ interactions all agents u have $\mathsf{phase}[u] = 0$ in Algorithm 4.5. The protocol then proceeds according to the variant of SimpleAlgorithm, which does not require an order among the opinions described in Section 4.9. By Statement (2) of Lemma 4.12, at most $\mathrm{O}\big(n/x_{\max}\big)$ opinions have at least one `collector` agent each, w.h.p. If no single `collector` agent is left for some opinion, this opinion cannot become a challenger in any of the tournaments. Therefore, the total number of tournaments executed in SimpleAlgorithmis bounded w.h.p. by $\mathrm{O}\big(n/x_{\max}\big)$. As argued in the proof of Statement (1) of Theorem 4.1, each tournament takes $\mathrm{O}\big(n \log n\big)$ interactions w.h.p., and the modified SimpleAlgorithmalso

needs to perform a leader-election, which takes $\mathrm{O}(n \log^2 n)$ interactions [51]. We conclude that ImprovedAlgorithmhas a runtime of $\mathrm{O}(n^2/x_{\max} \cdot \log n + n \log^2 n)$ interactions w.h.p. □

Proof: Space Complexity of Theorem 4.2.
ImprovedAlgorithmrequires the states used in the modified initialization (Algorithm 4.5) and the states used by SimpleAlgorithm(Statement (2) of Theorem 4.1). In Algorithm 4.5, all `collector` agents need to store the set of states $\mathcal{S}_c$ of size $\Theta(\log \log n)$ required to run the junta-based phase clocks. Additionally, the size of the phase variable is increased by a constant, starting now at $-c$. The remaining states have the same size as in SimpleAlgorithm. This gives us the claimed state space size of $\Theta(k \cdot \log \log n + \log n)$. □

4.9. Removing the Order

In this section, we explain how to remove the assumption that there is an order among the k opinions. Recall that in SimpleAlgorithm, we let opinion 1 be the first defender and opinion $i+1$ be the challenger of the i-th tournament. The number of tournaments was counted in the tcnt variable of `tracker`-agents. Instead, we now assign the `tracker` agent a slightly different task, and we use a unique leader agent (from the set of `tracker` agents) that randomly *samples* the next challenger before each match.

The leader agent interacts until it encounters an opinion that has not yet participated in a tournament. Then, the leader agent informs all `collector`-agents u with that opinion that they are the next challenger. Unfortunately, this cannot be done efficiently for each opinion. If $x_j = o(n)$ for some opinion j, it takes too long for the leader to interact with an agent of that opinion. To solve this, we use the remaining `tracker`-agents. `tracker` agents copy opinions that have not yet competed in a tournament (using the same number of states as for the counter tcnt before). This effectively amplifies the number of agents having an opinion that has not yet participated in a tournament, making this opinion *visible* to the leader agent.

Challenger and Defender Selection Assume for now that we have a unique leader agent. At the beginning of each tournament ℓ in Phase 0 the leader agent and `tracker`-agents sample until they meet a `collector`-agent with an opinion j that has not yet participated in a tournaments. As soon as the leader agent has sampled such an opinion j (either from a `collector`-agent directly or from a `tracker`-agent), it starts to broadcast among the `tracker`-agents, and the `collector`-agents that opinion j is the challenger of tournament ℓ. (This broadcast is done on a constant fraction of all agents and thus concludes w.h.p. within one phase.) Now, when a `collector`-agent u with $\mathsf{opinion}[u] = j$ interacts with an agent v that knows the challenger opinion, it sets $\mathsf{challenger}[u] \leftarrow \textsc{True}$ and becomes a challenger agent. We can implement this broadcast using one additional bit in the state space. We can use the same procedure to select the initial defender before the tournament starts.

Lemma 4.13. *Assume a unique leader agent exists. Then a challenger (defender) opinion is selected in* $\mathrm{O}(n \log n)$ *interactions w.h.p.*

Proof. The lemma follows essentially from the following observation. Let u be an arbitrary but fixed agent, and let A be a set of agents with $|A| = \Omega(n)$. Then it follows from Chernoff bounds that in $O(n \log n)$ interactions u interacts with an agent $v \in A$ at least once.

We now give a detailed proof of the correctness of the challenger selection. The same arguments follow the defender selection. We call an opinion j *remaining challenger candidate* if the opinion has not participated in a tournament yet. First, we show that the leader agent selects one of the remaining challenger candidates in $\mathrm{O}(n \log n)$ interactions w.h.p. Then we show that every `collector`-agent with opinion j sets its challenger bit in $\mathrm{O}(n \log n)$ interactions w.h.p.

Let agent w be the leader and let R be the set of agents whose opinions are among the remaining challenger candidates. If $|R| \geq n/10$, the probability that the leader w interacts in a fixed step with an agent $v \in R$ is at least constant. From Chernoff's bounds, the leader agent selects a challenger candidate in $\mathrm{O}(n \log n)$ interactions w.h.p. Assume $|R| < n/10$. In this case, we first argue that every `tracker`-agent u stores the opinion of one of the remaining challenger candidates w.h.p. This follows from the one-way epidemic process [10] where R is the set of infected agents and the `tracker`-agents are susceptible. By Lemma 4.3, it follows that the number of `tracker`-agents is at least $n/10$ w.h.p., and hence, the first claim holds.

For the second claim, we argue similarly. The first claim shows that the leader has chosen a challenger opinion j w.h.p. The one-way epidemic provides that every `tracker`-agent learns the identity of opinion j within $O(n \log n)$ interactions w.h.p. We utilize an additional Boolean flag to determine whether a `tracker`-agent has already stored the challenger opinion j. It now remains to show every `collector`-agent with opinion j interacts at least once with a `tracker`-agent w.h.p. Again, this follows from Chernoff's bounds; hence, the second claim holds.

The statement follows from a union bound over both claims. □

Regarding the leader agent, we use the leader election protocol from [51] with the phase clock from [28]. We run this protocol among the `tracker` agents. It requires $\mathrm{O}(\log \log n)$ states and computes a unique leader agent in $\mathrm{O}(n \log^2 n)$ interactions w.h.p. Note that the unique leader recognizes when the leader election protocol is concluded. This allows us to reuse the states from the leader election and integrate the leader-election protocol in an additional, special phase before the tournaments start.

We now describe how we modify SimpleAlgorithm. The leader election protocol from [51] determines a unique leader agent. We execute this protocol among the `tracker`-agents in a special phase as part of the preprocessing before the first tournament starts. When a unique leader is elected (w.h.p.), it broadcasts the end of the leader election and initiates the initial defender selection. The `clock`-agents wait in `phase` 0 until they receive the signal that a unique leader exists. The challenger selection is executed at the beginning of a tournament j in Phase 0 and replaces the original challenger selection of SimpleAlgorithmin Lines 2-3 of Algorithm 4.4.

Proof of Statement (2) of Theorem 4.1. The result mostly follows from the correctness of SimpleAlgorithm(Statement (1) of Theorem 4.1). Again, by Lemma 4.3, it holds that the population is partitioned into the roles `player`, `tracker`, `clock` and `collector`, where each role consists of at least $n/10$ agents w.h.p. The key modification affects the selection of a unique leader agent, the initial defender opinion, and the challenger opinion for each tournament.

Since the number of `tracker`-agents is at least $n/10$, the unique leader agent is computed in $\mathrm{O}(n \log^2 n)$ interactions w.h.p. by the leader election protocol from [51]. Then, by Lemma 4.13, we have a defender opinion at the beginning of the first tournament w.h.p. Similarly to the proof of Statement (1) of Lemma 4.6, we can argue with Lemma 4.13 that Statement (1) holds. It remains to show that the number of states is at most $\mathrm{O}(k + \log n)$. The `tracker`-agents require $\mathrm{O}(\log \log n)$ states to execute the leader election protocol. Until the end of this protocol, they do not store any other values. Once the leader election has concluded, they disregard the $O(\log \log n)$ states used for that protocol and instead use their states to store an opinion. Hence, $\mathrm{O}(k + \log n)$ many states are sufficient.

The overall state complexity follows from the proof of Statement (1) of Theorem 4.1 along with the observation that the broadcasts can be implemented using constantly many additional bits. □

Chapter 5.

Undecided State Dynamics with Stubborn Agents

In this part, we consider a variant of the approximate majority consensus problem with preference. We introduce the preference by considering a biased variant of the undecided state dynamics. The following generalization of the USD has a *preferred* opinion: w.l.o.g. Opinion 1. We call Opinion 2 *unpreferred*. To model the preference on Opinion 1, we assume that agents with Opinion 1 are *stubborn* in the following sense. Whenever an agent with Opinion 1 meets an agent with Opinion 2, it does not become undecided immediately. Instead, it draws a random number in $(0, 1]$ and keeps its opinion if this number is smaller or equal to a constant $p \in [0, 1]$. Otherwise, it becomes undecided. We call p the *stubbornness* in the following. More detail, if the random scheduler picks a pair of agents with states $(1, 2)$, the initiator remains unchanged with probability p. With probability $1 - p$, its new state is $\delta(1, 2) = \perp$. All other interactions remain as in the original version of the undecided state dynamics.

Formally, the transition function of the stubborn USD with stubbornness p is

$$(q, q') \mapsto \begin{cases} \perp \text{ if } q = 2, q' = 1 \\ \perp \text{ if } q = 1, q' = 2 \text{ with probability } 1 - p \\ q' \text{ if } q = \perp \\ q \text{ otherwise.} \end{cases}$$

In the following, we refer to the stubborn USD process with stubbornness p by USD_p and its transition function by δ_p. $\mathrm{USD}_p(\mathbf{x})$ is defined as the USD_p with initial configuration $\mathbf{x}$. Consequently, the standard USD process is USD_0 and its transition function δ_0.

Results and Methodology Our focus is on the convergence time of $\mathrm{USD}_p(\mathbf{x})$, defined as the number of interactions until all agents agree on one of the two opinions if USD_p is started on initial configuration $\mathbf{x} = (x_1, x_2, u)$. Let $T_i(p, \mathbf{x})$ be the convergence time of process $\mathrm{USD}_p(\mathbf{x})$ assuming opinion i survives.

Note that $\Omega\big(n \log n\big)$ is a trivial lower bound for the process since this is the time until each agent is, w.h.p., activated at least once.

Theorem 5.1. *Let $\epsilon, p \in (0, 1]$ be arbitrary constants and let $\mathbf{x} = (x_1, x_2, u)$ be a configuration with $x_1 \in [\epsilon \cdot n, x_2], u \leq \frac{n}{2}$. Let $p_s := 1 - x_1/x_2$. Then the following statements* *hold w.h.p.*

$$T_1(p,\mathbf{x}) = O(n \cdot \log n) \qquad \textit{if } p - p_s = \Omega\left(\sqrt{n^{-1} \cdot \log n}\right), \tag{5.1}$$

$$T_2(p,\mathbf{x}) = O(n \cdot \log n) \qquad \textit{if } p_s - p = \Omega\left(\sqrt{n^{-1} \cdot \log n}\right), \tag{5.2}$$

$$T_{1\vee 2}(p,\mathbf{x}) = O(n \cdot \log^2 n) \qquad \textit{otherwise.} \tag{5.3}$$

In Fig. 5.1, we visualize the main result to highlight the critical configuration regimes. We essentially show a phase transition around the *threshold probability* $p_s = 1 - x_1/x_2$. If p is sufficiently larger than p_s, the process will reverse the initial bias with high probability, and the agents will agree on Opinion 1 (blue area). For p sufficiently smaller than p_s, with high probability, all agents will agree on Opinion 2 instead (red area). In the intermediate cases (grey area), either of the two opinions might win, but we can still provide bounds for the convergence time.

Note that for initial configurations $\mathbf{x} = (x_1, x_2, 0)$ with $x_1(0) > x_2(0)$, it is known that Opinion 1 is more likely to win in $O(n \log n)$ interactions, even for $p = 0$ (see [35]). That is, the preferred opinion with stubbornness $p > 0$ has an advantage over opinion 2 to succeed. Therefore, we focus on scenarios where the initial majority opinion is 2.

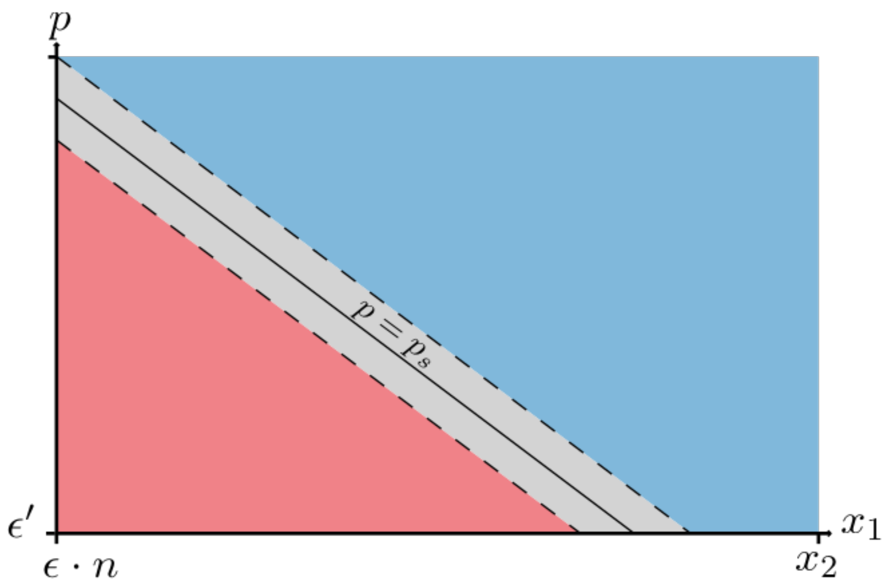

Figure 5.1.: Overview of the results: The x-coordinate indicates x_1 from $\varepsilon \cdot n$ to x_2. The y-coordinate indicates the stubbornness p. The black diagonal line represents $p = p_s$ and the dashed lines around it represent $p = p_s \pm \Theta\left(\sqrt{n^{-1} \cdot \log n}\right)$

As mentioned before, the case $p = 0$ is well-studied. A standard approach for this case is to track the evolution of the support of the opinions, e.g., using the (additive) *bias* $x_1(t) - x_2(t)$ or the (multiplicative) *gap* $x_1(t)/x_2(t)$ as a potential function (see [5, 9, 35]). As a first contribution, we generalized this approach by identifying the driving force of the biased version of the undecided state dynamics, which we call *weighted bias* (as a function of t) $\Delta_w(t)$. The weighted bias at time t is defined as $x_1(t) - (1-p)x_2(t)$. The initial weighted bias plays a similar role in determining the winning opinion as the initial bias does in the classical USD with $p = 0$. Note that an *initial weighted bias* $\Delta_w(0)$ of $c \cdot x_2(0)$ is equivalent to $p = 1 - x_1(0)/x_2(0) + c = p_s + c$. Weighted bias is a more practical way of looking at the problem.

The most difficult part of our analysis is those configurations where the stubbornness parameter balances out the initial support deficit of the preferred opinion, i.e., $p = p_s$ and equivalently $\Delta_w(0) = 0$. In contrast to the corresponding case with $p = 0$ and $x_1 = x_2$, where the initial bias is zero, creating a sufficiently large weighted bias is more involved. The standard way is to define a random walk on the integers and apply known anti-concentration and concentration bounds. This approach is not viable here, since $|\Delta_w(t+1) - \Delta_w(t)| \in \{0, 1-p, 1\}$ leads to non-integer states. Instead, we exploit the submartingale property of $\Delta_w(t)$ itself and the function $Y_t = \Delta_w^2(t) - r \cdot t$ for a suitably chosen value of r.

5.1. Cases in which Opinion 1 wins

In this section, we show Statement 5.1 of Theorem 5.1, namely that Opinion 1 wins if the stubbornness p is sufficiently larger than $1 - x_1(0)/x_2(0)$. We show in Lemma 5.3 that the gap increases w.h.p. such that there exists a time $T_1 = \mathrm{O}(n \log n)$ where each agent either has Opinion 1 or is undecided. From there on, it is easy to show that all agents agree on Opinion 1 after an additional $\mathrm{O}(n \log n)$ steps. As an auxiliary result for Lemma 5.3, we first show in Lemma 5.2 that the weighted bias does not halve during the first interactions.

Lemma 5.2. *Let $\boldsymbol{x}(t_0)$ be a configuration with weighted bias $\Delta_w(t_0) > 0$. Let $\xi(\tau)$ be the event that $\Delta_w(t) \geq \Delta_w(t_0)/2$ for all $t \in [t_0, \ldots, t_0 + \tau]$. Then, with probability at least $1 - n^{-6}$, $\xi(T)$ holds for all $\tau \leq \Delta_w^2(t_0)/(16 \ln n)$.*

Proof. We aim to apply an Azuma-Hoeffding-bound (Lemma A.20) to $\Delta_w(t)$ for each $t \in [t_0, t_0 + \tau]$ with $\tau = \Delta_w^2(t_0)/(16 \ln n)$. Fix an arbitrary $t \leq \tau$. To apply Lemma A.20, we need to show that $\mathbb{E}[\Delta_w(t+1) - \Delta_w(t) | \mathcal{F}_t] \geq 0$ and that $|\Delta_w(t+1) - \Delta_w(t)|$ is bounded.

First, we calculate the expected change in Δ_w by considering all possible interactions. For ease of presentation, we drop the parameter t whenever clear from the context. With probability $x_1 \cdot x_2/n^2$, the randomly chosen initiator has Opinion 1, and the responder has Opinion 2. In that case, the initiator is stubborn (it does not change its state) with probability p, resulting in $\Delta_w(t+1) = \Delta_w(t)$. With probability $1-p$ the initiator becomes undecided and $\Delta_w(t+1) = \Delta_w(t) - 1$. With probability $x_2 \cdot x_1/n^2$, the initiator loses its Opinion 2 and becomes undecided, resulting in $\Delta_w(t+1) = \Delta_w(t) + (1-p)$. Whenever an undecided agent initiates an interaction where the responder has either Opinion 1 or Opinion 2, it adopts the responder's opinion. Such interactions occur with probability $u \cdot x_1/n^2$ and $u \cdot x_2/n^2$, resulting in $\Delta_w(t+1) = \Delta_w(t) + 1$ and $\Delta_w(t+1) = \Delta_w(t) - (1-p)$.

At last, neutral interactions exist that do not change the potential, i.e., $\Delta_w(t+1) = \Delta_w(t)$. These interactions occur with the remaining probability $(x_1^2 + x_2^2 + n \cdot u)/n^2$.

$$\begin{aligned}
&\mathbb{E}\left[\Delta_w(t+1) - \Delta_w(t) | \mathcal{F}_t\right] \\
&\quad = \frac{x_1 \cdot x_2}{n^2} \cdot (p \cdot 0 + (1-p) \cdot (-1)) + \frac{x_2 \cdot x_1}{n^2} \cdot (1-p) + \frac{u \cdot x_1}{n^2} \cdot (+1) \\
&\qquad + \frac{u \cdot x_2}{n^2} \cdot (-1+p) + \frac{x_1^2 + x_2^2 + n \cdot u}{n^2} \cdot 0 \\
&\quad = -\frac{(1-p) \cdot x_1 \cdot x_2}{n^2} + \frac{(1-p) \cdot x_2 \cdot x_1}{n^2} + \frac{u \cdot x_1}{n^2} - \frac{(1-p) \cdot u \cdot x_2}{n^2} \\
&\quad = \frac{u}{n^2} \cdot \Delta_w(t) \geq 0.
\end{aligned}$$

It remains to show that $|\Delta_w(t+1) - \Delta_w(t)|$ is bounded. Here, we show that the term is bounded by 1. Consider every possible interaction at time $t+1$. If the initiator does not change its opinion, $|\Delta_w(t+1) - \Delta_w(t)| = 0$. If the support of Opinion 1 changes, we have $|\Delta_w(t+1) - \Delta_w(t)| = 1$. Otherwise, the support of Opinion 2 changes by one and $|\Delta_w(t+1) - \Delta_w(t)| = 1 - p \leq 1$.

Now we are ready to apply Lemma A.20 with $\lambda = \Delta_w(t_0)/2$.

$$\begin{aligned}
&\Pr\left[\Delta_w(t) < \frac{\Delta_w(t_0)}{2} \;\middle|\; \mathcal{F}_{t_0}\right] = \Pr\left[\Delta_w(t) - \Delta_w(t_0) < -\lambda \mid \mathcal{F}_{t_0}\right] \leq \exp\left(-\frac{2\lambda^2}{t}\right) \\
&\quad \leq \exp\left(-\frac{2\lambda^2}{\tau}\right) \leq \exp\left(-8 \ln n\right).
\end{aligned}$$

By application of the union bound over the first τ interactions, the statement holds with probability of at least $1 - \tau \cdot \exp(-8 \ln n) \geq 1 - n^2 \cdot n^{-8} = 1 - n^{-6}$. □

Note that Lemma 5.2 is an auxiliary result since it only shows that the weighted bias is not decreasing too much for $\Omega\big(\Delta_w^2(t_0)/\log n\big)$ time. We use this in the following lemma to show that the support of the initial majority opinion drops to zero during that time.

Lemma 5.3. *Let $\boldsymbol{x}(t_0)$ be a configuration with weighted bias $\Delta_w(t_0) \geq c_s \cdot n$ for an arbitrary constant c_s. Let $T_1 = \inf\{t \geq 0 \mid x_2(t) = 0\}$. Then,* $\Pr\left[T_1 \leq 20 \cdot c_s^{-1} \cdot n \log n\right] \geq 1 - n^{-2}$.

Proof. Let $\Psi(t) = x_2(t)/x_1(t)$ denote the inverse of the *gap*. The idea is to show that this potential function decreases exponentially and apply a known drift theorem. Similar to the proof of Lemma 5.2, we calculate the expected change in Ψ by considering all possible interactions.

$$
\begin{aligned}
&\mathbb{E}\left[\Psi(t+1)-\Psi(t)\mid\mathcal{F}_t\right]\\
&\quad=\frac{x_1\cdot x_2}{n^2}\left(p\cdot 0+(1-p)\cdot\left(\frac{x_2}{x_1-1}-\frac{x_2}{x_1}\right)\right)+\frac{x_2\cdot x_1}{n^2}\left(\frac{x_2-1}{x_1}-\frac{x_2}{x_1}\right)\\
&\qquad+\frac{u\cdot x_1}{n^2}\left(\frac{x_2}{x_1+1}-\frac{x_2}{x_1}\right)+\frac{u\cdot x_2}{n^2}\left(\frac{x_2+1}{x_1}-\frac{x_2}{x_1}\right)+\frac{x_1^2+x_2^2+n\cdot u}{n^2}\cdot 0\\
&\quad=\frac{(1-p)x_1\cdot x_2^2}{n^2}\left(\frac{1}{x_1(x_1-1)}\right)-\frac{x_2\cdot x_1}{n^2\cdot x_1}-\frac{u\cdot x_1\cdot x_2}{n^2\cdot x_1(x_1+1)}+\frac{u\cdot x_2}{n^2\cdot x_1}\\
&\quad=\frac{\Psi}{n^2}\cdot\left(\frac{(1-p)x_2}{x_1-1}-x_1-\frac{u\cdot x_1}{x_1+1}+u\right)\\
&\quad=-\frac{\Psi}{n^2}\cdot\left(x_1-(1-p)x_2+\frac{(1-p)x_2}{x_1-1}-\frac{u}{x_1+1}\right).
\end{aligned}
$$

Observe that the expected potential change in Ψ is a function of the weighted bias. We bound the weighted bias using Lemma 5.2: $\Delta_w(t)\geq\Delta_w(t_0)/2$ with probability at least $1-n^{-6}$ for all $t\in[t_0,t_0+c_3 n\log n]$. We trivially bound $(1-p)x_2/(x_1-1)\geq 0$. To bound the term $u/(x_1+1)$, we use $x_1\geq\Delta_w(t)\geq c_s\cdot n/2$ and $u\leq n$. Then, it holds w.h.p. that

$$
\begin{aligned}
\mathbb{E}\left[\Psi(t+1)-\Psi(t)\mid\mathcal{F}_t\right]&=-\frac{\Psi}{n^2}\cdot\left(\Delta_w(t)+\frac{(1-p)x_2}{x_1-1}-\frac{u}{x_1+1}\right)\\
&\leq-\frac{\Psi}{n^2}\cdot\left(\Delta_w(t)-\frac{u}{x_1+1}\right)\\
&\leq-\frac{\Psi}{n^2}\cdot\left(\frac{c_s\cdot n}{2}-\frac{n}{(c_s/2)\cdot n+1}\right)\\
&\leq-\frac{\Psi}{n}\cdot\left(\frac{c_s}{2}-\frac{2}{c_s\cdot n+2}\right)\\
&\leq-\frac{\Psi}{n}\cdot\frac{c_s}{4}.
\end{aligned}
$$

We now apply the multiplicative drift theorem (Theorem A.9, found in the appendix) to bound T_1 with $r=3\ln n$, $s_0=x_2(t_0)/x_1(t_0)\leq n$, $s_{\min}=(n-1)^{-1}$ and $\delta=n^{-1}\cdot c_s/4$. Then, we get

$$
\Pr\left[T_1>\frac{20n\cdot\ln n}{c_s}\right]\leq\Pr\left[T_1>\left\lceil\frac{r+\ln(s_0/s_{\min})}{\delta}\right\rceil\right]\leq e^{-r}=n^{-3}.
$$

Note that in order to apply Theorem A.9, we have to have $\mathbb{E}\left[\Psi(t)-\Psi(t+1)\mid\mathcal{F}_t\right]\geq\delta\cdot\Psi(t)$ for all $\Psi(t)\neq 0$ and all $t\geq t_0$. Lemma 5.2 asserts this only with high probability and for a limited time. But we can consider a process that deterministically jumps to configuration $(n,0,0)$ at time $t+1$ if for any t Lemma 5.2 is violated. We can apply Theorem A.9 to this process, and our original process behaves identical to it with probability $1-n^{-6}$. Thus, the lemma follows from the union bound. □

Recall that by assumption of Statement 5.1 Theorem 5.1, we have $\Delta_w(0) = \Omega\big(\sqrt{n} \cdot \log n\big)$. We now show that the weighted bias doubles every $\mathrm{O}\big(n\big)$ interaction until it is of size $\Theta\big(n\big)$ – the concrete bound of $n/10$ we show was chosen rather arbitrarily. Then we can apply Lemma 5.3. The proof of Lemma 5.4 requires bounds on the number of undecided agents (Lemma 5.9 and Lemma 5.10) that we postpone to Section 5.4.

Lemma 5.4. *Let* $\mathbf{x}(t_0)$ *be a configuration with* $\Delta_w(t_0) \geq \xi \cdot \sqrt{n \log n}$ *and let* $T = \inf \{\, t \geq t_0 \;\mid\; \Delta_w(t) \geq n/10 \,\}$. *Then* $\Pr\big[T \leq (\xi^2/6) \cdot n \log n\big] \geq 1 - n^{-3}$.

Proof. The proof idea is inspired by a repetition of the Gambler's ruin problem. For $1 \leq \ell \leq \log n$, we define the time intervals $\mathcal{I}_\ell = \{\, T_{\ell-1}, \ldots, T_\ell - 1 \,\}$ with $T_0 = t_0$ and

$$T_\ell \coloneqq \inf \{\, t \geq t_0 \;\mid\; \Delta_w(t) \geq \min(2^\ell \cdot \Delta_w(t_0), (n - u(t))/4) \,\}\,.$$

We show that the weighted bias leaves the interval $[\Delta_w(T_\ell)/2, 2\Delta_w(T_\ell)]$ within $\mathrm{O}\big(n\big)$ interactions, and that the value $2\Delta_w(T_\ell)$ is reached before $\Delta_w(T_\ell)/2$. We apply this result repeatedly until $\Delta_w(T_\ell) \geq (n - u(t))/4$. Recall from Lemma 5.2 that $(\Delta_w(t))_{t \geq t_0}$ is a submartingale with

$$\mathbb{E}\left[\Delta_w(t+1) \mid \mathcal{F}_t\right] \geq \Delta_w(t) + \frac{u(t) \cdot \Delta_w(t)}{n^2} \geq \Delta_w(t).$$

Assume that $\Delta_w(T_\ell) \geq \xi \cdot \sqrt{n \log n}$ and let $T_{\ell,\min} \coloneqq \inf \{\, t \geq T_\ell \;\mid\; \Delta_w(t) < \Delta_w(T_\ell)/2 \,\}$ and $\tau \coloneqq (\xi^2/6) \cdot n$. Then we bound $\Pr\left[T_{\ell,\min} > \tau\right]$ by using the Azuma-Hoeffding bound (Lemma A.20) with $\lambda = \Delta_w(T_\ell)/2$:

$$\begin{aligned}\Pr\left[T_{\ell,\min} < \tau\right] &= \Pr\left[\Delta_w(T_\ell + \tau) < \Delta_w(T_\ell)/2\right] = \Pr\left[\Delta_w(T_\ell + \tau) - \Delta_w(T_\ell) < -\lambda\right] \\ &\leq e^{-\frac{2\lambda^2}{(\xi^2/6)n}} \leq e^{-\frac{3\xi^2 n \log n}{\xi^2 n}} \leq n^{-3}.\end{aligned}$$

Next, let $T_{\ell,\max} \coloneqq \inf \{\, t \geq T_\ell \;\mid\; \Delta_w(t) \geq 2\Delta_w(T_\ell) \,\}$. We bound $\Pr\left[T_{\ell,\max} > \tau\right]$ assuming $\Delta_w(t) \geq \Delta_w(T_\ell)/2$ for all $t \in [T_\ell, T_\ell + \tau]$. Let

$$Z \coloneqq \sum_{t=T_\ell+1}^{T_\ell+\tau} \Delta_w(t) \text{ and } \mu \coloneqq \sum_{t=T_\ell+1}^{T_\ell+\tau} \mathbb{E}\left[\Delta_w(t) \mid \mathcal{F}_{t-1}\right].$$

In contrast to Lemma 5.2, to establish a sufficiently strong bias in the right direction, we have to exploit the additional positive term $u(t) \cdot \Delta_w(t)/n^2$.

From Lemma 5.10 it follows that w.h.p. $u(t) \geq c_u \cdot n$ for constant $c_u = (3/10) \cdot (1 - p)/(2 - p)$ as long as $\Delta_w(t) < (n - u(t))/4$ holds. Note that this requires additional $O(n)$ interactions if it is too small. Lemma 5.2 guarantees that the weighted bias does not lose too much support during this time. Therefore, as long as $\Delta_w(t) \geq \Delta_w(T_\ell)/2$ holds, we have

$$\mathbb{E}\left[\Delta_w(t+1) \mid \mathcal{F}_t\right] \geq \Delta_w(t) + \frac{u(t) \cdot \Delta_w(t)}{n^2} \geq \Delta_w(t) + \frac{c_u \cdot \Delta_w(T_\ell)}{2n} =: \Delta_w(t) + \varepsilon.$$

Then, it follows from the conditional Hoeffding bound (full version of [5]) for $\lambda = 2\xi \cdot \Delta_w(T_\ell)/2$

$$\begin{aligned}
&\Pr\left[\Delta_w(T_\ell+\tau) < 2\Delta_w(T_\ell)\right] \\
&\quad= \Pr\left[\Delta_w(T_\ell+\tau) < \Delta_w(T_\ell) + \tau\cdot\varepsilon - \lambda\right] \\
&\quad= \Pr\left[\sum_{t=T_\ell+1}^{T_\ell+\tau} \Delta_w(t) < \sum_{t=T_\ell}^{T_\ell+\tau-1} (\Delta_w(t)+\varepsilon) - \lambda\right] = \Pr\left[Z-\mu<\lambda\right] \le e^{-\frac{2\lambda^2}{4(\xi^2/6)n}} \le n^{-3}.
\end{aligned}$$

It follows by the union bound over the high probability events from above that

$$\Pr\left[\exists t \in [T_\ell, T_\ell+\tau] \colon \Delta_w(t) \ge \min\left\{\, 2\cdot\Delta_w(T_\ell), (n-u(t))/4 \,\right\}\right] \ge 1 - n^{-2}.$$

Applied to a fixed ℓ, this implies the length of the interval $\mathcal{I}_\ell$ is w.h.p. at most $(\xi^2/6)\cdot n$. From the union bound over all $\ell \le \log n$ intervals, we get there exists a time $t \in [T_0, T_{\log n}]$ such that $\Delta_w(t) \ge n - u(t)/4$. Otherwise, $\Delta_w(t) \ge 2^{\log n}\cdot\Delta_w(t_0) > n$, leading to a contradiction. The length of $[T_0, T_{(\log n)}]$ is at most $(\xi^2/6)\cdot n\cdot\log n$. At last, it follows from Lemma 5.9 that $(n-u(t))/4 \ge n/10$. □

We are ready to prove Statement 5.1 of Theorem 5.1.

Proof of Statement 5.1. Consider a configuration with $x_1(0) \in [\epsilon\cdot n, x_2], u(0) \le n/2$ and $p - p_s = \Omega\big(n^{-1/2}\cdot\log n\big)$. Then equivalently $\Delta_w(0) = \Omega\big(x_2\cdot n^{-1/2}\cdot\log n\big)$ and $x_1 + x_2 \ge n/2$. Since $x_1 \le x_2$, we have $x_2 = \Theta\big(n\big)$ and thus $\Delta_w(0) \ge \xi\cdot\sqrt{n\log n}$ for some constant ξ.

$$\begin{aligned}
T_a &= \inf\left\{\, t \ge 0 \;\middle|\; \Delta_w(t) \ge n/10 \,\right\} \\
T_b &= \inf\left\{\, t \ge 0 \;\middle|\; x_2(t) = 0 \,\right\}.
\end{aligned}$$

From Lemma 5.4, we have w.h.p. $T_a = \mathrm{O}\big(n\log n\big)$ and due to Lemma 5.3 $T_b = T_a + \mathrm{O}\big(n\log n\big)$.

Note that at no time can all agents be undecided since the last agent with Opinion 1 cannot encounter Opinion 2. Therefore, at time T_b, at least one agent with Opinion 1 exists.

With $x_2(T_b) = 0$, the process simplifies to a single productive rule: $\delta(1,\perp) = 1$. Let $T_1 = \inf\left\{\, t \ge T_b \;\middle|\; x_1(t) = n \,\right\}$. Assuming that $T_b < \infty$ and $x_1(T_1) = 1$, we have $T_1 \le T_b + 6n\log n$ with probability at least $1 - n^{-2}$. It is easy to see that $x_1(T_b) = 1$ gives an upper bound for T_1. The statement then follows from the union bound. □

5.2. Cases in which Opinion 2 wins

In this section, we prove Statement 5.2 of Theorem 5.1, namely that Opinion 2 wins if p is sufficiently smaller than $1 - x_1(0)/x_2(0)$. The general approach is identical to Section 5.1; given the asymmetric nature of the problem, some calculations differ slightly. For any $t \ge 0$, we call $\Delta_{\bar{w}} = (1-p)x_2(t) - x_1(t)$ the *negative weighted bias.* Note that for

$p = 1 - x_1(0)/x_2(0) - \gamma$, the initial negative weighted bias is $\gamma \cdot x_2(0)$. Analogous to Lemma 5.2, we show that the negative weighted bias does not decrease significantly for polynomial many interactions in this setting.

Lemma 5.5. *Let $\boldsymbol{x}(t_0)$ be a configuration with weighted bias $\Delta_{\bar{w}}(t_0) \geq c_s \cdot n$. Let $\xi(\tau)$ be the event that $\Delta_{\bar{w}}(t) \geq \Delta_{\bar{w}}(t_0)/2$ for all $t \in [t_0, \dots, t_0 + \tau]$. Then, with probability at least $1 - n^{-6}$, $\xi(T)$ holds for all $\tau \leq \Delta^2_{\bar{w}}(t_0)/(16 \ln n)$.*

Proof sketch. The proof follows along the lines of that of Lemma 5.2 with $\Delta_{\bar{w}}(t) = -\Delta_w(t)$. □

Lemma 5.6. *Let $\boldsymbol{x}(t_0)$ be a configuration with weighted bias $\Delta_{\bar{w}}(t_0) \geq c_s \cdot n$ for an arbitrary constant c_s. Let $T_1 = \inf\{t \geq 0 \mid x_2(t) = 0\}$. Then,* $\Pr\left[T_1 \leq 20 \cdot c_s^{-1} \cdot n \log n\right] \geq 1 - n^{-2}$.

Proof. The proof follows along the lines of that of Lemma 5.3 with the potential function $\Psi(t) = x_1(t)/x_2(t)$. Recall that the idea is to calculate the expected change of the potential function $\Psi(t)$ and apply a known drift theorem. From Lemma 5.5 and the initial size of $x_2(0)$ we get that w.h.p. $\Delta_{\bar{w}}(t) \geq \Delta_{\bar{w}}(t_0)/2$. In particular, we also get $x_2(t) \geq c_2 \cdot n$ for some constant c_2. This allows us to bound the expected change as follows.

$$
\begin{aligned}
&\mathbb{E}\left[\Psi(t+1) - \Psi(t) \mid \mathcal{F}_t\right] \\
&\quad = \frac{x_1 \cdot x_2}{n^2}\left(p\Psi + (1-p)\frac{x_1 - 1}{x_2}\right) + \frac{x_2 \cdot x_1}{n^2}\left(\frac{x_1}{x_2 - 1}\right) + \frac{u \cdot x_1}{n^2}\left(\frac{x_1 + 1}{x_2}\right) \\
&\quad\quad + \frac{u \cdot x_2}{n^2}\left(\frac{x_1}{x_2 + 1}\right) + \frac{x_1^2 + x_2^2 + (x_1 + x_2 + u)u}{n^2} \cdot \Psi - \Psi \\
&\quad = \frac{x_1 \cdot x_2}{n^2}\left(\Psi + \frac{1-p}{x_2}\right) + \frac{x_2 \cdot x_1}{n^2}\left(\Psi - \frac{x_1}{x_2} + \frac{x_1}{x_2 - 1}\right) + \frac{u \cdot x_1}{n^2}\left(\Psi + \frac{1}{x_2}\right) \\
&\quad\quad + \frac{u \cdot x_2}{n^2}\left(\Psi - \frac{x_1}{x_2} + \frac{x_1}{x_2 + 1}\right) + \frac{x_1^2 + x_2^2 + (x_1 + x_2 + u)u}{n^2} \cdot \Psi - \Psi \\
&\quad = \frac{\Psi}{n^2}\left(-(1-p)x_2 + \frac{x_1 \cdot x_2}{x_2 - 1} + u - \frac{u \cdot x_2}{x_2 + 1}\right) \\
&\quad = -\frac{\Psi}{n^2}\left((1-p)x_2 - x_1 - \frac{x_1}{x_2 - 1} - \frac{u}{x_2 + 1}\right) \\
&\quad \leq -\frac{\Psi}{n^2}\left(c_s \cdot n - \frac{n}{c_2 \cdot n - 1} - \frac{n}{c_2 \cdot n + 1}\right) \\
&\quad = -\frac{\Psi}{n}\left(c_s - \frac{1}{c_2 \cdot n - 1} - \frac{1}{c_2 \cdot n + 1}\right) \\
&\quad \leq -\frac{c_s}{2n} \cdot \Psi
\end{aligned}
$$

We now apply the multiplicative drift theorem (Theorem A.9) with $r = 3\ln n$, $s_0 = x_1(0)/x_2(0) \leq n$, $s_{\min} = (n-1)^{-1}$, $\delta = c_s/(2n)$ and get

$$\begin{aligned}\Pr\left[T_1 > 12c_s^{-1} n\ln n\right] &= \Pr\left[T_1 > \left\lceil \frac{12n\ln n}{c_s}\right\rceil\right] \leq \Pr\left[T_1 > \left\lceil \frac{3\ln n + \ln(n/(n-1)^{-1})}{c_s/(2n)}\right\rceil\right] \\ &\leq \Pr\left[T_1 > \left\lceil \frac{r + \ln(s_0/s_{\min})}{\delta}\right\rceil\right] \leq e^{-r} = n^{-3}\end{aligned}$$

□

Akin to Lemma 5.4, we now consider the case $\Delta_{\bar{w}} = o(n)$.

Lemma 5.7. *Let* $\mathbf{x}(t_0)$ *be a configuration with* $\Delta_{\bar{w}}(t_0) \geq \xi \cdot \sqrt{n \log n}$ *and let* $T = \inf\{\, t \geq t_0 \mid \Delta_{\bar{w}}(t) \geq n/10 \,\}$*. Then* $\Pr\left[T \leq (\xi^2/6)\cdot n\log n\right] \geq 1 - n^{-3}$*.*

Proof sketch. The proof follows along the lines of that of Lemma 5.4 for $\Delta_{\bar{w}}$ instead of Δ_w using Lemma 5.5 instead of Lemma 5.2. □

We now prove Statement 5.2 of Theorem 5.1.

Proof of Statement 5.2. Consider a configuration with $x_1(0) \in [\epsilon \cdot n, x_2], u(0) \leq n/2$ and $p_s - p = \Omega\big(n^{-1/2} \cdot \log n\big)$. Then equivalently $\Delta_{\bar{w}}(0) = \Omega\big(x_2 \cdot n^{-1/2} \cdot \log n\big)$ and $x_1 + x_2 \geq n/2$. Since $x_1 \leq x_2$, we have $x_2 = \Theta\big(n\big)$ and thus $\Delta_{\bar{w}}(0) \geq \xi \cdot \sqrt{n\log n}$ for some constant ξ.

$$\begin{aligned} T_a &= \inf\{\, t \geq 0 \mid \Delta_{\bar{w}}(t) \geq n/10 \,\} \\ T_b &= \inf\{\, t \geq 0 \mid x_1(t) = 0 \,\}\end{aligned}$$

By Lemma 5.7, we have w.h.p. $T_a = \mathrm{O}\big(n\log n\big)$ and then by Lemma 5.6 $T_b = T_a + \mathrm{O}\big(n \log n\big)$. Note that at no time can all agents be undecided since the last agent with Opinion 2 cannot encounter Opinion 2. Therefore, at time T_b, at least one agent with Opinion 2 exists.

With $x_1(T_b) = 0$, the process simplifies to a single productive rule: $\delta(2, \perp) = 1$. Let $T_2 = \inf\{\, t \geq T_b \mid x_2(t) = n \,\}$. Assuming that $T_b < \infty$ and $x_2(T_2) = 1$, we have $T_2 \leq T_b + 6n\log n$ with probability at least $1 - n^{-2}$. It is easy to see that $x_2(T_b) = 1$ gives an upper bound for T_2. The statement then follows from the union bound. □

5.3. Cases in which either Opinion 1 or Opinion 2 wins

In this section we consider the critical regime of $p \approx 1 - x_1(0)/x_2(0)$, i.e., the initial weighted bias ($\Delta_w(0) \coloneqq x_1(0) - (1-p)\cdot x_2(0)$) is small. We show that we reach a configuration after a while with a sufficiently large weighted bias ($|\Delta_w(t)| = \Omega\big(\sqrt{n\log n}\big)$). We do so by defining a submartingale $(Y_t)_{t\geq 0}$ as $Y_t = \Delta_w^2(t) - r \cdot t$ for a suitably chosen constant r and applying tail bounds (Lemma 5.8). At this point, either Statement 5.1 or Statement 5.3 of Theorem 5.1 applies.

We define T_w as the first time the process reaches such a weighted bias. More formally, $T_w \coloneqq \inf\{\, t \geq 0 : |\Delta_w(t)| \geq \xi \cdot \sqrt{n \cdot \log n}\,\}$. We show that $T_w = \mathrm{O}\big(n \cdot \log^2 n\big)$ (Lemma 5.8). One standard approach combines anti-concentration bounds and concentration bounds for random walks. Defining a random walk on the weighted bias $\Delta_w(t)$ is rather complicated since this results in a non-integer state space. Instead, we define two submartingales $(Z_t)_{t \geq 0}$ with $Z_t = \Delta_w(t)$ and $Y_t = Z_t^2 - r \cdot t$. We show that Y_t is a submartingale for a suitably chosen constant r and then we prove that $\mathbb{E}\left[Z_T - Z_0\right] = \sqrt{T}$ by considering Y_t. To bound $|Y_{t+1} - Y_t|$ we use tail bounds (see Lemma A.20). Unfortunately, there is one more challenge we have to address. We can only show that Y_t is a submartingale if the number of undecided agents is not too large (shown in Section 5.4).

Lemma 5.8. *Let $\boldsymbol{x}(0)$ be a configuration with $|\Delta_w(0)| < \xi \cdot \sqrt{n \cdot \log n}$ for an arbitrary constant ξ. Let $T_w \coloneqq \inf\{\, t \geq 0 : |\Delta_w(t)| \geq \xi \cdot \sqrt{n \cdot \log n}\,\}$. Then w.h.p. $T_w = \mathrm{O}\big(n \cdot \log^2 n\big)$.*

Proof. The idea of the proof is to apply the Azuma-Hoeffding bound to a suitable submartingale. Let $T_u \coloneqq \inf\{\, t \geq 0 : u(t) > x_1(t) + x_2(t) + 6\xi \cdot \sqrt{n \cdot \log n}\,\}$. We define $Y_t = \Delta_w(t)^2 - r \cdot t$ for $t < \min\{\, T_w, T_u \,\}$ where the constant r is chosen later. Otherwise, $Y_t = Y_{t-1}$. Note that by Lemma 5.9, w.h.p. $T_u = \omega\big(n \log^2 n\big)$. We show that $Y_0, Y_1, \ldots$ is a submartingale. The calculation is similar to that in the proof of Lemma 5.2. We consider every possible interaction and the resulting change. Assume that $t < \min\{\, T_w, T_u \,\}$. Then,

$$
\begin{aligned}
&\mathbb{E}\left[Y_{t+1} | \mathcal{F}_t, \Delta_w(t) = \Delta_w\right] \\
&\quad = \mathbb{E}\left[\Delta_s^2(t+1) - r \cdot (t+1) | \mathcal{F}_t, \Delta_w(t) = \Delta_w\right] \\
&\quad = \frac{x_1 \cdot x_2}{n^2} \cdot \big(p \cdot \Delta_w^2 + (1-p) \cdot (\Delta_w - 1)^2\big) + \frac{x_2 \cdot x_1}{n^2} \cdot (\Delta_w + 1 - p)^2 \\
&\quad\quad + \frac{u \cdot x_1}{n^2} \cdot (\Delta_w + 1)^2 + \frac{u \cdot x_2}{n^2} \cdot (\Delta_w - 1 + p)^2 \\
&\quad\quad + \frac{n^2 - 2x_1 \cdot x_2 - u \cdot (x_1 + x_2)}{n^2} \cdot \Delta_w^2 - r \cdot (t+1) \\
&\quad = \frac{x_1 \cdot x_2}{n^2} \cdot \big(2\Delta_w^2 + (2-p)(1-p)\big) + \frac{u \cdot x_1}{n^2} \cdot \big(\Delta_w^2 + 2\Delta_w + 1\big) \\
&\quad\quad + \frac{u \cdot x_2}{n^2} \cdot \big(\Delta_w^2 - 2(1-p)\Delta_w + (1-p)^2\big) \\
&\quad\quad + \frac{n^2 - 2x_1 \cdot x_2 - u \cdot (x_1 + x_2)}{n^2} \cdot \Delta_w^2 - r \cdot (t+1) \\
&\quad = \Delta_w^2 + \frac{x_1 \cdot x_2}{n^2} \cdot (2-p) \cdot (1-p) \\
&\quad\quad + \frac{u \cdot 2\Delta_w}{n^2} \cdot (x_1 - (1-p)x_2) + \frac{u}{n^2} \cdot \big(x_1 + (1-p)^2 x_2\big) - r \cdot (t+1) \\
&\quad = Y_t + \frac{x_1 \cdot x_2}{n^2} \cdot (2-p) \cdot (1-p) + \frac{u \cdot 2 \cdot \Delta_w^2}{n^2} + \frac{u}{n^2} \cdot (x_1 + (1-p)^2 x_2) - r \\
&\quad \geq Y_t + \frac{x_1 \cdot (1-p)x_2}{n^2} - r.
\end{aligned}
$$

Note that p is constant and for $t < \min\{T_w, T_u\}$ it holds that $u < x_1 + x_2 + o(n)$ and $x_1 - (1-p)x_2 = o(n)$. Then $x_1, (1-p)x_2 = \Theta(n)$. So there exists a worst case bound on $(x_1 \cdot (1-p)x_2)/n^2$ that is in $\Theta(1)$. Using that bound for r yields the desired result $\mathbb{E}[Y_{t+1}|\mathcal{F}_t] \geq Y_t$. For $t \geq \min\{T_w, T_u\}$, we have $\mathbb{E}[Y_{t+1}|\mathcal{F}_t] = Y_t$ by definition.

Next, we apply the Azuma-Hoeffding bound (Lemma A.20) with $\tau = \alpha \cdot n \cdot \log^2 n$ for some constant α that we determine later and $\lambda = \Delta_w^2(0) + \xi \cdot \sqrt{10 \cdot \tau \cdot n} \cdot \log n \geq \xi \cdot \sqrt{10 \cdot \tau \cdot n} \cdot \log n$. Furthermore, for $t < \min\{T_w, T_u\}$ we have

$$\begin{aligned}|Y_{t+1} - Y_t| &\leq (\Delta_w(t) + 1)^2 - r \cdot (t+1) - \Delta_w^2(t) + r \cdot t = 2\Delta_w(t) + 1 - r \\ &\leq 2 \cdot \xi \cdot \sqrt{n \cdot \log n} + 1.\end{aligned}$$

Thus we have $(b-a)^2 \leq \left(2 \cdot (2 \cdot \xi \cdot \sqrt{n \cdot \log n} + 1)\right)^2 \leq 5 \cdot \xi^2 \cdot n \cdot \log n$. This yields probability

$$\exp\left(-\frac{2\lambda^2}{\tau(b-a)^2}\right) \leq \exp\left(-\frac{2 \cdot \xi^2 \cdot 10 \cdot \tau \cdot n \cdot \log^2 n}{\tau \cdot 5 \cdot \xi^2 \cdot n \cdot \log n}\right) = \exp(-4 \cdot \log n) = n^{-4}$$

for the event

$$\begin{aligned}&\Delta_w^2(\tau) - r \cdot \tau - (\Delta_w^2(0) - 0) < -\lambda \\ \Leftrightarrow &|\Delta_w(\tau)| < \sqrt{r \cdot \tau + \Delta_w^2(0) - \lambda} \\ \Leftrightarrow &|\Delta_w(\tau)| < \sqrt{r \cdot \tau - \xi \cdot \sqrt{10 \cdot \tau \cdot n} \cdot \log n} \\ \Leftrightarrow &|\Delta_w(\tau)| < \sqrt{r \cdot \alpha \cdot n \cdot \log^2 n - \xi \cdot \sqrt{10 \cdot \alpha} \cdot n \cdot \log^2 n} \\ \Leftrightarrow &|\Delta_w(\tau)| < \sqrt{r \cdot \alpha - \xi \cdot \sqrt{10 \cdot \alpha}} \cdot \sqrt{n} \cdot \log n.\end{aligned}$$

It is now clear that there exists a constant $\alpha = \alpha(\xi, r) > 0$ such that $r \cdot \alpha - \xi \cdot \sqrt{10 \cdot \alpha} \geq \xi^2$. Then, $|\Delta_w(\tau)| \geq \sqrt{r \cdot \alpha - \sqrt{r \cdot \alpha}} \cdot \sqrt{n} \cdot \log n$ implies $|\Delta_w(\tau)| \geq \xi \cdot \sqrt{n \cdot \log n}$, i.e., $T_w \leq \tau$. $\square$

5.4. Bounds on the number of undecided agents

In this section, we show two results. First, we show that from an arbitrary initial state, we quickly reach a configuration $\mathbf{X}(t)$ with

$$\min\{x_1(t), (1-p)x_2(t)\} \leq u(t) \leq x_1(t) + x_2(t).$$

After entering that region, we show that the number of undecided agents will w.h.p. remain in that region for $\Omega(n \log^2 n)$ interactions.

Lemma 5.9. *Let* $\Phi_{up}(t) \coloneqq u(t) - x_1(t) - x_2(t)$ *and* $T_u \coloneqq \inf\{t \geq 0 : \Phi_{up}(t) > 6\xi \cdot \sqrt{n \log n}\}$. *Let* $\boldsymbol{x}(0)$ *be an arbitrary configuration with* $\Phi_{up}(0) < 2\xi \cdot \sqrt{n \log n}$. *Then w.h.p. it holds that* $T_u = \omega(n \log^2 n)$.

Proof. The idea is to show for any time $t > 0$ where $u(t)$ crosses the threshold $x_1(t) + x_2(t) + \Delta$ (for $\Delta := 2\xi \cdot \sqrt{n \cdot \log n}$) from below, it only ever exceeds it by another 2Δ due to a negative drift that we now calculate. We say the t'th interaction is *productive*, if $\mathbf{X}(t+1) \neq \mathbf{X}(t)$. We denote this event by prod_t. Note that any productive interaction alters Φ_{up} by ± 2. Then,

$$\Pr\left[\Phi_{up}(t+1) = \Phi_{up}(t) + 2 \mid \mathcal{F}_t\right] = \frac{(1-p)x_1 \cdot x_2}{n^2} + \frac{x_2 \cdot x_1}{n^2} \qquad = \frac{(2-p)x_1 \cdot x_2}{n^2},$$
$$\Pr\left[\Phi_{up}(t+1) = \Phi_{up}(t) - 2 \mid \mathcal{F}_t\right] = \frac{u \cdot x_1}{n^2} + \frac{u \cdot x_2}{n^2} \qquad = \frac{u \cdot (x_1 + x_2)}{n^2}.$$

Note that for a fixed u, the product $x_1 \cdot x_2$ is maximal for $x_1 = x_2 = (x_1 + x_2)/2$. Then,

$$\Pr\left[\Phi_{up}(t+1) = \Phi_{up}(t) + 2 \mid \mathcal{F}_t, \mathrm{prod}_t\right] = \frac{(2-p)x_1 \cdot x_2}{(2-p)x_1 \cdot x_2 + u \cdot (x_1 + x_2)}$$
$$= \frac{1}{2} + \frac{(2-p)x_1 \cdot x_2 - u \cdot (x_1 + x_2)}{2\left((2-p)x_1 \cdot x_2 + u \cdot (x_1 + x_2)\right)} \leq \frac{1}{2} + \frac{(2-p)\frac{x_1+x_2}{2} \cdot \frac{x_1+x_2}{2} - u \cdot (x_1 + x_2)}{2\left((2-p)x_1 \cdot x_2 + u \cdot (x_1 + x_2)\right)}$$
$$= \frac{1}{2} + \frac{(x_1 + x_2) \cdot \left(\frac{2-p}{4} \cdot (x_1 + x_2) - u\right)}{2\left((2-p)x_1 \cdot x_2 + u \cdot (x_1 + x_2)\right)} \leq \frac{1}{2} - \frac{(x_1 + x_2) \cdot \Phi_{up}(t)}{2\left((2-p)x_1 \cdot x_2 + u \cdot (x_1 + x_2)\right)}$$

So as long as $\Phi_{up}(t) \geq \Delta > 0$ holds at time t, we can bound the denominator by

$$\begin{aligned} &2\left((2-p)x_1 \cdot x_2 + u \cdot (x_1 + x_2)\right) \\ &\quad \leq 2\left(2x_1 \cdot x_2 + u \cdot (x_1 + x_2)\right) \\ &\quad \leq 2((x_1 + x_2)^2 + u \cdot (x_1 + x_2)) \\ &\quad = 2(u + x_1 + x_2) \cdot (x_1 + x_2) \\ &\quad = 2n \cdot (x_1 + x_2). \end{aligned}$$

Therefore, if $\Phi_{up}(t) \geq \Delta$ and we condition on a productive step, we have

$$\Pr\left[\Phi_{up}(t+1) = \Phi_{up} + 2 \mid \mathcal{F}_t, \mathrm{prod}_t\right] \geq \frac{1}{2} - \frac{\Delta \cdot (x_1 + x_2)}{2n \cdot (x_1 + x_2)} \geq \frac{1}{2} - \frac{\Delta}{2n}.$$

Now, we consider a sequence of $m = \omega\big(n \cdot \log^2 n\big)$ interactions. We let T_i for $i > 0$ denote the first time after T_{i-1} where $\Phi_{up}(T_i - 1) < \Delta$ and $\Phi_{up}(T_i) \geq \Delta$. We show that in every sequence of at most m interactions, starting at time T_i, Φ_{up} does not exceed the threshold by more than 3Δ before falling back below Δ with probability at least $1 - n^{-4\xi^2}$. To show that, we consider a sequence of independent Bernoulli trials $(Z_i)_{i>0}$ with success probability $\tilde{p} = 1/2 + \Delta/(2n)$. Each trial corresponds to a productive interaction where $\Phi_{up}(t) \geq \Delta$ and a success to a -2-step in Φ_{up}. The number of failed trials exceeds the number of successful trials by more than Δ trials – and thus Φ_{up} exceeds $\Delta + 2\Delta = 6\xi \cdot \sqrt{n \cdot \log n}$ with probability at most (see Lemma A.18):

$$\left(\frac{1-\tilde{p}}{\tilde{p}}\right)^{\Delta} = \left(1 - \frac{2\Delta}{n + \Delta}\right)^{\Delta} \leq \exp\left(-\frac{2\Delta^2}{n + \Delta}\right) \leq \exp\left(-\frac{2 \cdot 4\xi^2 \cdot n \cdot \log n}{2n}\right) = n^{-4\xi^2}.$$

Finally, every such sequence has at least one interaction; thus, there are at most m sequences in m interactions. The claim then follows from the union bound. $\square$

We complete this section by showing that after creating at least $\min\{x_1, (1-p)x_2\}$ undecided agents, that number does not drop significantly while the weighted bias is not large enough. The proof idea is the same as in Lemma 5.9.

Lemma 5.10. *Let $T = \inf\{t \geq 0 \mid u(t) \geq \min\{x_1(t), (1-p)x_2(t)\}\}$. Let $\mathbf{x}(0)$ be an arbitrary initial configuration with $u(0) \leq x_1(0)+x_2(0)$. Then, $\Pr[T \leq 144 \cdot n] \geq 1-n^{-3}$. Furthermore, as long as $\sqrt{n \log n} \leq \Delta_w(t) \leq (n-u(t))/4$ for $t \geq T$, it holds that $u(t) \geq (3/10) \cdot (1-p)/(2-p) \cdot n$ w.h.p.*

Proof. We track the evolution of the undecided agents over time to show the result. Let $T = \inf\{t \geq 0 \mid u(t) \geq \min\{x_1(t), (1-p)x_2(t)\}/2\}$. We distinguish between $\min\{x_1(t), (1-p)x_2(t)\} = x_1(t)$ and $\min\{x_1(t), (1-p)x_2(t)\} = (1-p)x_2(t)$. In the first case, as long as $t < T$, it holds that

$$\begin{aligned}
\Pr[U(t+1) = u(t)+1 \mid \mathcal{F}_t, \text{prod}_t] &= \frac{1}{2} + \frac{(2-p)x_1x_2 - u\cdot(x_1+x_2)}{2((2-p)x_1x_2 + u\cdot(x_1+x_2))} \\
&\geq \frac{1}{2} + \frac{(2-p)(x_1+x_2) - x_1/2\cdot(x_1+x_2)}{2((2-p)x_1x_2 + u\cdot(x_1+x_2))} \\
&\geq \frac{1}{2} + \frac{x_1((x_1+x_2)/2 - \Delta_w)}{3x_1(x_1+x_2)} \\
&\geq \frac{1}{2} + \frac{1}{12}.
\end{aligned}$$

In the second case, as long as $t < T$, it holds that

$$\begin{aligned}
\Pr[U(t+1) = u(t)+1 \mid \mathcal{F}_t, \text{prod}_t] &= \frac{1}{2} + \frac{(2-p)x_1x_2 - u\cdot(x_1+x_2)}{2((2-p)x_1x_2 + u\cdot(x_1+x_2))} \\
&\geq \frac{1}{2} + \frac{(2-p)(x_1+x_2) - (1-p)/2x_2\cdot(x_1+x_2)}{2((2-p)x_1x_2 + u\cdot(x_1+x_2))} \\
&\geq \frac{1}{2} + \frac{x_2((1-p)(x_1+x_2)/2 - \Delta_w)}{3x_2(x_1+x_2)} \\
&\geq \frac{1}{2} + \frac{1}{12}.
\end{aligned}$$

Now, we bound the number of productive interactions in the first $\mathrm{O}(n)$ interactions. An interaction at time t is productive with at least constant probability

$$\frac{(2-p)x_1(t)x_2(t) + u(t)(n-u(t))}{n^2} \geq \varepsilon$$

due to Lemma 5.9 and the assumption $\Delta_w(t) \leq (n-u(t))/4$. From Chernoff bound, it follows $\Omega(n)$ productive interactions in $\mathrm{O}(n)$ interactions w.h.p. Now we show that the number of undecided agents does not drop below $\min\{x_1(t)/2, (1-p)x_2/2\} - \Delta$

with $\Delta = \sqrt{n \log n}$. This allows us to relate this sequence of productive interactions and the evolution of the undecided agents to a biased random walk on $\mathbb{N}_0$ with a reflective barrier at position 0. The probability of moving to the right is $\tilde{p} = 1/2 + 1/12$ and the probability of moving to the left is $1 - \tilde{p}$. Then Lemma A.19 implies $T = \mathrm{O}\big(n\big)$ w.h.p.

Now, we show that the number of undecided agents remains relatively large. Similar to Lemma 5.9, we consider the time steps T_i where for the first time after T_{i-1} we have $u(T_i - 1) > \min\{ x_1(T_{i-1})/2, (1-p)x_2(T_{i-1})/2 \}$ and $u(T_i) \leq \min\{ x_1(T_i)/2, (1-p)x_2(T_i)/2 \}$. We do a case study on $u(t)$. As long as $u(t) \leq x_1(t)/2$ for $t \geq T_i$ it holds that

$$
\begin{aligned}
\Pr\left[U(t+1) = u(t) - 1 \mid \mathcal{F}_t, \mathrm{prod}_t\right] &= \frac{1}{2} - \frac{(2-p)x_1x_2 - u \cdot (x_1 + x_2)}{2((2-p)x_1x_2 + u \cdot (x_1 + x_2))} \\
&\leq \frac{1}{2} - \frac{(2-p)(x_1 + x_2) - x_1/2 \cdot (x_1 + x_2)}{2((2-p)x_1x_2 + u \cdot (x_1 + x_2))} \\
&\leq \frac{1}{2} - \frac{x_1((x_1 + x_2)/2 - \Delta_w)}{3x_1(x_1 + x_2)} \\
&\leq \frac{1}{2} - \frac{1+p}{12} \\
&\leq \frac{1}{2} - \frac{1-p}{12}
\end{aligned}
$$

where in the last inequalities we use $\Delta_w < (1-p)(x_1 + x_2)/4$. Similarly, as long as $u(t) \leq (1-p)x_2(t)/2$ for $t \geq T_i$ it holds that

$$
\begin{aligned}
\Pr\left[U(t+1) = u(t) - 1 \mid \mathcal{F}_t, \mathrm{prod}_t\right] &= \frac{1}{2} - \frac{(2-p)x_1x_2 - u \cdot (x_1 + x_2)}{2((2-p)x_1x_2 + u \cdot (x_1 + x_2))} \\
&\leq \frac{1}{2} - \frac{(2-p)(x_1 + x_2) - (1-p)/2x_2 \cdot (x_1 + x_2)}{2((2-p)x_1x_2 + u \cdot (x_1 + x_2))} \\
&\leq \frac{1}{2} - \frac{x_2((1-p)(x_1 + x_2)/2 - \Delta_w)}{3x_2(x_1 + x_2)} \\
&\leq \frac{1}{2} - \frac{(1-p)(x_1 + x_2)/4)}{3(x_1 + x_2)} \\
&\leq \frac{1}{2} - \frac{1-p}{12}
\end{aligned}
$$

where in the last inequalities we use $\Delta_w < (1-p)(x_1 + x_2)/4$. Now we show that the number of undecided agents does not drop below $\min\{ x_1(t)/2, (1-p)x_2/2 \} - \Delta$ with $\Delta = \sqrt{n \log n}$. We consider a sequence of independent Bernoulli trials $(Z_i)_{i \geq 0}$ with success probability $\tilde{p} = 1/2 + (1-p)/12$. Each trial corresponds to a productive interaction under the assumption from above. From Lemma A.18, it follows the number of failed trials exceeds the number of successful trials by more than $\Delta/2$ trials with probability at most

$$
\left(\frac{1-\tilde{p}}{\tilde{p}}\right)^{\Delta/2} = \left(\frac{6-(1-p)}{6+(1-p)}\right)^{\Delta/2}
$$

At last, from $u \geq x_1/2 - \Delta$, $x_1 - (1-p)x_2 > 0$ and $n = x_1 + x_2 + u$ it follows $u \geq (3/10) \cdot (1-p)/(2-p) \cdot n$. On the other hand, from $u \geq (1-p)x_2/2 - \Delta$, $x_1 - (1-p)x_2 > 0$ and $n = x_1 + x_2 + u$ it follows $u \geq (3/10) \cdot (1-p)/(2-p) \cdot n$. The claim then follows from the union bound. $\square$

Chapter 6.

Conclusion

In this thesis, we show fast convergence of plurality consensus problems in the population protocol model. We demonstrate that the convergence rates depend on the magnitude of support of the initial largest opinion and the type of bias in the initial configuration. In fact, we especially observe a typical tradeoff behavior between time and space with variants of the Undecided State Dynamics.

Regarding the synchronized Undecided State Dynamics, one open question is whether our result is tight. The main reason for a running time of $\mathrm{O}(\log^2 n)$ is that our algorithm needs $\mathrm{O}(\log n)$ phases of length $\mathrm{O}(\log n)$ for breaking the ties in the case of several opinions with roughly the same support. It might be possible to work with a phase length as a function of k, resulting in a refined running time of $\mathrm{O}(\log k \log n)$. Moreover, it may be possible to interleave consecutive phases in order to reduce the running time even further. Regarding the (unsynchronized) Undecided State Dynamics, it remains open to prove convergence of the $k > 2$ opinion USD with *no initial bias* in the gossip model and, moreover, to understand whether a unified framework exists for analyzing the process in both models simultaneously.

Separately, we leave as future work analyzing the k-opinion USD in the presence of adversarial nodes or communication noise. Recent results of d'Amore et al. [42, 44] and of Cruciani et al. [40], which analyze the 2-state USD process (as well as other majority dynamics for $k > 2$) under such settings, suggest that the k-opinion USD is also robust to these noise models. Quantifying the effect of such noise on the convergence rate of the k-opinion USD is thus an interesting open question.

Regarding the Undecided State Dynamics with stubborn agents, it remains an open question whether a tight runtime of $\mathrm{O}(n \log n)$ is also achievable in the hard regime. An adaption to the synchronous gossip model is an interesting question as well. We believe that our model would show the same behavior in the parallel gossip model. Interestingly, we believe this is not the case for a variant where undecided agents can also be stubborn. This is due to the fact that undecided agents have different effects in sequential and parallel models. A generalization for more than two opinions seems natural. Describing the phase transition for this case looks non-trivial. It has certain similarities to measure the convergence time of consensus protocols.

Regarding the exact plurality consensus problem, our protocols use majority, leader election, and junta election protocols as black boxes. Improving the guarantees of these black boxes would also carry over to our protocols. For example, a leader election protocol that has a *with high probability* runtime of $\mathrm{O}(\log n)$ would immediately improve our runtime. Similarly, a constant state majority protocol and a constant state junta election protocol (that works *with high probability*) would immediately improve our state space bounds. Furthermore, we believe that $\Omega(n/x_{\max})$ is a natural lower bound for the runtime, and thus, the possible improvements mentioned above would lead to a state- and time-optimal exact plurality consensus protocol.

In our main result we prune small opinions in order to reduce the number of tournaments. We conjecture that this yields almost optimal protocols. We believe that additional techniques are required to further improve the runtime (possibly at the expense of slightly increased state complexity). In particular, it would be interesting to find another, more efficient way than pairwise comparison of opinions via tournaments to identify the plurality opinion.

Bibliography

[1] D. Alistarh, J. Aspnes, and R. Gelashvili. "Space-Optimal Majority in Population Protocols". In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*. Ed. by A. Czumaj. SIAM, 2018, pp. 2221–2239. DOI: 10.1137/1.9781611975031.144.

[2] D. Alistarh, B. Dudek, A. Kosowski, D. Soloveichik, and P. Uznanski. "Robust Detection in Leak-Prone Population Protocols". In: *DNA Computing and Molecular Programming - 23rd International Conference, DNA*. Vol. 10467. Lecture Notes in Computer Science. Springer, 2017, pp. 155–171. DOI: 10.1007/978-3-319-66799-7_11.

[3] D. Alistarh, R. Gelashvili, and M. Vojnovic. "Fast and Exact Majority in Population Protocols". In: *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*. Ed. by C. Georgiou and P. G. Spirakis. ACM, 2015, pp. 47–56. DOI: 10.1145/2767386.2767429.

[4] D. Alistarh, M. Töpfer, and P. Uznanski. "Comparison Dynamics in Population Protocols". In: *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*. Ed. by A. Miller, K. Censor-Hillel, and J. H. Korhonen. ACM, 2021, pp. 55–65. DOI: 10.1145/3465084.3467915.

[5] T. Amir, J. Aspnes, P. Berenbrink, F. Biermeier, C. Hahn, D. Kaaser, and J. Lazarsfeld. "Fast Convergence of k-Opinion Undecided State Dynamics in the Population Protocol Model". In: *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, PODC 2023, Orlando, FL, USA, June 19-23, 2023*. Ed. by R. Oshman, A. Nolin, M. M. Halldórsson, and A. Balliu. ACM, 2023, pp. 13–23. DOI: 10.1145/3583668.3594589.

[6] T. Amir, J. Aspnes, and J. Lazarsfeld. "Approximate Majority with Catalytic Inputs". In: *24th International Conference on Principles of Distributed Systems, OPODIS 2020, December 14-16, 2020, Strasbourg, France (Virtual Conference)*. Ed. by Q. Bramas, R. Oshman, and P. Romano. Vol. 184. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 19:1–19:16. DOI: 10.4230/LIPIcs.OPODIS.2020.19.

[7] A. Anagnostopoulos, L. Becchetti, E. Cruciani, F. Pasquale, and S. Rizzo. "Biased opinion dynamics: when the devil is in the details". In: *Inf. Sci.* 593 (2022), pp. 49–63. DOI: 10.1016/J.INS.2022.01.072.

[8] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. "Computation in networks of passively mobile finite-state sensors". In: *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Distributed Computing, PODC 2004*. Ed. by S. Chaudhuri and S. Kutten. ACM, 2004, pp. 290–299. DOI: `10.1145/1011767.1011810`.

[9] D. Angluin, J. Aspnes, and D. Eisenstat. "A simple population protocol for fast robust approximate majority". In: *Distributed Computing* 21.2 (2008), pp. 87–102. DOI: `10.1007/s00446-008-0059-z`.

[10] D. Angluin, J. Aspnes, and D. Eisenstat. "Fast computation by population protocols with a leader". In: *Distributed Comput.* 21.3 (2008), pp. 183–199. DOI: `10.1007/s00446-008-0067-z`.

[11] D. Angluin, J. Aspnes, and D. Eisenstat. "Stably computable predicates are semilinear". In: *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing, PODC*. ACM, 2006, pp. 292–299. DOI: `10.1145/1146381.1146425`.

[12] A. Antoniadis, F. Biermeier, A. Cristi, C. Damerius, R. Hoeksma, D. Kaaser, P. Kling, and L. Nölke. "On the Complexity of Anchored Rectangle Packing". In: *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*. Ed. by M. A. Bender, O. Svensson, and G. Herman. Vol. 144. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 8:1–8:14. DOI: `10.4230/LIPICS.ESA.2019.8`.

[13] R. B. Ash. *Information Theory*. Dover books on advanced mathematics. Dover Publications, 1990.

[14] H. Attiya and J. Welch. *Distributed computing: fundamentals, simulations, and advanced topics*. Vol. 19. John Wiley & Sons, 2004.

[15] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling. "Fast Consensus via the Unconstrained Undecided State Dynamics". In: *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. SIAM, 2022, pp. 3417–3429. DOI: `10.1137/1.9781611977073.135`.

[16] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling. "Population Protocols for Exact Plurality Consensus: How a small chance of failure helps to eliminate insignificant opinions". In: *PODC '22: ACM Symposium on Principles of Distributed Computing*. ACM, 2022, pp. 224–234. DOI: `10.1145/3519270.3538447`.

[17] G. Bankhamer, R. Elsässer, D. Kaaser, and M. Krnc. "Positive Aging Admits Fast Asynchronous Plurality Consensus". In: *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*. See `https://arxiv.org/abs/1806.02596` for the full version. ACM, 2020, pp. 385–394. DOI: `10.1145/3382734.3406506`.

[18] L. Becchetti, A. Clementi, A. Korman, F. Pasquale, L. Trevisan, and R. Vacus. "On the Role of Memory in Robust Opinion Dynamics". In: *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*. ijcai.org, 2023, pp. 29–37. DOI: `10.24963/IJCAI.2023/4`.

[19] L. Becchetti, A. E. F. Clementi, and E. Natale. "Consensus Dynamics: An Overview". In: *SIGACT News* 51.1 (2020), pp. 58–104. DOI: `10.1145/3388392.3388403`.

[20] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, and R. Silvestri. "Plurality Consensus in the Gossip Model". In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*. Ed. by P. Indyk. SIAM, 2015, pp. 371–390. DOI: `10.1137/1.9781611973730.27`.

[21] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, R. Silvestri, and L. Trevisan. "Simple dynamics for plurality consensus". In: *Distributed Computing* 30.4 (2017), pp. 293–306. DOI: `10.1007/s00446-016-0289-4`.

[22] S. Ben-Nun, T. Kopelowitz, M. Kraus, and E. Porat. "An $O(\log^{3/2} n)$ Parallel Time Population Protocol for Majority with $O(\log n)$ States". In: *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event*. Ed. by Y. Emek and C. Cachin. ACM, 2020, pp. 191–199. DOI: `10.1145/3382734.3405747`.

[23] P. Berenbrink, F. Biermeier, and C. Hahn. *Undecided State Dynamics with Stubborn Agents*. 2024. arXiv: `2406.07335 [cs.DC]`.

[24] P. Berenbrink, F. Biermeier, C. Hahn, and D. Kaaser. "Loosely-Stabilizing Phase Clocks and The Adaptive Majority Problem". In: *1st Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2022, March 28-30, 2022, Virtual Conference*. Ed. by J. Aspnes and O. Michail. Vol. 221. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 7:1–7:17. DOI: `10.4230/LIPICS.SAND.2022.7`.

[25] 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: *Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC 2017, Washington, DC, USA, July 25-27, 2017*. Ed. by E. M. Schiller and A. A. Schwarzmann. ACM, 2017, pp. 335–344. DOI: `10.1145/3087801.3087817`.

[26] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. "Balanced Allocations: The Heavily Loaded Case". In: *SIAM J. Comput.* 35.6 (2006), pp. 1350–1385. DOI: `10.1137/S009753970444435X`.

[27] P. Berenbrink, R. Elsässer, T. Friedetzky, D. Kaaser, P. Kling, and T. Radzik. "A Population Protocol for Exact Majority with $O(\log^{5/3} n)$ Stabilization Time and $\Theta(\log n)$ States". In: *32nd International Symposium on Distributed Computing, DISC 2018*. Ed. by U. Schmid and J. Widder. Vol. 121. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 10:1–10:18. DOI: `10.4230/LIPIcs.DISC.2018.10`.

[28] P. Berenbrink, R. Elsässer, T. Friedetzky, D. Kaaser, P. Kling, and T. Radzik. "Time-space trade-offs in population protocols for the majority problem". In: *Distributed Comput.* 34.2 (2021), pp. 91–111. DOI: 10.1007/s00446-020-00385-0.

[29] P. Berenbrink, T. Friedetzky, G. Giakkoupis, and P. Kling. "Efficient Plurality Consensus, Or: the Benefits of Cleaning up from Time to Time". In: *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy.* Ed. by I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi. Vol. 55. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, 136:1–136:14. DOI: 10.4230/LIPIcs.ICALP.2016.136.

[30] P. Berenbrink, T. Friedetzky, D. Kaaser, and P. Kling. "Tight & Simple Load Balancing". In: *2019 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2019.* IEEE, 2019, pp. 718–726. DOI: 10.1109/IPDPS.2019.00080.

[31] P. Berenbrink, G. Giakkoupis, A.-M. Kermarrec, and F. Mallmann-Trenn. "Bounds on the Voter Model in Dynamic Networks". In: *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy.* Ed. by I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi. Vol. 55. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, 146:1–146:15. DOI: 10.4230/LIPIcs.ICALP.2016.146.

[32] P. Berenbrink, G. Giakkoupis, and P. Kling. "Optimal time and space leader election in population protocols". In: *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020.* Ed. by K. Makarychev, Y. Makarychev, M. Tulsiani, G. Kamath, and J. Chuzhoy. ACM, 2020, pp. 119–129. DOI: 10.1145/3357713.3384312.

[33] S. P. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. "Randomized gossip algorithms". In: *IEEE Transactions on Information Theory* 52.6 (2006), pp. 2508–2530. DOI: 10.1109/TIT.2006.874516.

[34] 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: *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK.* Ed. by I. Potapov, P. G. Spirakis, and J. Worrell. Vol. 117. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 28:1–28:15. DOI: 10.4230/LIPIcs.MFCS.2018.28.

[35] A. Condon, M. Hajiaghayi, D. G. Kirkpatrick, and J. Manuch. "Simplifying Analyses of Chemical Reaction Networks for Approximate Majority". In: *DNA Computing and Molecular Programming - 23rd International Conference, DNA 23.* Vol. 10467. Lecture Notes in Computer Science. Springer, 2017, pp. 188–209. DOI: 10.1007/978-3-319-66799-7_13.

[36] C. Cooper, R. Elsässer, H. Ono, and T. Radzik. "Coalescing random walks and voting on graphs". In: *ACM Symposium on Principles of Distributed Computing, PODC '12, Funchal, Madeira, Portugal, July 16-18, 2012.* Ed. by D. Kowalski and A. Panconesi. ACM, 2012, pp. 47–56. DOI: 10.1145/2332432.2332440.

[37] C. Cooper, R. Elsässer, and T. Radzik. "The Power of Two Choices in Distributed Voting". In: *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II.* Ed. by J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias. Vol. 8573. Lecture Notes in Computer Science. Springer, 2014, pp. 435–446. DOI: 10.1007/978-3-662-43951-7_37.

[38] C. Cooper, R. Elsässer, T. Radzik, N. Rivera, and T. Shiraga. "Fast Consensus for Voting on General Expander Graphs". In: *Distributed Computing - 29th International Symposium, DISC 2015, Tokyo, Japan, October 7-9, 2015, Proceedings.* Ed. by Y. Moses. Vol. 9363. Lecture Notes in Computer Science. Springer, 2015, pp. 248–262. DOI: 10.1007/978-3-662-48653-5_17.

[39] C. Cooper, T. Radzik, N. Rivera, and T. Shiraga. "Fast Plurality Consensus in Regular Expanders". In: *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria.* Ed. by A. W. Richa. Vol. 91. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 13:1–13:16. DOI: 10.4230/LIPIcs.DISC.2017.13.

[40] E. Cruciani, H. A. Mimun, M. Quattropani, and S. Rizzo. "Phase transition of the k-majority dynamics in biased communication models". In: *Distributed Comput.* 36.2 (2023), pp. 107–135. DOI: 10.1007/S00446-023-00444-2.

[41] E. Cruciani, E. Natale, A. Nusser, and G. Scornavacca. "Phase transition of the 2-Choices dynamics on core-periphery networks". In: *Distributed Comput.* 34.3 (2021), pp. 207–225. DOI: 10.1007/S00446-021-00396-5.

[42] F. D'Amore, A. E. F. Clementi, and E. Natale. "Phase Transition of a Non-linear Opinion Dynamics with Noisy Interactions - (Extended Abstract)". In: *Structural Information and Communication Complexity - 27th International Colloquium, SIROCCO 2020, Paderborn, Germany, June 29 - July 1, 2020, Proceedings.* Ed. by A. W. Richa and C. Scheideler. Vol. 12156. Lecture Notes in Computer Science. Springer, 2020, pp. 255–272. DOI: 10.1007/978-3-030-54921-3_15.

[43] F. D'Amore, A. E. F. Clementi, and E. Natale. "Phase transition of a nonlinear opinion dynamics with noisy interactions". In: *Swarm Intell.* 16.4 (2022), pp. 261–304. DOI: 10.1007/S11721-022-00217-W.

[44] F. D'Amore and I. Ziccardi. "Phase Transition of the 3-Majority Dynamics with Uniform Communication Noise". In: *Structural Information and Communication Complexity - 29th International Colloquium, SIROCCO 2022, Paderborn, Germany, June 27-29, 2022, Proceedings.* Ed. by M. Parter. Vol. 13298. Lecture Notes in Computer Science. Springer, 2022, pp. 98–115. DOI: 10.1007/978-3-031-09993-9_6.

[45] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. "Stabilizing Consensus with the Power of Two Choices". In: *Algorithmic Methods for Distributed Cooperative Systems, 06.09. - 11.09.2009.* Ed. by S. P. Fekete, S. Fis-

cher, M. A. Riedmiller, and S. Suri. Vol. 09371. Dagstuhl Seminar Proceedings. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2009.

[46] B. Doerr, L. A. Goldberg, L. Minder, T. Sauerwald, and C. Scheideler. "Stabilizing consensus with the power of two choices". In: *SPAA 2011: Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures, San Jose, CA, USA, June 4-6, 2011 (Co-located with FCRC 2011)*. Ed. by R. Rajaraman and F. M. auf der Heide. ACM, 2011, pp. 149–158. DOI: 10.1145/1989493.1989516.

[47] D. Doty, M. Eftekhari, L. Gasieniec, E. E. Severson, P. Uznanski, and G. Stachowiak. "A time and space optimal stable population protocol solving exact majority". In: *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*. IEEE, 2021, pp. 1044–1055. DOI: 10.1109/FOCS52979.2021.00104.

[48] B. Dudek and A. Kosowski. "Universal protocols for information dissemination using emergent signals". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*. Ed. by I. Diakonikolas, D. Kempe, and M. Henzinger. ACM, 2018, pp. 87–99. DOI: 10.1145/3188745.3188818.

[49] W. Feller. *An introduction to probability theory and its applications*. 3rd edition, revised printing. Wiley series in probability and mathematical statistics: Probability and mathematical statistics. Wiley, 1968.

[50] L. Gasieniec, D. D. Hamilton, R. Martin, P. G. Spirakis, and G. Stachowiak. "Deterministic Population Protocols for Exact Majority and Plurality". In: *20th International Conference on Principles of Distributed Systems, OPODIS 2016, December 13-16, 2016, Madrid, Spain*. Ed. by P. Fatourou, E. Jiménez, and F. Pedone. Vol. 70. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, 14:1–14:14. DOI: 10.4230/LIPIcs.OPODIS.2016.14.

[51] L. Gasieniec and G. Stachowiak. "Enhanced Phase Clocks, Population Protocols, and Fast Space Optimal Leader Election". In: *J. ACM* 68.1 (2021), 2:1–2:21. DOI: 10.1145/3424659.

[52] M. Ghaffari and J. Lengler. "Nearly-Tight Analysis for 2-Choice and 3-Majority Consensus Dynamics". In: *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018, Egham, United Kingdom, July 23-27, 2018*. Ed. by C. Newport and I. Keidar. ACM, 2018, pp. 305–313.

[53] M. Ghaffari and M. Parter. "A Polylogarithmic Gossip Algorithm for Plurality Consensus". In: *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*. Ed. by G. Giakkoupis. ACM, 2016, pp. 117–126. DOI: 10.1145/2933057.2933097.

[54] S. Greenberg and M. Mohri. "Tight Lower Bound on the Probability of a Binomial Exceeding its Expectation". In: *CoRR* abs/1306.1433 (2013). arXiv: 1306.1433.

[55] Y. Hassin and D. Peleg. "Distributed Probabilistic Polling and Applications to Proportionate Agreement". In: *Inf. Comput.* 171.2 (2001), pp. 248–268. DOI: 10.1006/inco.2001.3088.

[56] W. Hoeffding. "Probability Inequalities for Sums of Bounded Random Variables". In: *Journal of the American Statistical Association* 58.301 (1963), pp. 13–30. DOI: 10.1080/01621459.1963.10500830.

[57] S. Janson. "Tail bounds for sums of geometric and exponential variables". In: *Statistics & Probability Letters* 135 (2018), pp. 1–6. DOI: 10.1016/j.spl.2017.11.017.

[58] S. Janson, T. Luczak, and A. Rucinski. *Random graphs.* Wiley-Interscience series in discrete mathematics and optimization. Wiley, 2000. DOI: 10.1002/9781118032718.

[59] V. Kanade, F. Mallmann-Trenn, and T. Sauerwald. "On coalescence time in graphs: When is coalescing as fast as meeting?: Extended Abstract". In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019.* Ed. by T. M. Chan. SIAM, 2019, pp. 956–965. DOI: 10.1137/1.9781611975482.59.

[60] P. N. Klein and N. E. Young. "On the Number of Iterations for Dantzig-Wolfe Optimization and Packing-Covering Approximation Algorithms". In: *SIAM J. Comput.* 44.4 (2015), pp. 1154–1172. DOI: 10.1137/12087222X.

[61] A. Kosowski and P. Uznanski. "Brief Announcement: Population Protocols Are Fast". In: *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018.* Ed. by C. Newport and I. Keidar. ACM, 2018, pp. 475–477.

[62] A. D. Kshemkalyani and M. Singhal. *Distributed computing: principles, algorithms, and systems.* Cambridge University Press, 2011.

[63] J. Lengler. "Drift Analysis". In: *Theory of Evolutionary Computation - Recent Developments in Discrete Optimization.* Ed. by B. Doerr and F. Neumann. Natural Computing Series. Springer, 2020, pp. 89–131. DOI: 10.1007/978-3-030-29414-4_2.

[64] J. Lengler and A. Steger. "Drift analysis and evolutionary algorithms revisited". In: *Combinatorics, Probability and Computing* 27.4 (2018), pp. 643–666.

[65] D. Levin and Y. Peres. *Markov Chains and Mixing Times.* 2nd edition. American Mathematical Society, 2017.

[66] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press, 2005. DOI: 10.1017/CBO9780511813603.

[67] M. Mobilia. "Does a single zealot affect an infinite group of voters?" In: *Physical review letters* 91.2 (2003), p. 028701.

[68] M. Mobilia, A. Petersen, and S. Redner. "On the role of zealotry in the voter model". In: *Journal of Statistical Mechanics: Theory and Experiment* 2007.08 (2007), P08029.

[69] Y. Mocquard, F. Robin, B. Sericola, and E. Anceaume. "Stochastic analysis of average-based distributed algorithms". In: *J. Appl. Probab.* 58.2 (2021), pp. 394–410. DOI: 10.1017/jpr.2020.97.

[70] A. Mukhopadhyay. "Phase transitions in biased opinion dynamics with 2-choices rule". In: *Probability in the Engineering and Informational Sciences* (2023), pp. 1–18. DOI: 10.1017/S0269964823000098.

[71] T. Nakata, H. Imahayashi, and M. Yamashita. "A probabilistic local majority polling game on weighted directed graphs with an application to the distributed agreement problem". In: *Networks* 35.4 (2000), pp. 266–273. DOI: 10.1002/1097-0037(200007)35:4%3C266::AID-NET5%3E3.0.CO;2-4.

[72] E. Natale and I. Ramezani. "On the Necessary Memory to Compute the Plurality in Multi-agent Systems". In: *Algorithms and Complexity - 11th International Conference, CIAC 2019, Rome, Italy, May 27-29, 2019, Proceedings.* Ed. by P. Heggernes. Vol. 11485. Lecture Notes in Computer Science. Springer, 2019, pp. 323–338. DOI: 10.1007/978-3-030-17402-6_27.

[73] F. Neumann, ed. *Theory of Evolutionary Computation Recent Developments in Discrete Optimization.* 1st ed. 2020. Natural Computing Series: Springer eBooks: Computer Science. Springer, 2020, p. 506.

[74] Y. Peres, K. Talwar, and U. Wieder. "Graphical balanced allocations and the $(1+\beta)$-choice process". In: *Random Struct. Algorithms* 47.4 (2015), pp. 760–775. DOI: 10.1002/rsa.20558.

[75] E. Perron, D. Vasudevan, and M. Vojnovic. "Using Three States for Binary Consensus on Complete Graphs". In: *INFOCOM 2009. 28th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 19-25 April 2009, Rio de Janeiro, Brazil.* IEEE, 2009, pp. 2527–2535. DOI: 10.1109/INFCOM.2009.5062181.

[76] G. Schoenebeck and F.-Y. Yu. "Consensus of Interacting Particle Systems on Erdös-Rényi Graphs". In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018.* Ed. by A. Czumaj. SIAM, 2018, pp. 1945–1964. DOI: 10.1137/1.9781611975031.127.

[77] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. "Binary opinion dynamics with stubborn agents". In: *ACM Transactions on Economics and Computation (TEAC)* 1.4 (2013), pp. 1–30.

Appendix A.

Appendix

A.1. Concentration Results

Theorem A.1 ([66], Theorem 4.4, 4.5)**.** *Let $X_1, \ldots, X_n$ be independent Poisson trials with $\Pr[X_i = 1] = p_i$ and let $X = \sum X_i$ with $\mathbb{E}[X] = \mu$. Then the following Chernoff bounds hold:*
For $0 < \delta' \leq 1$:

$$\Pr[X > (1 + \delta')\mu] \leq e^{-\mu\delta'^2/3}.$$

For $0 < \delta' < 1$:

$$\Pr[X < (1 - \delta')\mu] \leq e^{-\mu\delta'^2/2},$$

Theorem A.2 (General Chernoff upper Bound)**.** *Let $X_1 \cdots, X_n$ be independent 0-1 random variables. Let $X = \sum_{i=1}^n X_i$ and $\mu_u \geq 0$ such that $\mathbb{E}[X] \leq \mu_u$. Then, for any $\delta' > 0$*

$$\Pr[X \geq (1 + \delta') \cdot \mu_u] \leq e^{-\frac{\delta'^2 \cdot \mu_u}{2+\delta'}}.$$

Lemma A.3 (Super-exponential Chernoff Bound)**.** *Let $X_1, ..., X_n$ be n independent random variables taking value in $\{0, 1\}$ and $X = \sum_{i=1}^n X_i$. Then, for $\mathbb{E}[X] = \mu$ and $\delta > 0$ it holds that*

$$\Pr\left[X > \mu + \delta\sqrt{\mu}\right] < \exp(-c \cdot \delta^2) \qquad \textit{for } \delta^2 \leq \mu$$
$$\Pr\left[X > \mu + \frac{\delta^2}{1 + \ln(\frac{\delta^2}{\mu})}\right] < \exp(-c \cdot \delta^2) \qquad \textit{for } \delta^2 > \mu$$

where $c > 0$ is a universal constant.

Proof. Let X be defined as in the lemmas statement. From the Chernoff bound we have for $\lambda > 0$ that $\Pr[X > \mu(1 + \lambda)] < \exp(-\min\{\lambda, \lambda^2\} \cdot \mu/3)$. It is easy to see that, for fitting constants $c_1, c_2 > 0$, this implies for $\delta > 0$ that

$$\Pr\left[X > \mu + \delta\sqrt{\mu}\right] < \exp(-c_1 \cdot \delta^2) \text{ if } \delta^2 \leq \mu, \text{ and} \tag{A.1}$$
$$\Pr\left[X > \mu + \delta^2/7\right] < \exp(-c_2 \cdot \delta^2) \text{ if } \delta^2 > \mu. \tag{A.2}$$

Inequality (A.1) corresponds directly to the first statement of the lemma. However, observe that (A.2) only implies the second desired inequality in case $\mu < \delta^2 \leq \mu \cdot e^6$. This results from the fact that $\delta^2/7 \leq \delta^2/(1 + \ln(\delta^2/\mu))$ in this setting.

In order to tackle the case of $\delta^2 > \mu \cdot e^6$, we employ a different version of the Chernoff bound, which is tighter for large values of δ. That is, by inequality (1.10.8) of [73] we have for $\lambda > 0$ that $Pr[X > \mu(1+\lambda)] < (e/\lambda)^{(\lambda \cdot \mu)}$. We now define $y = \ln(\delta^2/\mu) > 6$ and set $\lambda = \delta^2(1+y)^{-1}\mu^{-1} = e^y(1+y)^{-1}$. This way, we get

$$\Pr\left[X > \mu + \frac{\delta^2}{1+\ln(\delta^2/\mu)}\right] < \left(\frac{e(1+y)}{e^y}\right)^{\frac{\delta^2}{(1+y)}}. \tag{A.3}$$

For $y \geq 6$, it holds that $(1+y) < e^{y/2-1}$. This implies that $e(1+y)e^{-y} < e^{-y/2}$ and allows us to upper-bound the term on the right-hand side of (A.3) as follows

$$\left(\frac{e(1+y)}{e^y}\right)^{\frac{\delta^2}{(1+y)}} < e^{-\frac{y}{2}\frac{\delta^2}{1+y}} < e^{-\frac{\delta^2}{4}}.$$

The results follow when setting $c = \min\{c_1, c_2, 1/4\}$. □

Next we consider tail bounds for sums of geometrically distributed random variables.

Theorem A.4 ([57, Theorem 2.1])**.** *Let $X = \sum_{i=1}^n X_i$ where $X_i, i = 1, \ldots, n$, are independent geometric random variables with $X_i \sim Geo(p_i)$ for $p_i \in (0,1]$. For any $\lambda \geq 1$,*

$$\Pr\left[X \geq \lambda \cdot \mathbb{E}\left[X\right]\right] \leq \exp(-\min_i\{p_i\} \cdot \mathbb{E}\left[X\right] \cdot (\lambda - 1 - \ln\lambda)).$$

Lemma A.5 (Full version of [46])**.** *Suppose that $X_1, \ldots, X_n$ are independent random variables on $\mathbb{N}$, such that there is a constant $\gamma > 0$ with $\Pr\left[X_i = k\right] \leq \gamma(1-\delta)^{k-1}$ for every $k \in \mathbb{N}$. Let $X = \sum_{i=1}^n X_i$, $\mu = \mathbb{E}\left[X\right]$. Then it holds for all $\varepsilon > 0$ that*

$$\Pr\left[X \geq (1+\varepsilon)\mu + O(n)\right] \leq e^{-\frac{\varepsilon^2 n}{2(1+\varepsilon)}}.$$

Theorem A.6 (rephrased, based on [56])**.** *Let $X_1, X_2, \ldots, X_n$ be independent random variables and $a_i \leq X_i \leq b_i$ $(i = 1, 2, \ldots, n)$, then for $\lambda > 0$*

$$\Pr\left[\sum_{i=1}^n X_i - \mathbb{E}\left[\sum_{i=1}^n X_i\right] \geq \lambda\right] \leq e^{-\frac{2\lambda^2}{\sum_{i=1}^n (b_i - a_i)^2}}.$$

Lemma A.7. *Consider a sequence of τ random variables $Z_1, \ldots, Z_\tau$ w.r.t. a sequence of random vectors $\mathbf{X}(0), \ldots, \mathbf{X}(\tau-1)$. Let $\mathbf{X}_{<i} = \{\mathbf{X}(0), \ldots, \mathbf{X}(i-1)\}$ for all $i \leq \tau$. Let $Z = \sum_{i=1}^\tau Z_i$ and $\mu = \sum_{i=1}^\tau \mu_i$ with $\mu_i = E[Z_i \mid X_{<i}]$ for $i \leq \tau$. Assume $a \leq Z_i \leq b$ for all $i \leq \tau$. Then for all $\lambda > 0$*

$$\Pr[Z - \mu < -\lambda] \leq e^{-\frac{2\lambda^2}{\tau(b-a)^2}}$$

Proof. We follow the standard proof technique for Hoeffding bounds. For any $t < 0$ we have

$$\Pr[Z - \mu < -\lambda] = \Pr\left[e^{t(Z-\mu)} > e^{-t\lambda}\right] \leq \frac{E\left[e^{t(Z-\mu)}\right]}{e^{-t\lambda}}$$

where in the last inequality we apply Markov's inequality. First we consider the term $E\left[e^{t(Z-\mu)}\right]$. Since we do not assume any independence among the Z_i's we utilize the concept of conditional independence via the law of total expectation. That is,

$$\begin{aligned}
E\left[e^{t(Z-\mu)}\right] &= E\left[e^{\sum_{i=1}^{\tau} t(Z_i-\mu_i)}\right] \\
&= E\left[E\left[e^{\sum_{i=1}^{\tau} t(Z_i-\mu_i)} \mid \mathbf{X}_{<\tau}\right]\right] \\
&= E\left[e^{\sum_{i=1}^{\tau-1} t(Z_i-\mu_i)} \cdot E\left[e^{t(Z_\tau-\mu_\tau)} \mid \mathbf{X}_{<\tau}\right]\right] \\
&= E\left[e^{\sum_{i=1}^{\tau-1} t(Z_i-\mu_i)}\right] \cdot E\left[e^{t(Z_\tau-\mu_\tau)} \mid \mathbf{X}_{<\tau}\right] \\
&= E\left[e^{\sum_{i=1}^{\tau-2} t(Z_i-\mu_i)} \cdot E\left[e^{t(Z_{\tau-1}-\mu_{\tau-1})} \mid \mathbf{X}_{<\tau-1}\right]\right] \cdot E\left[e^{t(Z_\tau-\mu_\tau)} \mid \mathbf{X}_{<\tau}\right] \\
&= E\left[e^{\sum_{i=1}^{\tau-2} t(Z_i-\mu_i)}\right] \cdot E\left[e^{t(Z_{\tau-1}-\mu_{\tau-1})} \mid \mathbf{X}_{<\tau-1}\right] \cdot E\left[e^{t(Z_\tau-\mu_\tau)} \mid \mathbf{X}_{<\tau}\right] \\
&= \prod_{i=1}^{\tau} E\left[e^{t(Z_i-\mu_i)} \mid \mathbf{X}_{<i}\right]
\end{aligned}$$

Due to the conditional expected value we cannot directly apply Hoeffding's lemma to yield an upper bound on this expression. Recall that this result states for any real valued random variable W such that $a \le W \le b$ almost surely that for all $\lambda \in \mathbb{R}$

$$E\left[e^{\lambda(W)}\right] \le e^{\lambda \cdot E[W] \frac{\lambda^2(b-a)^2}{8}}.$$

Fortunately we can derive a conditional version as well. The key is to define new random variables $W_i = Z_i - \mu_i$ for all $i \le \tau$ and observe that $E[W_i \mid X_{<i}] = 0$. Using the convexity of $e^{\lambda x}$ we get

$$\begin{aligned}
E\left[e^{tW_i} \mid X_{<i}\right] &\le E\left[\frac{b-W_i}{b-a} \cdot e^{ta} + \frac{W_i-a}{b-a} \cdot e^{tb} \mid X_{<i}\right] \\
&= \frac{b-E[W_i \mid X_{<i}]}{b-a} \cdot e^{ta} + \frac{E[W_i \mid X_{<i}]-a}{b-a} \cdot e^{tb} \\
&= \frac{b}{b-a} \cdot e^{ta} + \frac{-a}{b-a} \cdot e^{tb}
\end{aligned}$$

The remaining steps to proof the conditional version of Hoeffding's lemma are identical to the original proof. Thus,

$$E\left[e^{\lambda(Z_i-\mu_i)} \mid X_{<i}\right] \le e^{\frac{\lambda^2(b-a)^2}{8}}.$$

At last we combine this with the calculation from the beginning and obtain

$$\Pr[Z-\mu < -\varepsilon] \le \frac{E\left[e^{t(Z-m)}\right]}{e^{-t\varepsilon}} \le e^{\frac{\tau \cdot t^2(b-a)^2}{8} + t\varepsilon}$$

By optimizing the choice of $t < 0$ we set $t = -4\lambda/(\tau(b-a)^2)$ and the desired result follows. □

A.2. Drift results

The next result is a modified version of a drift result in [46]. We adapted the proof slightly. The original proof can be found in the full version [45].

Theorem A.8 ([46], Modified version of Claim 2.9)**.** *Consider a Markov Chain* $(W(t))_{t=1}^{\infty}$ *with the state space* $\{0, \dots, c_4\sqrt{\log n}\}$ *for an arbitrary constant* $c_4 > 0$*. For some constants* $c_2 > 0$ *and* $\varepsilon > 0$ *it has the following properties:*

- $\Pr[W(t+1) \geq 1 | W(t) = 0] = \Omega(1)$
- $\Pr[W(t+1) \geq \min\{(1+\varepsilon)W(t), m\}] \geq 1 - e^{-c_2 W(t)}$

Then it holds for $t = \mathrm{O}(\log n)$ *that*

$$\Pr[W(t) \geq m] \geq 1 - n^{-2}$$

Proof. We follow the outline of the original proof in the full version [45]. Let $B \in \mathbb{N} \cup \{0, \infty\}$ be a random variable that denotes the number of consecutive successful rounds (abb.: winning streak) when starting at round t_0 with $W_{t_0} = 0$ until the first failure similar to a geometrically distributed random variable. Let $\ell^* \in \mathbb{N}$ be the smallest number such that $W(\ell^*) \geq c_4\sqrt{\log n}$. We know that $\Pr[B = 0] \leq 1 - p$ and for any $1 \leq \ell \leq \ell^*$

$$\Pr[B = \ell] \leq p \cdot \prod_{j=1}^{\ell-1} (1 - e^{-c_2(1+\varepsilon)^{2j}}) \cdot e^{-c_2(1+\varepsilon)^{2\ell}} \leq p' \cdot e^{-c_2(1+\varepsilon)^{2\ell}} \leq c_3 \cdot \delta^{\ell}$$

for some constant $p', \delta < 1$. It is easy to see that for some constant $c_q < 1$

$$\Pr[B = \ell | B < \infty] \leq c_q \cdot \delta^{\ell}$$

Thus, $\mathbb{E}[B | B < \infty] = \Theta(1)$. In a similar way, it also follows for any starting value $W(t_0) = w_0 \geq 0$ that $\Pr[B = \ell | W(t_0) = w_0] \leq c_q \cdot \delta^{\ell}$, i.e., the probability holds irrespective of $W(t_0)$. If a winning streak holds for more than $t' = \Theta(\log\log n)$ phases, then we reach a phase where $W(t_0 + t') = c_4\sqrt{\log n}$ with probability

$$\Pr[B \geq t'] \geq p \cdot \prod_{j=1}^{t'} (1 - e^{-c_2(1+\varepsilon)^{2j}}) \geq c_5$$

for some constant $c_5 < 1$. Thus, by a standard Chernoff bound, we have to consider $\Theta(\log n)$ attempts such that at least one streak lasts for more than t' phases w.h.p.
As stated in the original proof, at most $\Theta(\log n)$ attempts requires at most $\Theta(\log n)$ phases w.h.p. which finishes the proof. □

Theorem A.9 (Theorem 18 of [63])**.** *Let* $(X_t)_{t \geq 0}$ *be a sequence of non-negative random variables with a finite state space* $\mathcal{S} \subseteq \mathbb{R}_0^+$ *such that* $0 \in \mathcal{S}$*. Let* $s_{min} := \min(\mathcal{S} \setminus \{0\})$, *and let $T := \inf\{t \geq 0 \mid X_t = 0\}$. Suppose that $X_0 = s_0$, and that there exists $\delta > 0$ such that for all $s \in \mathcal{S} \setminus \{0\}$ and all $t \geq 0$,*

$$E[X_t - X_{t+1} \mid X_t = s] \geq \delta s.$$

Then, for all $r \geq 0$,

$$Pr\left[T > \left\lceil \frac{r + \ln(s_0/s_{min})}{\delta} \right\rceil\right] \leq e^{-r}.$$

The next result is another modified version of a drift result in [46]. For convenience, we give a slightly adapted and condensed version of the proof.

Lemma A.10 (Modified version of [46])**.** *Let $W(t)$ be the random variable at time t of a random walk on the state space $[0, \log\log n]$ with a reflective state 0 and absorbing state $\log\log n$ and initially $W(0) = 0$ The transition probabilities are defined for every $t \in \mathbb{N}$ and $\ell \in [1, \log\log n - 1]$ as follows*

$$\begin{aligned}
&\Pr[W(t+1) = 1 \mid W(t) = 0] = p \\
&\Pr[W(t+1) = \ell + 1 \mid W(t) = \ell] = 1 - e^{-2^\ell} \\
&\Pr[W(t+1) = 0 \mid W(t) = \ell] = e^{-2^\ell}
\end{aligned}$$

where $p \leq 1$ is an arbitrary constant. Let T be the first time that $W(T) = \log\log n$, i.e., W reaches the absorbing state. Then $T = O(\log n)$ w.h.p.

Proof. We consider a sequence of attempts $Z_1, Z_2, \ldots$ such that W reaches the absorbing state $\log\log n$.

The attempts are identical distributed and each (unsuccessful) attempt can be described by a random variable B that denotes the number of consecutive successes (right steps of W) starting with $W(t) = 0$ before its first fail (falling back to state 0). Note that a successful attempt ends up in the absorbing state $\log\log n$. We show that each attempt is successful with at least constant probability and then apply Chernoff bounds to conclude that $O(\log n)$ attempts are sufficient to provide at least one successful attempt. Additionally, we show that $t = O(\log n)$, i.e., the total number of trials sum up over all attempts is $O(\log n)$. We start with the first statement. For any $\ell \in [1, \log\log n - 1]$ we have

$$\Pr[B = \ell] = p \cdot \prod_{j=1}^{\ell-1} (1 - e^{-2^j}) \cdot e^{-2^\ell} \leq p \cdot e^{-2^\ell} \leq e^{-2^\ell}$$

and hence,

$$\begin{aligned}
\Pr[B < \log\log n] &= \sum_{\ell=0}^{\log\log n-1} \Pr[B = \ell] \\
&= \Pr[B = 0] + \sum_{\ell=1}^{\log\log n-1} \Pr[B = \ell] \\
&\le (1-p) + \sum_{\ell=1}^{\log\log n-1} p \cdot e^{-2^{\ell}} \\
&\le (1-p) + p \cdot \sum_{\ell=1}^{\infty} e^{-2^{\ell}} \\
&\le (1-p) + 0.2 \cdot p \\
&= 1 - 0.8 \cdot p
\end{aligned}$$

Therefore each attempt Z_i is successful with probability at least $1 - (1 - 0.8p) = 0.8p$. Now consider $r = c \log n$ random variables $S_1, \ldots, S_r$ each indicates whether the attempt Z_i is successful. We know $\Pr[S_i = 1] \ge 0.8p$ for every $i \le r$. An application of Chernoff bounds (Theorem A.1) yields at least one successful attempt w.h.p.

Now we continue with the second part of the statement. From the first part we already know that $r = c \log n$ attempts are sufficient. We upper bound the total number of steps of the random walk until it reaches the absorbing state by upper bound the total number of steps of $r = c \log n$ unsuccessful attempts. In order to do that we define new independent random variables $Z'_i = Z_i + 1$ for each $i \le r$ and $Z' = \sum_{i=1}^{r} Z'_i$. Observe that $\Pr[Z'_i = \ell] = \Pr[Z_i = \ell - 1]$ for every $\ell \in [1, \log\log n - 1]$. Using our results from the first part we know that

$$\Pr\left[Z'_i = \ell\right] \le \begin{cases} p \cdot e^{-2^{\ell-1}} & , \ell \in [2, \log\log n - 1] \\ p & , \ell = 1. \end{cases}$$

By simple calculation it is easy to see that $e^{-2^{\ell-1}} \le e^{-2(\ell-1)}$ for $\ell \ge 2$ and hence, $\Pr[Z'_i = \ell] \le p \cdot e^{-2(\ell-1)}$ for all $\ell \in [1, \log\log n - 1]$. Therefore

$$\mathbb{E}\left[Z'\right] = \sum_{i=1}^{r} \mathbb{E}\left[Z'_i\right] = c \log n \cdot \sum_{\ell=1}^{\log\log n-1} \ell \cdot \Pr\left[Z'_i = \ell\right] \le c \log n \cdot p \cdot \sum_{\ell=1}^{\log\log n-1} \ell \cdot e^{-2^{\ell-1}} \le p \cdot c \cdot \log n.$$

This allows us to apply the following Chernoff bound (Lemma A.5) that yields for $Z' = \sum_{i=1}^{r} Z'_i$, $\mu = \mathbb{E}[Z']$, $\varepsilon = 2$ and $c = 6$

$$\Pr\left[Z' \ge c' \log n\right] \le \Pr\left[Z' \ge (1+\varepsilon)p \cdot c \cdot \log n + O(r)\right] \le e^{-\frac{\varepsilon^2 \cdot c \cdot \log n}{2(1+\varepsilon)}} \le n^{-4}$$

and hence, the second part of the statement holds. $\square$

A.3. Pólya-Eggenberger Distribution

The Pólya-Eggenberger process is a simple urn process that consists of n steps. Initially, the urn contains a red and b blue balls, where $a, b \in \mathbb{N}_0$. One fixed step of the process can be described as follows. First, a ball is drawn from the urn uniformly at random with replacement. Second, an additional ball that matches the color of the drawn ball is added to the urn. The corresponding *Pólya-Eggenberger distribution*, denoted by $\mathrm{PE}(a, b, n)$, describes the number of *total* red balls that are contained in the urn after all n steps. Alongside a more detailed discussion of this process, the following tail inequalities have been shown in [17] [1].

Theorem A.11 (Theorem 1 of [17])**.** *Let $A \sim \mathrm{PE}(a, b, n-(a+b))$, $\mu = (a/(a+b))n$ and $a + b \geq 1$. Then, for any δ with $0 < \delta < \sqrt{a}$ and some small constant $1 > \varepsilon_p > 0$ it holds that*

$$\Pr\left(A < \mu - \sqrt{a} \cdot \frac{n}{a+b} \cdot \delta\right) < 4\exp(-\varepsilon_p \cdot \delta^2)$$
$$\Pr\left(A > \mu + \sqrt{a} \cdot \frac{n}{a+b} \cdot \delta\right) < 4\exp(-\varepsilon_p \cdot \delta^2)$$

Theorem A.12 (simplified Theorem 47 of [17])**.** *Let $A \sim \mathrm{PE}(a, b, n-(a+b))$ with $1 \leq a \leq b$. Then, for some large constant $c_p > 1$ it holds that*

$$P\Big(A > \frac{n}{a+b} \cdot (3a + c_p \log n)\Big) < 2n^{-2},$$

A.4. Anti-Concentration Results

Lemma A.13. *Let $X \sim \mathrm{Bin}(n, p)$ with $\mu = np$. Then, for δ with $n/2 > (1+\delta)\mu > \mu$, it holds that*

$$\Pr[X \geq (1+\delta)\mu] \geq \frac{1}{\sqrt{8(1+\delta)\mu}} \cdot \left(1 - \frac{\delta^2\mu^2}{n-(1+\delta)\mu}\right) \cdot \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right)^\mu.$$

Proof. We start by considering some k with $n/2 > k > \mu$. By Lemma 4.7.2 of [13] we have that

$$\Pr[X \geq k] \geq \frac{1}{\sqrt{8k}} \exp\left(-nD\left(\frac{k}{n} \,\Big|\Big|\, p\right)\right), \tag{A.4}$$

where $D(\cdot||\cdot)$ denotes the Kullback-Leibler divergence with

$$D\left(\frac{k}{n} \,\Big|\Big|\, p\right) = \frac{k}{n}\ln\left(\frac{k}{\mu}\right) + \left(1 - \frac{k}{n}\right)\ln\left(\frac{n-k}{n-\mu}\right).$$

Hence, it follows that

$$\exp\left(-nD\left(\frac{k}{n} \,\Big|\Big|\, p\right)\right) = \left(\frac{\mu}{k}\right)^k \cdot \left(\frac{n-\mu}{n-k}\right)^{n-k}. \tag{A.5}$$

[1] In [17] the Pólya-Eggenberger distribution is defined to describe the number of *added* instead of *total* red balls at the end of the process. We adapted Theorem A.11 and Theorem A.12 accordingly.

Next, we use that $(1 + x/m)^m \geq e^x(1 - x^2/m)$ for $m > 1$ and $|x| < m$, which can be derived with the help of the well-known inequality $(1+\frac{1}{x})^{x+1} \geq e$ as well as the Bernoulli inequality. This implies

$$\left(\frac{n-\mu}{n-k}\right)^{n-k} = \left(1 + \frac{k-\mu}{n-k}\right)^{n-k} \geq e^{k-\mu}\left(1 - \frac{(k-\mu)^2}{n-k}\right). \tag{A.6}$$

When combining (A.4) with (A.5) and then (A.6), the statement follows for $k = (1 + \delta)\mu$. $\square$

Lemma A.14. *Let $X \sim \mathrm{Bin}(n,p)$ with $\mu = np$. Then, for δ with $n/2 > (1+\delta)\mu > \mu$* ***and*** *$\delta\mu < \sqrt{n}/2$, it holds that*

$$\Pr[X \geq (1+\delta)\mu] \geq \frac{1}{6 \cdot \sqrt{(1+\delta)\mu}} \cdot \exp\left(-\delta^2\mu\right)$$

Proof. The result is implied by Lemma A.13. We lower bound some factors involved in the right-hand side of Lemma A.13. It follows from $\delta\mu < \sqrt{n}/2$ and $\mu < n/2$ that

$$\left(1 - \frac{\delta^2\mu^2}{n-(1+\delta)\mu}\right) > \frac{1}{2}(1 - \mathrm{o}(1)).$$

Additionally, when using the well-known inequality $e^x \geq (1+x)$ twice, we get

$$\left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right) \geq \frac{1}{(1+\delta)^\delta} = \left(\frac{1}{1+\delta}\right)^\delta \geq \left(\frac{1}{e^\delta}\right)^\delta = e^{-\delta^2}.$$ $\square$

Lemma A.15 (Lemma 4 of [60])**.** *Let $X \sim \mathrm{Bin}(n,p)$ with $\mu = np$. For any $\delta \in (0, 1/2]$ and $p \in (0, 1/2]$, assuming $\delta^2\mu \geq 3$, it holds that*

$$\Pr[X \geq (1+\delta)\mu] \geq e^{-9\delta^2\mu}$$
$$\Pr[X \leq (1-\delta)\mu] \geq e^{-9\delta^2\mu}$$

Theorem A.16 (Theorem 1 of [54])**.** *Let $X \sim \mathrm{Bin}(n,p)$ with $\mu = np$. If $1/n < p$, then*

$$\Pr[X \geq \mu] > 1/4.$$

A.5. Random Walks

The following statement bounds the hitting time for biased random walks. Similar results have already been shown, e.g., in [26, 49, 65]. For convenience, we give here a combined version of these standard results that fits our needs.

Lemma A.17 (Random Walk Hitting Time)**.** *Let $(W_t)_{t\in\mathbb{N}}$ be a biased random walk on state space $\mathbb{N}_0$, initially at 0. Let $0 < p < 1$ denote the probability for the walk to move to the right (increase its current position by 1). Conversely, let $q = (1-p)$ denote the probability that it moves to the left (or stays in position in case it currently resides at position 0). Then, for any $N > 0$ and hitting time $\tau_N = \min\{t \mid W_t = N\}$ the following holds:*

1. *If $p > q$ then $\tau_N \le (\frac{2}{p-q})^2 \cdot N$ with probability at least $1 - \exp(-N)$.*

2. *If $p < q$ then $\tau_N \ge (q/p)^{N/2}$ with probability at least $1 - (p/q)^{N/2}$.*

Proof. We start with the first statement and assume $p > q$. We use a similar idea as in Lemma 3.3 of [26]. That is, we let X_i denote a random variable with $X_i = -1$ if the random walk moves to the left, and $X_i = 1$ if it moves to the right in step i. Observe that $S_m = \sum_{i=1}^{m} X_i$ minorizes the position W_m of the random walk for any $m \ge 0$. We set $m = (2/p-q)^2 N$ and apply Hoeffding's bound (Theorem 4.12 of [66]). As $-1 \le X_i \le 1$ this yields for any $t \ge 0$ that

$$\Pr\left[S_m \le \mathbb{E}\left[S_m\right] - t\right] \le \exp(-2t^2/4m).$$

Setting $t = \mathbb{E}\left[S_m\right] - N = m(p-q) - N \ge 0$ this yields that

$$\begin{aligned}\Pr\left[S_m \le N\right] &\le \exp\left(-\frac{(m(p-q)-N)^2}{2m}\right) = \exp\left(-\frac{m(p-q)^2}{2} + N(p-q) - \frac{N^2}{2m}\right)\\ &= \exp\left(-2N + N(p-q) - \frac{N^2}{2m}\right) \le \exp(-N).\end{aligned}$$

As S_m minorizes W_m, this implies that the random walk must have hit N before step m with probability at least $1 - \exp(-N)$.

In order to show the second statement, we assume $q < p$ and couple our process with a sequence of gamblers ruin instances. The gambler starts with 1 money and repeatedly gambles: either it wins 1 money with probability p or loses 1 money with probability q. The gambler continues until it either runs out of money or reaches a budget of $N + 1$. Assume our random walk currently resides at position 0. We couple its next moves with a gamblers ruin process as follows: the random walk moves to the right each time the gambler wins, otherwise it moves to the left. If the gambler reaches budget $N + 1$, then this implies that the random walk hit N before going back to 0. Otherwise, the gambler runs broke which implies that the random walk is again back at 0. Hence, we may lower-bound τ_N by the number of gamblers ruin instances required for the gambler to hit $N + 1$ for the first time. According to [49] the player reaches the desired budget with probability

$$\frac{\frac{q}{p} - 1}{(\frac{q}{p})^{N+1} - 1} \le (\frac{p}{q})^N.$$

We now apply union bounds, which implies that the gambler wins in any of the first $(q/p)^{N/2}$ instances with probability at most

$$(q/p)^{N/2} \cdot (\frac{p}{q})^N = \frac{p}{q}^{N/2}. \qquad \square$$

Each gamblers ruin instance corresponds to at least one move of the random walk. Therefore, the number of required gamblers ruin instances serves as a lower bound for the hitting time.

Lemma A.18 ([49])**.** *If we run an arbitrarily long sequence of independent trials, each with success probability at least* p*, then the probability that the number of failures ever exceeds the number of successes by* b *is at most* $((1-p)/p)^b$*.*

Lemma A.19 ([5])**.** *Let* $W(t)$ *be the random variable at time* t *of a random walk on the positive integers with a reflective border at* 0 *and* $W(0) = 0$*. Let* p *be the probability of a* $+1$*-step. Let* $q > p$ *be the probability of a* -1*-step everywhere except for the origin. Let* $r = 1 - p - q$ *be the probability of remaining in place (*$1 - p$ *for the origin). Let* $T_m = \inf\{\, t \geq 0 \mid W(t) \geq m \,\}$*. Then* $\Pr[T_m \leq n^c] \leq n^c \cdot (p/q)^m$*.*

Lemma A.20 ([58])**.** *Consider a submartingale* $Z_0, Z_1, \ldots$ *w.r.t. a filtration* $\mathcal{F} = (\mathcal{F}_i)_{i=0}^{\tau-1}$*. Assume* $a \leq Z_i - Z_{i-1} \leq b$ *for all* $i \geq 1$*. Then for all positive integers* τ *and* $\lambda > 0$

$$\Pr[Z_\tau - Z_0 < -\lambda] \leq e^{-\frac{2\lambda^2}{\tau(b-a)^2}}$$