



HAL
open science

Early Adapting to Trends: Self-Stabilizing Information Spread using Passive Communication

Amos Korman, Robin Vacus

► **To cite this version:**

Amos Korman, Robin Vacus. Early Adapting to Trends: Self-Stabilizing Information Spread using Passive Communication. ACM Symposium on Principles of Distributed Computing, Jul 2022, Salerno, Italy. hal-03615861

HAL Id: hal-03615861

<https://hal.science/hal-03615861>

Submitted on 21 Mar 2022

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

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

Early Adapting to Trends: Self-Stabilizing Information Spread using Passive Communication*

Amos Korman[†] and Robin Vacus[‡]

March 21, 2022

Abstract

How to efficiently and reliably spread information in a system is one of the most fundamental problems in distributed computing. Recently, inspired by biological scenarios, several works focused on identifying the minimal communication resources necessary to spread information under faulty conditions. Here we study the self-stabilizing *bit-dissemination* problem, introduced by Boczkowski, Korman, and Natale in [SODA 2017]. The problem considers a fully-connected network of n agents, with a binary world of *opinions*, one of which is called *correct*. At any given time, each agent holds an opinion bit as its public output. The population contains a *source* agent which knows which opinion is correct. This agent adopts the correct opinion and remains with it throughout the execution. We consider the basic *PULL* model of communication, in which each agent observes relatively few randomly chosen agents in each round. The goal of the non-source agents is to quickly converge on the correct opinion, despite having an arbitrary initial configuration, i.e., in a self-stabilizing manner. Once the population converges on the correct opinion, it should remain with it forever. Motivated by biological scenarios in which animals observe and react to the behavior of others, we focus on the extremely constrained model of *passive communication*, which assumes that when observing another agent the only information that can be extracted is the opinion bit of that agent. We prove that this problem can be solved in a poly-logarithmic in n number of rounds with high probability, while sampling a logarithmic number of agents at each round. Previous works solved this problem faster and using fewer samples, but they did that by decoupling the messages sent by agents from their output opinion, and hence do not fit the framework of passive communication. Moreover, these works use complex recursive algorithms with refined clocks that are unlikely to be used by biological entities. In contrast, our proposed algorithm has a natural appeal as it is based on letting agents estimate the current tendency direction of the dynamics, and then adapt to the emerging trend.

1 Introduction

1.1 Background and motivation

Disseminating information from one or several sources to the whole population is a fundamental building block in a myriad of distributed systems Censor-Hillel et al. (2012); Demers et al. (1987); Karp et al. (2000); Chierichetti et al. (2018); Giakkoupis and Woelfel (2011), including in multiple natural systems Sumpter et al. (2008); Razin et al. (2013); Feinerman and Korman (2017). This task becomes particularly challenging when the system is prone to faults Boczkowski et al. (2018b); Georgiou et al. (2013, 2011); Dutta et al. (2013), or when the agents or their interactions are constrained Boczkowski et al. (2019); Angluin et al. (2006). These issues find relevance in a variety of systems, including insect populations Razin et al. (2013), chemical reaction networks Chen et al. (2014), and mobile sensor networks Yick et al. (2008). In particular, in many biological systems, the internal computational abilities of individuals are impressively diverse, whereas the communication capacity is highly limited Barclay (1982); Razin et al. (2013); Feinerman

*This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 648032).

[†]CNRS, located at the French-Israeli Laboratory on Foundations of Computer Science, UMI FILOFOCS, CNRS, UP7, TAU, HUJI, WIS International Joint Research Unit, Tel-Aviv, Israel.

[‡]CNRS, located at the Research Institute on the Foundations of Computer Science (IRIF), Paris, France.

and Korman (2017). An extreme situation, often referred to as *passive communication* Wilkinson (1992), is when information is gained by observing the behavior of other animals, which, in some cases, may not even intend to communicate Cvikel et al. (2015); Giraldeau and Caraco (2018). Such public information can reflect on the quality of possible behavioral options, hence allowing to improve fitness when used properly Danchin et al. (2004).

Consider, for example, the following scenario that serves as an inspiration for our model. A group of n animals is scattered around an area searching for food. Assume that one side of the area, say, either the eastern side or the western side, is preferable (e.g., because it contains more food or because it is less likely to attract predators). However, only a few animals know which side is preferable. These knowledgeable animals will therefore spend most of their time in the preferable side of the area. Other animals would like to exploit the knowledge held by the knowledgeable animals, but they are unable to distinguish them from others. Instead, what they can do, is to scan the area in order to roughly estimate how many animals are on each side, and, if they wish, move between the two sides. Can the group of non-knowledgeable animals manage to locate themselves on the preferable side relatively fast, despite initially being spread in an arbitrary way while being completely uncorrelated?

The scenario above illustrates the notion of passive communication. The decision that an animal must make at any given time is to specify on which side of the area it should forage. This choice would be visible by others, and would in fact be the only information that an animal could reveal. Moreover, it cannot avoid revealing it. In particular, even the knowledgeable animals, who do not necessarily wish to communicate intentionally, cannot avoid revealing their “correct” choice. Assuming these animals do not actively try to harm others, they would simply reside on the preferable side, and promote this choice passively. The unknowledgeable animals, on the other hand, have a clearer incentive to cooperate, and they could, in principle, manipulate their choices to enhance the convergence process towards the correct choice. However, if their algorithm is also required to self-stabilize, then such a manipulation would be highly limited because, eventually, the choices should converge on one particular choice.

1.2 The problem

This paper studies the *self-stabilizing bit-dissemination* problem, introduced by Boczkowski, Korman, and Natale in Boczkowski et al. (2019), with the aim of solving it using passive communication. The problem considers a fully-connected network of n agents, and a binary world of *opinions*, say $\{0, 1\}$. One of these opinions is called *correct* and the other is called *wrong*. Execution proceeds in synchronous rounds (though agents do not have knowledge about the round number). At any given round t , each agent i holds an opinion bit $Y_t^{(i)} \in \{0, 1\}$ (viewed as its output variable). The population contains one *source agent* which knows which opinion is correct. This agent adopts the correct opinion and remains with it throughout the execution. Each agent knows whether or not it is the source, which can be formalized by assuming a designated *source-bit* in the state of an agent indicating this fact. We study the basic *PULL* model of communication Demers et al. (1987); Karp et al. (2000); Boczkowski et al. (2018b), in which in each round, each agent sees the information held by ℓ other agents, chosen uniformly at random (with replacement), where ℓ is small compared to n . We consider the *passive communication* model which assumes that the only information that can be obtained by sampling an agent is its opinion bit. Hence, sampling ℓ agents is equivalent to receiving an integer between 0 and ℓ corresponding to the number of agents with opinion 1 among the sampled agents.

In the *self-stabilizing* framework, the goal of the non-source agents is to quickly converge on the correct opinion, despite having an arbitrary initial configuration, that is set by an adversary (and despite not being able to distinguish the source from non-sources). We note that although the initial states of agents are arbitrary, we assume that their source-bit is not corrupted, and therefore, they reliably know whether or not they are the source. Moreover, by our assumption, it is guaranteed that there is only one source agent in the system¹. The adversary may initially set a different opinion to the source, but then the value of the correct bit would change, and the convergence should be guaranteed with respect to the new value.

The *running time* of the protocol corresponds to the first round t_{con} that the configuration of opinions

¹Our framework and proofs can be extended to allow for a constant number of sources, however, in this case it must be guaranteed that all sources agree on which opinion is the correct one. Indeed, as mentioned below, when there are conflicts between sources, the problem cannot be solved efficiently in the passive communication model, even if significantly more agents support one opinion.

reached a consensus on the correct opinion, and remained unchanged forever after. We say that a protocol converges in time T *with high probability (w.h.p)* if $t_{con} \leq T$ with probability at least $1 - 1/n^c$, for some constant $c > 1$. Note that we do not require agents to irrevocably commit on their final opinion, but rather that they eventually converge on the correct opinion without necessarily being aware that convergence has happened.

On the difficulties resulting from using passive communication. Previous works on the self-stabilizing bit-dissemination problem focused on identifying the minimal number of bits per interaction (message size) that need to be revealed in order to solve the problem in a short time. Boczkowski, Korman, and Natale showed in Boczkowski et al. (2019) that the problem can efficiently be solved in $\tilde{O}(\log n)$ rounds w.h.p, by sampling $\ell = 2$ agents at each round and using messages of size 3 bits. The protocol therein is based on agents having clocks that tick at each round. However, because of the self-stabilizing setting, the values of these clocks may initially be completely uncorrelated. The idea in that paper was to use 3 auxiliary bits in the messages to synchronize the clocks of agents in a self-stabilizing manner. These clocks were then used to facilitate the convergence of the opinion bit, which encapsulated another bit in the message. These 4 bits were then compressed to 3 bits using a recursive message-size-reduction mechanism. By reducing the message size of that clock-synchronization scheme to 1, and using a version of the aforementioned recursive message-size-reduction mechanism, a recent work by Bastide, Giakkoupis, and Saribekyan allows to reduce the message size of the self-stabilizing bit-dissemination problem to 1 bit Bastide et al. (2021). Moreover, stabilization is achieved in $O(\log n)$ rounds w.h.p., by sampling a single agent at each round. Importantly, however, both of these works decoupled the opinion of an agent from the message it uses, and hence, do not fit the framework of passive communication. Moreover, these works use complex recursive algorithms with refined clocks that are unlikely to be used by biological entities. Instead, we are interested in identifying algorithms that have a more natural appeal.

To illustrate the difficulty of self-stabilizing information spread under passive communication, let us consider a more generalized problem than bit-dissemination called *majority bit-dissemination*. In this problem, the population contains $k \geq 1$ source agents which may not necessarily agree on which opinion is correct. Specifically, in addition to its opinion, each source-agent stores a *preference* bit $\in \{0, 1\}$. Let k_i be the number of source agents whose preference is i . Assume that sufficiently more source agents share a preference i over $1 - i$ (e.g., at least twice as many), and call i the *correct bit*. Then, w.h.p., all agents (including the sources that might have the opposite preference) should converge their opinions on the correct bit in poly-logarithmic time, and remain with that opinion for polynomial time². The authors of Boczkowski et al. (2019) showed that the self-stabilizing majority bit-dissemination problem can be solved in logarithmic time, using messages of size 3 bits, and the authors of Bastide et al. (2021) showed how to reduce the message size to 1. As mentioned, the messages in these protocols were different than the opinions, which were stored as internal variables, and hence the protocols in Boczkowski et al. (2019); Bastide et al. (2021) are not based on passive communication. In fact, the following simple argument implies that this problem could not be solved in poly-logarithmic time under the model of passive communication, even if the sample size is n (i.e., all agents are being observed in each round)!

Assume by contradiction that there exists a self-stabilizing algorithm that solves the majority bit-dissemination problem using passive communication. Let us run this algorithm on a scenario with $k_1 = n/2$ and $k_0 = n/4$. Since $k_1 \gg k_0$, then after a poly-logarithmic time, w.h.p., all agents would have opinion 1, and would remain with that opinion for polynomial time. Denote by t_0 the first time after convergence, and let s denote the internal state of one of the $n/4$ non-source agents at time t_0 . Similarly, let s' denote the internal state at time t_0 of one of the $n/4$ source agents with preference 0. Now consider a second scenario, where we have $k_0 = n/4$ and $k_1 = 0$. An adversary sets the internal states of agents (including their opinions) as follows. The internal states of the k_0 source agents (with preference 0) are all set to be s' . Moreover, their opinions (that these sources must publicly declare on) are all 1. Next, the adversary sets the internal states of all non-source agents to be s , and their opinions to be 1. We now compare the execution of the algorithm in the first scenario (starting at time t_0) with the execution in the second scenario (starting at time 0, i.e., after the adversary manipulated the states). Note that both scenarios start with all opinions

²Observe that the case $k = 1$ consists of a slightly weaker version than the bit-dissemination problem because the latter problem requires that after convergence is guaranteed w.h.p., the configuration remains correct for an indefinite time with probability 1.

being 1. Hence, since we consider the passive communication model, all observations in the first round of the corresponding executions, would be unanimously 1. Moreover, as long as no agent changes its opinion in both scenarios, all observations would continue to be unanimously 1. Furthermore, it is given that from time t_0 , w.h.p., all agents in the first scenario remain with opinion 1 for a polynomially long period. Therefore, using a union bound argument, it is easy to see that also in the second scenario, w.h.p., all agents would remain with opinion 1 for polynomial time. This contradicts the fact that in the second scenario, w.h.p., the agents should converge on the opinion 0 in poly-logarithmic time.

Note that the aforementioned impossibility result does not preclude the possibility of solving the self-stabilizing bit dissemination problem in the passive communication model, which does not involve a conflict between sources. Indeed, the authors of Boczkowski et al. (2019) have suggested several candidate protocols which worked well in simulations³, however, as mentioned therein, their analysis appears to be beyond the reach of currently known techniques regarding randomly-interacting agent systems in self-stabilizing contexts.

1.3 Our results

We propose a simple algorithm that efficiently solves the self-stabilizing bit-dissemination problem in the passive communication model. The algorithm has a natural appeal as it is based on letting agents estimate the current tendency direction of the dynamics, and then adapt to the emerging trend. More precisely (but still, informally), each non-source agent counts the number of agents with opinion 1 it observes in the current round and compares it to the number observed in the previous round. If more 1's are observed now, then the agent adopts the opinion 1, and similarly, if more 0's are observed now, then it adopts the opinion 0 (if the same number of 1's is observed in both rounds then the agent does not change its opinion). Intuitively, on the global level, this behavior creates a persistent movement of the average opinion of the non-source agents towards either 0 or 1, which “bounces” back when hitting the wrong opinion.

More formally, we first describe the following algorithm. In addition to the opinion bit, the algorithm at an agent uses two internal variables at round t , called count_t and count_{t-1} . The former variable aims to store the number of 1's observe in the previous round, while the latter aims to store the number of 1's observe in the current round. Hence, both encode a number in $\{0, \dots, \ell\}$, and can be stored using $O(\log \ell)$ bits of memory. The following procedure is executed at each round separately by each agent. As an input, a given agent observes the opinions of ℓ randomly chosen agents. Let J_t denote the set of agents it sampled at round t , and let $S_t(J_t) = \{Y_t^{(j)}\}_{j \in J_t}$ denote the set containing their opinions at round t . Finally, for any set A of opinions, let $\text{COUNT}(A)$ denote the number of 1-opinions in A .

Input: $S_t(J_t)$
 $\text{count}_t \leftarrow \text{COUNT}(S_t(J_t))$;
if $\text{count}_t > \text{count}_{t-1}$ **then** $Y_{t+1} \leftarrow 1$;
else if $\text{count}_t < \text{count}_{t-1}$ **then** $Y_{t+1} \leftarrow 0$;
else $Y_{t+1} \leftarrow Y_t$;

As it turns out, one feature of the aforementioned protocol will make the analysis difficult – that is, that Y_t and Y_{t+1} are dependent, even when conditioning on (x_{t-1}, x_t) . This is because count_{t-1} is used to compute both Y_t and Y_{t+1} . For example, if the set J_{t-1} sampled at round $t - 1$ happens to contain more 1's, then count_{t-1} is larger. In this case, Y_t has a higher chance of being 1, and Y_{t+1} has a higher chance of being 0. For this reason we introduce a modified version of the protocol that solves this dependence issue. The idea is to partition the set of opinions sampled at round t into 2 sets of equal size. One subset will be used to compare with a subset of round $t - 1$, and the other subset will be used to compare with a subset of round $t + 1$. Note that this implies that the set of agents sampled in a round is 2ℓ rather than ℓ , however, since we are interested in the case $\ell = O(\log n)$, this does not cause a problem. This modified protocol, called *Follow the Emerging Trend (FET)*, is the one we shall actually analyze. Note that, although we used time indices for clarity, the protocol does not require the agents to know t . The following consists of the main result in the paper.

³Note that simulations results may be deceiving in self-stabilizing contexts, since the worst initial conditions for a given protocol are not always evident.

Protocol 1: Follow the Emerging Trend (FET) at round t

Input: $S_t(J_t)$ Partition $S_t(J_t)$ into two sets S'_t, S''_t of equal size uniformly at random ; $\text{count}'_t \leftarrow \text{COUNT}(S'_t)$; $\text{count}''_t \leftarrow \text{COUNT}(S''_t)$;**if** $\text{count}'_t > \text{count}''_{t-1}$ **then** $Y_{t+1} \leftarrow 1$;**else if** $\text{count}'_t < \text{count}''_{t-1}$ **then** $Y_{t+1} \leftarrow 0$;**else** $Y_{t+1} \leftarrow Y_t$;

Theorem 1. *Algorithm FET solves the self-stabilizing bit-dissemination problem in the passive communication model. It converges in $O(\log^{5/2} n)$ rounds on the correct opinion, with high probability, while relying on $\ell = O(\log n)$ samples in each round, and using $O(\log \ell)$ bits of memory per agent.*

1.4 Related works

In recent years, the study of *population protocols* has attracted significant attention in the distributed computing community Alistarh and Gelashvili (2018, 2015); Alistarh et al. (2016); Angluin et al. (2008); Aspnes and Ruppert (2007). These models often consider agents that interact under random meeting patterns while being restricted in both their memory and communication capacities. By now, we understand the computational power of such systems rather well, but apart from a few exceptions, this understanding is limited to non-faulty scenarios.

The framework of *opinion dynamics* corresponds to settings of multiple agents, where in each round, each agent samples one or more agents at random, extracts their opinions, and employs a certain rule for updating its opinion. The study of opinion dynamics crosses disciplines, and is highly active in physics and computer science, see review e.g., in Becchetti et al. (2020). Many of the models of opinion dynamics can be considered as following passive communication, since the information an agent reveals coincides with its opinion. Generally speaking, however, the typical scientific approach in opinion dynamics is to start with some simple update rule, and analyze the resulting dynamics, rather than tailoring an updating rule to solve a given distributed problem. For example, perhaps the most famous dynamics in the context of interacting particles systems concerns the *voter* model Liggett and Liggett (1985). In theoretical computer science, extensive research has been devoted to analyzing the time to reach consensus, following different updating rules including the *3-majority* Doerr et al. (2011), *Undecided-State Dynamics* Angluin et al. (2008), and others. In these works, consensus should be reached either on an arbitrary value, or on the majority (or plurality) opinion, as evident in the initial configuration.

In many natural settings, however, the group must converge on a particular consensus value that is a function of the environment. Moreover, agents have different levels of knowledge regarding the desired value, and the system must utilize the information held by the more knowledgeable individuals Sumpter et al. (2008); Ayalon et al. (2021); Korman et al. (2014); Rajendran et al. (2022). As explained in more detail below, when communication is restricted, and the system is prone to faults, this task can become challenging.

Propagating information from one or more sources to the rest of the population has been the focus of a myriad of works in distributed computing. This dissemination problem has been studied under various models taking different names, including *rumor spreading*, *information spreading*, *gossip*, *broadcast*, and others, see e.g., Giakkoupis (2014); Censor-Hillel et al. (2012); Demers et al. (1987); Karp et al. (2000); Chierichetti et al. (2018); Giakkoupis and Woelfel (2011). A classical algorithm in the *PULL* model spreads the opinion of the source to all others in $2 \log n$ rounds, by letting each uninformed agent copy the opinion of an informed agent whenever seeing one for the first time Karp et al. (2000). Unfortunately, this elegant algorithm does not suit all realistic scenarios, since its soundness crucially relies on the absence of misleading information. To address such issues, rumor spreading has been studied under different models of faults. One line of research investigates the case that messages may be corrupted with some fixed probability Feinerman et al. (2017); Boczkowski et al. (2018a). Another model of faults is *self-stabilization* Dijkstra (1974), where the system must eventually converge on the opinion of the source regardless of the initial configuration of states Dijkstra (1974). For example, the algorithm in Karp et al. (2000) fails in this setting, since non-source agents may be initialized to “think” that they have already been informed by the correct opinion, while they

actually hold the wrong opinion.

More details on previous works regarding the self-stabilizing bit-dissemination problem. The self-stabilizing bit-dissemination problem was introduced in Boczkowski et al. (2019), with the aim of minimizing the message size. As mentioned therein, if all agents share the same notion of global time, then convergence can be achieved in $\mathcal{O}(\log n)$ time w.h.p. even under passive communication. The idea is that agents divide the time horizon into phases of length $T = 4 \log n$, and that each phase is further subdivided into 2 subphases of length $2 \log n$ each. In the first subphase of each phase, if a non-source agent observes an opinion 0, then it copies it as its new opinion, but if it sees 1 it ignores it. In the second subphase, it does the opposite, namely, it adopts the output bit 1 if and only if it sees an opinion 1. Now, consider the first phase. If the source supports opinion 0 then at the end of the first subphase, every output bit would be 0 w.h.p., and the configuration would remain that way forever. Otherwise, if the source supports 1, then at the end of the second subphase all output bits would be 1 w.h.p., and remain 1 forever.

The aforementioned protocol indicates that the self-stabilizing bit-dissemination problem could be solved efficiently by running a self-stabilizing *clock-synchronization* protocol in parallel to the previous example protocol. This parallel execution amounts to adding one bit to the message size of the clock synchronization protocol. The main technical contribution of Boczkowski et al. (2019), as well as the focus of subsequent work in Bastide et al. (2021), was solving the self-stabilizing clock-synchronization using as few as possible bits per message. In fact, the authors in Bastide et al. (2021) managed to do so using 1-bit messages. This construction thus implies a solution to the self-stabilizing bit-dissemination problem using 2 bits per message. A recursive procedure, similar to the one established in Boczkowski et al. (2019), then allowed to further compress the 2 bits into 1-bit messages. Importantly, however, the 1-bit message revealed by an agent is different from its opinion bit, which is kept in the protocols of Boczkowski et al. (2019); Bastide et al. (2021) as an internal variable. At first glance, to adhere to the passive communication model, one may suggest that agents simply choose their opinion to be this 1-bit message used in Bastide et al. (2021), just for the purpose of communication, until a consensus is reached, and then switch the opinion to be the correct bit, once it is identified. There are, however, two difficulties to consider regarding this approach. First, in our setting, the source agent does not change its opinion (which, in the case of Bastide et al. (2021), may prevent the protocol from reaching a consensus at all). Second, even assuming that the protocol functions properly despite the source having a stable opinion, it is not clear how to transition from the “communication” phase (where agents use their opinion to operate the protocol, e.g., for synchronizing clocks) to the “consensus” phase (where all opinions must be equal to the correct bit at every round). For instance, the first agents to make the transition may disrupt other agents still in the first phase.

2 Proof of Theorem 1: General Overview

The goal of this section is to prove Theorem 1. The $\mathcal{O}(\log \ell)$ bits upper bound on the memory complexity clearly follows from the fact that the only variables kept by the FET algorithm (Protocol 1) are count' and count'' , which are used to count the number of 1’s in a sample (of size ℓ).

Since the protocol is symmetric with respect to the opinion of the source, we may assume without loss of generality that the source has opinion 1. Our goal would therefore be to show that the FET algorithm converges to 1 fast, w.h.p., regardless of the initial configuration of non-source agents. Note that in order to achieve running time of $\mathcal{O}(T)$ w.h.p guarantee, is it sufficient to show that the algorithm stabilizes in T rounds with probability at least $1 - 1/n^\epsilon$, for some $\epsilon > 0$. Indeed, because of the self-stabilizing property of the algorithm, the probability that the algorithm does not stabilize within $2T/\epsilon$ rounds is at most $(1/n^\epsilon)^{2/\epsilon} = 1/n^2$.

For the sake of analysis, let x_t denote the fraction of agents with opinion 1 at round t among the whole population of agents (including the source). We shall extensively use the two dimensional grid $\mathcal{G} := \{0, \frac{1}{n}, \dots, \frac{n-1}{n}, 1\}^2$. When analyzing what happens at round $t + 2$, the x -axis of \mathcal{G} would represent x_t , and the y -axis would represent x_{t+1} .

Observation 1. *For any round t , conditioning on $x_t = \mathbf{x}_t$, and $x_{t+1} = \mathbf{x}_{t+1}$, the probability that a non-*

source agent i has opinion 1 on round $t + 2$ is

$$\mathbb{P}\left(Y_{t+2}^{(i)} = 1 \mid \begin{array}{l} x_t = \mathbf{x}_t \\ x_{t+1} = \mathbf{x}_{t+1} \end{array}\right) = \mathbb{P}(B_\ell(\mathbf{x}_{t+1}) > B_\ell(\mathbf{x}_t)) + \mathbb{1}_{\{Y_{t+1}^{(i)}=1\}} \cdot \mathbb{P}(B_\ell(\mathbf{x}_{t+1}) = B_\ell(\mathbf{x}_t)). \quad (1)$$

Moreover, there are independent binary random variables X_1, \dots, X_n such that x_{t+2} is distributed as $\frac{1}{n} \sum_{i=1}^n X_i$. Eventually,

$$\mathbb{E}\left(x_{t+2} \mid \begin{array}{l} x_t = \mathbf{x}_t \\ x_{t+1} = \mathbf{x}_{t+1} \end{array}\right) = \mathbb{P}(B_\ell(\mathbf{x}_{t+1}) > B_\ell(\mathbf{x}_t)) + \mathbf{x}_{t+1} \cdot \mathbb{P}(B_\ell(\mathbf{x}_{t+1}) = B_\ell(\mathbf{x}_t)) + \frac{1}{n}(1 - \mathbb{P}(B_\ell(\mathbf{x}_{t+1}) \geq B_\ell(\mathbf{x}_t))). \quad (2)$$

The proof of Observation 1 is deferred to Appendix B. A consequence of Observation 1, is that the execution of the algorithm induces a Markov chain on \mathcal{G} . This Markov chain has a unique absorbing state, $(1, 1)$, since we assumed the source to have opinion 1. To prove Theorem 1 we therefore only need to bound the time needed to reach $(1, 1)$.

2.1 Partitioning the grid into domains

Let us fix $0 < \delta < 1/2$ and $\lambda_n = \frac{1}{\log^{1/2+\delta} n}$. We partition \mathcal{G} into domains as follows (see illustration on Figure 1a).

$$\begin{aligned} \text{GREEN}_1 &= \left\{ (x_t, x_{t+1}) \text{ s.t. } x_{t+1} \geq x_t + \delta \right\}, \\ \text{PURPLE}_1 &= \left\{ (x_t, x_{t+1}) \text{ s.t. } \frac{1}{\log n} \leq x_t < \frac{1}{2} - 3\delta \text{ and } (1 - \lambda_n) \cdot x_t \leq x_{t+1} < x_t + \delta \right\}, \\ \text{RED}_1 &= \left\{ (x_t, x_{t+1}) \text{ s.t. } \frac{1}{\log n} \leq x_{t+1} \text{ and } x_t < \frac{1}{2} - 3\delta \text{ and } x_t - \delta \leq x_{t+1} < (1 - \lambda_n) \cdot x_t \right\}, \\ \text{CYAN}_1 &= \left\{ (x_t, x_{t+1}) \text{ s.t. } 0 \leq \min(x_t, x_{t+1}) < \frac{1}{\log n} \text{ and } x_t - \delta < x_{t+1} < x_t + \delta \right\}, \\ \text{YELLOW} &= \left\{ (x_t, x_{t+1}) \text{ s.t. } \frac{1}{2} - 3\delta \leq x_t < \frac{1}{2} \leq 3\delta \text{ and } \frac{1}{2} - 4\delta \leq x_{t+1} \leq \frac{1}{2} + 4\delta \text{ and } |x_{t+1} - x_t| < \delta \right\}. \end{aligned}$$

Similarly, for the former 4 domains, we define $\text{GREEN}_0, \text{PURPLE}_0, \text{RED}_0$ and CYAN_0 to be their symmetric equivalents (w.r.t the point $(\frac{1}{2}, \frac{1}{2})$), and finally define: $\text{GREEN} = \text{GREEN}_0 \cup \text{GREEN}_1$, $\text{PURPLE} = \text{PURPLE}_0 \cup \text{PURPLE}_1$, $\text{RED} = \text{RED}_0 \cup \text{RED}_1$, and $\text{CYAN} = \text{CYAN}_0 \cup \text{CYAN}_1$. We shall analyze each area separately, conditioning on the Markov chain to be at any point in that area, and focusing on the number of rounds required to escape the area, and the probability that this escape is made to a particular other area. Figure 1b represents a sketch of the proof of Theorem 1, which may help to navigate the intermediate results.

As it happens, the dynamics starting from a point (x_t, x_{t+1}) highly depends on the difference between x_t and x_{t+1} . Roughly speaking, the larger $|x_{t+1} - x_t|$ is the faster is the convergence. For this reason, we refer to $|x_{t+1} - x_t|$ as the *speed* of the point (x_t, x_{t+1}) . (This could also be viewed as the “derivative” of the process at time t .)

2.2 Analyzing the Markov chain at different domains

Due to lack of space, most proofs are deferred to the appendix. We have decided to include in the extended abstract the core of the proof of the main lemma that concerns the Yellow domain (Section 3.2), which was one of the more challenging results obtained in this paper. In addition, we also provide the core of the proof of the main lemma that concerns the Cyan domain (Section 4), because this part of the state-space is essential for understanding the dynamics of the protocol.

Let us now give an overview of the intermediate results. First we consider GREEN , in which the speed of points is large. In Lemma 1 we show that from points in that domain, non-source agents reach a consensus in just one round w.h.p. In particular, if the Markov chain is at some point in GREEN_1 , then the consensus




Figure 1: (a) Partitioning the state space into domains. The x-axis (resp., y-axis) represents the proportion of agents with opinion 1 in round t (resp., $t + 1$). The thick dashed line at the frontier between PURPLE₁ and RED₁ is defined by $x_{t+1} = (1 - \lambda_n)x_t$. (b) Sketch of the proof of Theorem 1. All transitions are w.p. at least $1 - 1/n^{\Omega(1)}$. The numbers next to the self-loops indicate that the process stays in the corresponding domain for this number of rounds w.p. at least $1 - 1/n^{\Omega(1)}$. The source is assumed to have opinion 1.

will be on 1, and we are done. If, on the other hand, the Markov chain is in GREEN₀, then the consensus of non-source agents would be on 0. As we show later, in that case the Markov chain would reach CYAN₁ in one round w.h.p.

Lemma 1 (Green area). *Assume that c is sufficiently large. If $(x_t, x_{t+1}) \in \text{GREEN}_1$, then w.h.p., for every non-source agent i , $Y_i^{(t+2)} = 1$. Similarly, if $(x_t, x_{t+1}) \in \text{GREEN}_0$, then w.h.p., for every non-source agent i , $Y_i^{(t+2)} = 0$.*

The proof of Lemma 1 follows from a simple application of Hoeffding's inequality, and is deferred to Appendix C.1. Next, we consider the area PURPLE, and show that the population goes from PURPLE to GREEN in just one round, w.h.p. In PURPLE, the speed is relatively low, and x_t and x_{t+1} are quite far from $1/2$. On the next round, we expect x_{t+2} to be close to $1/2$, thus gaining enough speed in the process to join GREEN. The proof of the following lemma is rather straightforward, and is deferred to Appendix C.2.

Lemma 2 (Purple area). *Assume that c is sufficiently large. If $(x_t, x_{t+1}) \in \text{PURPLE}_1$, then $(x_{t+1}, x_{t+2}) \in \text{GREEN}_1$ w.h.p. Similarly, if $(x_t, x_{t+1}) \in \text{PURPLE}_0$, then $(x_{t+1}, x_{t+2}) \in \text{GREEN}_0$ w.h.p.*

Next, we bound the time that can be spent in RED, by using the fact that as long as the process is in RED₁ (resp., RED₀), x_t (resp., $(1 - x_t)$) decreases (deterministically) by at least a multiplicative factor of $(1 - \lambda_n)$ at each round. After a poly-logarithmic number of rounds, the Markov chain must leave RED and in this case, we can show that it cannot reach YELLOW right away. The proof of the following lemma is again relatively simple, and is deferred to Appendix C.3.

Lemma 3 (Red area). *Consider the case that $(x_{t_0}, x_{t_0+1}) \in \text{RED}$ for some round t_0 , and let $t_1 = \min\{t \geq t_0, (x_t, x_{t+1}) \notin \text{RED}\}$. Then $t_1 < t_0 + \log^{1/2+2\delta} n$, and $(x_{t_1}, x_{t_1+1}) \notin \text{YELLOW} \cup \text{RED}$.*

Next, we bound the time that can be spent in CYAN₁. (A similar result holds for CYAN₀.) Roughly speaking, this area corresponds to the situation in which, over the last two consecutive rounds, the population

is in an almost-consensus over the wrong opinion. In this case, many agents (a constant fraction) see only 0 in their corresponding samples in the latter round. As a consequence, everyone of them who will see at least one opinion 1 in the next round, will adopt opinion 1. We can expect this number to be of order $\ell = O(\log n)$. This means that, as long as the Markov chain is in CYAN_1 , the value of x_t would grow by a logarithmic factor in each round. This implies that within $\log(n)/\log(\log n)$ rounds, the Markov chain will leave the CYAN_1 area and go to $\text{GREEN}_1 \cup \text{PURPLE}_1$. Informally, this phenomenon can be viewed as a form of “bouncing” — the population of non-sources reaches an almost consensus on the wrong opinion, and “bounces back”, by gradually increasing the fraction of agents with the correct opinion, up to an extent that is sufficient to enter $\text{GREEN}_1 \cup \text{PURPLE}_1$. The proof of the following lemma is given in Section 4.

Lemma 4 (Cyan area). *Consider the case that $(x_{t_0}, x_{t_0+1}) \in \text{CYAN}_1$ for some round t_0 , and let $t_1 = \min\{t \geq t_0, (x_t, x_{t+1}) \notin \text{CYAN}_1\}$. Then with probability at least $1 - 1/n^{\Omega(1)}$ we have (1) $t_1 < t_0 + \log(n)/\log(\log n)$, and (2) $(x_{t_1}, x_{t_1+1}) \in \text{GREEN}_1 \cup \text{PURPLE}_1$.*

Eventually, we consider the central area, namely, YELLOW , where the speed is very low, and bound the time that can be spent there. The proof of the following lemma is more complex than the previous ones, and it appears in Section 3.

Lemma 5 (Yellow area). *Consider the case that $(x_{t_0}, x_{t_0+1}) \in \text{YELLOW}$. Then, w.h.p.,*

$$\min\{t > t_0 \text{ s.t. } (x_t, x_{t+1}) \notin \text{YELLOW}\} < t_0 + O(\log^{5/2} n).$$

2.3 Assembling the lemmas

Given the aforementioned lemmas, we have everything we need to prove our main result.

Proof of Theorem 1. Recall that without loss of generality, we assumed the source to have opinion 1. The reader is strongly encouraged to refer to Figure 1b to follow the ensuing arguments more easily.

- Let $t_1 = \min\{t \geq 0, (x_t, x_{t+1}) \notin \text{YELLOW}\}$. If $(x_0, x_1) \in \text{YELLOW}$, we apply Lemma 5 to get that

$$t_1 < O(\log^{5/2} n) \text{ w.h.p. and } (x_{t_1}, x_{t_1+1}) \in \text{RED} \cup \text{CYAN} \cup \text{PURPLE} \cup \text{GREEN}. \quad (3)$$

Else, $(x_0, x_1) \notin \text{YELLOW}$ so $t_1 = 0$, and Eq. (3) also holds.

- Let $t_2 = \min\{t \geq t_1, (x_t, x_{t+1}) \notin \text{RED}\}$. If $(x_{t_1}, x_{t_1+1}) \in \text{RED}$, we apply Lemma 3 to get that

$$t_2 < t_1 + \log^{1/2+2\delta} n \text{ w.h.p. and } (x_{t_2}, x_{t_2+1}) \in \text{CYAN} \cup \text{PURPLE} \cup \text{GREEN}. \quad (4)$$

Else, $(x_{t_1}, x_{t_1+1}) \notin \text{RED}$ so $t_1 = t_2$, and by Eq. (3), it implies that Eq. (4) also holds.

- Let $t_3 = \min\{t \geq t_2, (x_t, x_{t+1}) \notin \text{CYAN}\}$. If $(x_{t_2}, x_{t_2+1}) \in \text{CYAN}$, we apply Lemma 4 to get that

$$t_3 < t_2 + \log(n)/\log(\log n) \text{ and } (x_{t_3}, x_{t_3+1}) \in \text{PURPLE} \cup \text{GREEN} \text{ with probability at least } 1 - 1/n^{\Omega(1)}. \quad (5)$$

Else, $(x_{t_2}, x_{t_2+1}) \notin \text{CYAN}$ so $t_2 = t_3$, and by Eq. (4), it implies that Eq. (5) also holds.

- Let $t_4 = \min\{t \geq t_3, (x_t, x_{t+1}) \in \text{GREEN}\}$. By Lemma 2, and by Eq. (5), we have that $t_4 = t_3$ or $t_4 = t_3 + 1$ w.h.p.

If $(x_{t_4}, x_{t_4+1}) \in \text{GREEN}_1$, then by Lemma 1 the consensus is reached on round t_4+1 . Otherwise, if $(x_{t_4}, x_{t_4+1}) \in \text{GREEN}_0$, by Lemma 1, we obtain that $x_{t_4+2} = 1/n$ w.h.p. (meaning that all agents have opinion 0 except the source). Therefore, in this case, either $(x_{t_4+1}, x_{t_4+2}) \in \text{GREEN}_0$ or $(x_{t_4+1}, x_{t_4+2}) \in \text{CYAN}_1$ (because for a point (x_t, x_{t+1}) to be in any other area, it must be the case that $x_{t+1} \geq 1/\log(n)$, by definition). In the former case, we apply Lemma 1 again to get that $x_{t_4+3} = 1/n$ w.h.p., which implies that $(x_{t_4+2}, x_{t_4+3}) = (1/n, 1/n) \in \text{CYAN}_1$. As we did before, we apply Lemma 4, 2 and 1 to show that, with probability at least $1 - 1/n^{\Omega(1)}$, the system goes successively to $\text{PURPLE}_1 \cup \text{GREEN}_1$, then to GREEN_1 , and eventually reaches the absorbing state $(1, 1)$ in less than $\log(n)/\log(\log n) + 2$ rounds.

Altogether, the convergence time is dominated by t_1 , and is hence $O(\log n)^{5/2}$ with probability at least $1 - 1/n^\epsilon$, for some $\epsilon > 0$. As mentioned, this implies that for any given $c > 1$, the algorithm reaches consensus in $O(\log n)^{5/2}$ time with probability at least $1 - 1/n^c$. This concludes the proof of Theorem 1. \square

3 Escaping the Yellow Area

The goal of this section is to prove Lemma 5. It might be easier for the reader to think of the Yellow area as a square. Formally, let us define YELLOW' as the following square bounding box around YELLOW :

$$\text{YELLOW}' = \left\{ (x_t, x_{t+1}) \text{ s.t. } 1/2 - 4\delta \leq x_t, x_{t+1} \leq 1/2 + 4\delta \right\}.$$

Obviously, $\text{YELLOW} \subset \text{YELLOW}'$, so in order to prove Lemma 5 it suffices to prove Lemma 6 below.

Lemma 6. *Consider that $(x_{t_0}, x_{t_0+1}) \in \text{YELLOW}'$. Then, w.h.p., $\min\{t > t_0 \text{ s.t. } (x_t, x_{t+1}) \notin \text{YELLOW}'\} < t_0 + O(\log^{5/2} n)$.*

3.1 General structure of the proof

In order to prove Lemma 6, we first partition YELLOW' , as follows (for an illustration, see Figure 2):

$$\begin{aligned} \mathbf{A}_1 &= \{(x_t, x_{t+1}) \mid \text{(i) } x_{t+1} \geq 1/2 \text{ and (ii) } x_{t+1} - x_t \geq x_t - 1/2\} \cap \text{YELLOW}', \\ \mathbf{B}_1 &= \{(x_t, x_{t+1}) \mid \text{(i) } x_{t+1} \geq x_t \text{ and (ii) } x_{t+1} - x_t < x_t - 1/2\} \cap \text{YELLOW}', \\ \mathbf{C}_1 &= \{(x_t, x_{t+1}) \mid \text{(i) } x_{t+1} < 1/2 \text{ and (ii) } x_{t+1} \geq x_t\} \cap \text{YELLOW}'. \end{aligned}$$

Similarly, we define $\mathbf{A}_0, \mathbf{B}_0, \mathbf{C}_0$ their symmetric equivalents (w.r.t the point $(\frac{1}{2}, \frac{1}{2})$), and $\mathbf{A} = \mathbf{A}_0 \cup \mathbf{A}_1$, $\mathbf{B} = \mathbf{B}_0 \cup \mathbf{B}_1$, and $\mathbf{C} = \mathbf{C}_0 \cup \mathbf{C}_1$.




Figure 2: Partitioning the YELLOW' domain.

Remark 1. *At various times throughout our analysis, we would like to calculate different statistical properties of the system at round $t + 2$, conditioning on $(\mathbf{x}_t, \mathbf{x}_{t+1}) \in \mathcal{G}$, as was done in e.g., Observation 1. For the sake of clarity of presentation, in all subsequent cases, we shall omit the conditioning notation. The reader should therefore keep in mind, that whenever such properties are calculated, they are actually done while conditioning on $x_t = \mathbf{x}_t$ and $x_{t+1} = \mathbf{x}_{t+1}$, where the point $(\mathbf{x}_t, \mathbf{x}_{t+1})$ would always be clear from the context. For example, in Lemma 7 below, the probability $\mathbb{P}((x_{t+1}, x_{t+2}) \notin \text{YELLOW}')$ is calculated conditioning on $x_t = \mathbf{x}_t$, $x_{t+1} = \mathbf{x}_{t+1}$, where $(\mathbf{x}_t, \mathbf{x}_{t+1}) \in \mathbf{A}$, yet this conditioning is omitted.*

In the next lemma, we study the distribution of the future location of any point $(x_t, x_{t+1}) \in \mathbf{A}$. This area happens to be ideal to escape YELLOW' , because it allows the Markov chain to quickly build up “speed”. Item (a) in the next lemma says that, with some probability that depends on the current speed the following occur: (1) the speed in the following round increases by a factor of two, and (2) the process in the next round either remains in \mathbf{A} , or goes outside of YELLOW' . Note that when the current speed is not too low, that is, when $x_{t+1} - x_t > 1/\sqrt{n}$, this combined event happens with constant probability. Item (b) says that with constant probability, (1) the speed in the next round would not be too low, and (2) the process either remains in \mathbf{A} , or goes outside of YELLOW' .

Lemma 7. *If $(x_t, x_{t+1}) \in \mathbf{A}$, and provided that δ is small enough and n is large enough,*

- (a) *We have $\mathbb{P}((x_{t+1}, x_{t+2}) \notin \text{YELLOW}' \setminus \mathbf{A} \cap |x_{t+2} - x_{t+1}| > 2|x_{t+1} - x_t|) > 1 - \exp(-3n \cdot (x_{t+1} - x_t)^2)$.*
- (b) *There exists a constant $c_2 = c_2(c) > 0$ s.t. $\mathbb{P}((x_{t+1}, x_{t+2}) \notin \text{YELLOW}' \setminus \mathbf{A} \cap |x_{t+2} - x_{t+1}| > 1/\sqrt{n}) > c_2$.*

Now, we can iteratively use the previous result to prove that any state in \mathbf{A} has a reasonable chance to escape YELLOW' . The proofs of both Lemmas 7 and 8 are deferred to Appendix C.5.1.

Lemma 8. *There is a constant $c_3 = c_3(c)$ s.t. if $(x_{t_0}, x_{t_0+1}) \in \mathbf{A}$, then $\mathbb{P}(\exists t_1 < t_0 + \log n, (x_{t_1}, x_{t_1+1}) \notin \text{YELLOW}') > c_3$.*

We are left with proving that the system cannot be stuck in \mathbf{B} or \mathbf{C} for too long. We start with \mathbf{B} . The analysis of this area is relatively complex, because it is difficult to rule out the possibility that the Markov chain remains there at a low speed. We prove that any state in \mathbf{B} must either make a step towards escaping YELLOW' , or have a good chance of leaving \mathbf{B} . The proof of the following lemma is given in Section 3.2.

Lemma 9. *There are constants $c_4, c_5 > 0$ such that if $(x_t, x_{t+1}) \in \mathbf{B}$, then either*

- (a) *$|x_{t+1} - 1/2| > (1 + c_4/\sqrt{\ell})|x_t - 1/2|$, or*
- (b) *$\mathbb{P}((x_{t+1}, x_{t+2}) \notin \mathbf{B}) > c_5$.*

Now, we can iteratively use the previous result to prove that any state in \mathbf{B} either leaves \mathbf{B} or escapes YELLOW' in a reasonable amount of time. The proof of Lemma 10 is deferred to Section C.5.2.

Lemma 10. *If $(x_{t_0}, x_{t_0+1}) \in \mathbf{B}$, then, w.h.p., $\min\{t \geq t_0, (x_t, x_{t+1}) \notin \mathbf{B}\} < t_0 + \frac{\sqrt{c}}{c_4} \cdot \log^{3/2} n$.*

We are left with proving that the system cannot stay in \mathbf{C} for too long. Fortunately, from this area, the Markov chain is naturally pushed towards \mathbf{A} , which makes the analysis simple. The proof of Lemma 11 is deferred to Section C.5.3.

Lemma 11. *There is a constant $c_6 > 0$ such that if $(x_t, x_{t+1}) \in \mathbf{C}$, then*

$$\max\{\mathbb{P}((x_{t+1}, x_{t+2}) \notin \text{YELLOW}' \setminus \mathbf{A}), \mathbb{P}((x_{t+2}, x_{t+3}) \notin \text{YELLOW}' \setminus \mathbf{A})\} > c_6.$$

Eventually, we have all the necessary results to conclude the proof regarding the Yellow area.

Proof of Lemma 6. By Lemma 8, if $(x_{t_0}, x_{t_0+1}) \in \mathbf{A}$, then $\mathbb{P}(\exists t_1 < t_0 + \log n, (x_{t_1}, x_{t_1+1}) \notin \text{YELLOW}') > c_3 > 0$. By Lemma 11, this implies that if $(x_{t_0}, x_{t_0+1}) \in \mathbf{A} \cup \mathbf{C}$,

$$\mathbb{P}(\exists t_1 < t_0 + \log n + 2, (x_{t_1}, x_{t_1+1}) \notin \text{YELLOW}') > \min(c_3, c_3 \cdot c_6) = c_3 \cdot c_6 > 0. \quad (6)$$

By Lemma 10, w.h.p., whenever the process is at \mathbf{B} , it does not spend more than $(\sqrt{c}/c_4) \cdot \log^{3/2} n$ consecutive rounds there. This means, that for any constant $c' > 0$, during $c' \log^{5/2} n$ consecutive rounds, w.h.p., we must either leave YELLOW' or be at $\mathbf{A} \cup \mathbf{C}$ on at least $\frac{c' \log^{5/2} n}{(\sqrt{c}/c_4) \cdot \log^{3/2} n} = \frac{c' c_4}{\sqrt{c}} \cdot \log n$ distinct rounds. By Eq. (6), the probability that the system fails to escape YELLOW' in each of these occasions is at most $(1 - c_3 \cdot c_6)^{(c' c_4/\sqrt{c}) \cdot \log n}$. Taking c' to be sufficiently large concludes the proof of Lemma 6. \square

3.2 Proof of Lemma 9 — The main lemma regarding area \mathbf{B}

The goal of this section is to prove Lemma 9, which concerns Area \mathbf{B} inside the Yellow domain. Without loss of generality, we may assume that $(x_t, x_{t+1}) \in \mathbf{B}_1$ (the same arguments apply to \mathbf{B}_0 symmetrically).

The proof shall use the following lemma which concerns competition between two coins, both being tossed k times. One coin has a probability of p to be “heads”, and the other has probability q . We assume that $p < q$, although the difference is very small, $q - p \leq 1/\sqrt{k}$. The lemma upper bounds the probability that the more likely coin wins, that is, that the coin with parameter q falls on “heads” more times than the other coin. The proof is deferred to Appendix A.2.2.

Lemma 12. *There exists a constant $\alpha > 1$, s.t. for every integer k , every $p, q \in [1/3, 2/3]$ with $p < q$ and $q - p \leq 1/\sqrt{k}$, we have*

$$\mathbb{P}(B_k(p) < B_k(q)) < \frac{1}{2} + \alpha(q - p)\sqrt{k} - \frac{1}{2}\mathbb{P}(B_k(p) = B_k(q)).$$

Next, let us define

$$g(x, y) = \mathbb{P}(B_\ell(y) > B_\ell(x)) + y \cdot \mathbb{P}(B_\ell(y) = B_\ell(x)) + \frac{1}{n}(1 - \mathbb{P}(B_\ell(y) \geq B_\ell(x))), \quad (7)$$

so that, conditioning on (x_t, x_{t+1}) , $\mathbb{E}(x_{t+2}) = g(x_t, x_{t+1})$ by Observation 1. We start with the following claim, which will be used to prove the subsequent claim. The proof of Claim 1 is deferred to Appendix A.3.

Claim 1. *Let $x \in [1/3, 2/3]$. On the interval $[x, x + 1/\sqrt{\ell}]$, and for ℓ large enough, $y \mapsto g(x, y) - y$ is a strictly increasing function of y .*

The next claim concerns the fixed points of $g(x, y)$ as a function of y .

Claim 2. *For any given $x \in [1/2 + 4/n, 1/2 + 4\delta]$, as a function of y , the equation $y = g(x, y)$ has at most one solution on the interval $[x, x + 1/\sqrt{\ell}]$. Moreover, in the case that it has no solution, then $g(x, x + 1/\sqrt{\ell}) < x + 1/\sqrt{\ell}$.*

Proof. First, we claim that $g(x, x) < x$. Let $p = \mathbb{P}(B_\ell(x) > B_\ell(x))$ and $q = \mathbb{P}(B_\ell(x) = B_\ell(x))$. We rearrange the definition of g slightly to obtain $g(x, x) = p + x \cdot q + \frac{p}{n}$. Moreover, $x = x \cdot (2p + q) \geq (1 + 8/n) \cdot p + x \cdot q > g(x, x)$, where the first inequality is because $x \geq 1/2 + 4/n$. Next, let $h(y) = g(x, y) - y$. Function h is continuous, and what we just showed implies $h(x) < 0$. Moreover, by Claim 1, we know that h is strictly increasing on $[x, x + 1/\sqrt{\ell}]$. Therefore, either $h(x + 1/\sqrt{\ell}) \geq 0$, in which case there is a unique $y^* \in [x, x + 1/\sqrt{\ell}]$ such that $h(y^*) = 0$; or $h(x + 1/\sqrt{\ell}) < 0$, i.e., $g(x, x + 1/\sqrt{\ell}) < x + 1/\sqrt{\ell}$, in which case the equation $y = g(x, y)$ has no solution on the interval. \square

For every $x \in [1/2 + 4/n, 1/2 + 4\delta]$, let $f(x)$ be the solution of the equation $y = g(x, y)$ in the interval $[x, x + 1/\sqrt{\ell}]$ if it exists, and $f(x) = x + 1/\sqrt{\ell}$ otherwise. Note that by Claim 2, with this definition we always have

$$g(x, f(x)) \leq f(x). \quad (8)$$

Claim 3. *For any $x \in [1/2 + 4/n, 1/2 + 4\delta]$, it holds that*

$$f(x) - x > \frac{1}{2\alpha\sqrt{\ell}} \left(x - \frac{1}{2} \right),$$

where $\alpha > 1$ is the constant stated in Lemma 12.

Proof. If $f(x)$ is not a solution to $y = g(x, y)$, then by definition $f(x) = x + 1/\sqrt{\ell}$, i.e.,

$$f(x) - x = \frac{1}{\sqrt{\ell}} > \frac{1}{2\alpha\sqrt{\ell}} \left(x - \frac{1}{2} \right),$$

and so the statement holds. Otherwise, then $f(x) = g(x, f(x))$ and belongs to $[x, x + 1/\sqrt{\ell}]$. By Lemma 12, there exists $\alpha > 0$ s.t.

$$\mathbb{P}(B_\ell(f(x)) > B_\ell(x)) < \frac{1}{2} + \alpha(f(x) - x)\sqrt{\ell} - \frac{1}{2}\mathbb{P}(B_\ell(f(x)) = B_\ell(x)).$$

This can be plugged into the definition of f (Eq. (7)) to give

$$f(x) < \frac{1}{2} + \alpha(f(x) - x)\sqrt{\ell} + \left(f(x) - \frac{1}{2} \right) \mathbb{P}(B_\ell(f(x)) = B_\ell(x)) + \frac{1}{n}$$

which we can rewrite,

$$(1 - \mathbb{P}(B_\ell(f(x)) = B_\ell(x))) \left(f(x) - \frac{1}{2} \right) < \alpha(f(x) - x)\sqrt{\ell} + \frac{1}{n}.$$

This gives

$$f(x) - x > \frac{1 - \mathbb{P}(B_\ell(f(x)) = B_\ell(x))}{\alpha\sqrt{\ell}} \left(f(x) - \frac{1}{2} \right) - \frac{1}{\alpha \cdot n\sqrt{\ell}} > \frac{1}{2\alpha\sqrt{\ell}} \left(x - \frac{1}{2} - \frac{2}{n} \right),$$

where the last inequality comes from the upper bound $\mathbb{P}(B_\ell(f(x)) = B_\ell(x)) < 1/2$ (which is true when ℓ is large enough), and from the fact that $f(x) > x$. Since $(x - 1/2) \geq 4/n$, this implies

$$f(x) - x > \frac{1}{4\alpha\sqrt{\ell}} \left(x - \frac{1}{2} \right),$$

as desired. This completes the proof of Claim 3. \square

Next, rewriting $f(x) - x = (f(x) - 1/2) - (x - 1/2)$, we get from Claim 3 that for every $x \in [1/2 + 4/n, 1/2 + 4\delta]$,

$$\left(f(x) - \frac{1}{2} \right) > \left(1 + \frac{1}{4\alpha\sqrt{\ell}} \right) \cdot \left(x - \frac{1}{2} \right). \quad (9)$$

We are now ready to conclude the proof of Lemma 9. Let $c_4 = 1/4\alpha$.

- If $x_t \in [1/2, 1/2 + 4/n]$, then by definition of \mathbf{B} , $x_{t+1} \in [1/2, 1/2 + 8/n]$. For the same reason, for (x_{t+1}, x_{t+2}) to be in \mathbf{B} , it is necessary that $x_{t+2} \in [1/2, 1/2 + 16/n]$. By Lemma 17 (see Appendix refsec:noise), there is a constant probability that it is not the case, and so (b) (in the statement of the lemma) holds.
- Otherwise, if $x_t \in [1/2 + 4/n, 1/2 + 4\delta]$ and $x_{t+1} > f(x_t)$, then by Eq. (9),

$$x_{t+1} - \frac{1}{2} > f(x_t) - \frac{1}{2} > \left(1 + \frac{c_4}{\sqrt{\ell}} \right) \left(x_t - \frac{1}{2} \right),$$

and so (a) holds.

- Else, $f(x_t) \geq x_{t+1}$. Moreover, by the definitions of f and \mathbf{B}_1 , we have the following relation:

$$x_t + 1/\sqrt{\ell} \geq f(x_t) \geq x_{t+1} \geq x_t. \quad (10)$$

By Eq. (8), $g(x_t, f(x_t)) \leq f(x_t)$, i.e., $g(x_t, f(x_t)) - f(x_t) \leq 0$. By Claim 1, for ℓ large enough, function $y \mapsto g(x_t, y) - y$ is strictly increasing on $[x_t, x_t + 1/\sqrt{\ell}]$. Eq. (10) ensures that x_{t+1} and $f(x_t)$ are within this interval, so $g(x_t, x_{t+1}) - x_{t+1} \leq g(x_t, f(x_t)) - f(x_t) \leq 0$, i.e., $g(x_t, x_{t+1}) \leq x_{t+1}$. Recall that $\mathbb{E}(x_{t+2}) = g(x_t, x_{t+1})$ – therefore, $\mathbb{E}(x_{t+2}) \leq x_{t+1}$. By Lemma 16 (see Appendix B.2), there is a constant probability c_5 that $x_{t+2} < x_{t+1}$. If this the case, since $x_{t+1} > 1/2$ and by the definition of \mathbf{B} , we get that $(x_{t+1}, x_{t+2}) \notin \mathbf{B}$ and (b) holds.

This concludes the proof of Lemma 9. \square

4 Proof of Lemma 4 — Cyan area

The goal of this section is to prove Lemma 4, which concerns Area CYAN. We distinguish between two cases.

Case 1. $x_{t_0} \geq 1/\log(n)$. In this case, by definition of CYAN_1 , we must have $x_{t_0+1} < 1/\log(n)$. Note that in this case, for n large enough, $x_{t_0+1} - \delta < 1/\log(n) - \delta < 0$. Then,

- either $x_{t_0+2} < x_{t_0+1} + \delta$. In this case, $x_{t_0+1} - \delta < 0 < x_{t_0+2} < x_{t_0+1} + \delta$, and so $(x_{t_0+1}, x_{t_0+2}) \in \text{CYAN}_1$ (but this time Case 2 applies).
- or $x_{t_0+2} \geq x_{t_0+1} + \delta$, and so $(x_{t_0+1}, x_{t_0+2}) \in \text{GREEN}_1$.
- (We can't have $x_{t_0+2} = 0$ because the source is assumed to have opinion 1.)

Case 2. $x_{t_0} < 1/\log(n)$. Let $\gamma = \gamma(c) = (1 - 1/e) \cdot \exp(-2c)/2$ and let $K = K(c) = c \cdot \exp(-2c)/2$. We will study separately three ranges of value for x_{t+1} . Claim 4 below concerns small values of x_{t+1} , Claim 5 concerns intermediate values of x_{t+1} , and Claim 6 concerns large values of x_{t+1} . Their proofs follow from simple applications of Chernoff's bound, and so they are deferred to Appendix C.4.

Claim 4. *If $x_t < 1/\log(n)$, and if $0 < x_{t+1} \leq 1/\ell$, then $\mathbb{P}(x_{t+2} > \frac{K}{2}x_{t+1} \log n) > 1 - \exp(-\frac{K}{8} \log n)$.*

Claim 5. *If $x_t < 1/\log(n)$, and if $1/\ell < x_{t+1} \leq \gamma$, then $\mathbb{P}(x_{t+2} > \gamma) > 1 - \exp(-\frac{\gamma n}{8}) > 1 - \exp(-\frac{K}{8} \log n)$.*

Claim 6. *If $x_t < 1/\log(n)$, and if $x_{t+1} > \gamma$, then $\mathbb{P}(x_{t+2} > \frac{1}{2}) > 1 - \exp(-\frac{n}{18}) > 1 - \exp(-\frac{K}{8} \log n)$.*

We say that a round t is *successful* if $(x_t, x_{t+1}) \in \text{CYAN}_1$, and the event of either Claim 4, 5 or 6 happens. Formally,

$$\begin{cases} x_{t+1} \leq 1/\ell \text{ and } x_{t+2} > Kx_{t+1} \log n \text{ (corresponding to Claim 4), or} \\ 1/\ell < x_{t+1} \leq \gamma \text{ and } x_{t+2} > \gamma \text{ (corresponding to Claim 5), or} \\ \gamma < x_{t+1} \text{ and } (x_{t+1}, x_{t+2}) \notin \text{CYAN}_1 \text{ (corresponding to Claim 6).} \end{cases}$$

Let X be the number of successful rounds starting from t_0 . This definition implies that necessarily,

$$X < \frac{\log(n/\ell)}{\log(K \cdot \log n/2)} + 2 := X_{\max}.$$

Indeed, since $x_{t_0+1} > 1/n$ (by definition of CYAN_1),

- $\log(n/\ell)/\log(K \cdot \log n/2)$ rounds are always enough to get $x_{t+1} > 1/\ell$
- one more round is enough to get $x_{t+1} > \gamma$
- one more round is enough to get $x_{t+1} > 1/2$, in which case $(x_t, x_{t+1}) \notin \text{CYAN}_1$.

Therefore, by Claims 4, 5 and 6, the probability that, starting from t_0 , all rounds are successful until the system is out of CYAN_1 is at least $(1 - \exp(-\frac{K}{8} \log n))^{X_{\max}} \geq 1 - X_{\max} \cdot \exp(-\frac{K}{8} \log n) = 1 - 1/n^{\Omega(1)}$. Moreover, for any successful round t , $x_{t+2} > x_{t+1}$ (by definition of a successful round) and $x_{t+1} < \delta + 1/\log(n)$ (this is a straightforward consequence of the definition of CYAN_1). Thus, by construction of the partition, we must have $(x_{t+1}, x_{t+2}) \in \text{CYAN}_1 \cup \text{GREEN}_1 \cup \text{PURPLE}_1$. This implies that $(x_t, x_{t+1}) \in \text{GREEN}_1 \cup \text{PURPLE}_1$, which concludes the proof of Lemma 4.

5 Discussion and Future Work

This paper considers a natural problem of information spreading in a self-stabilizing context, where it is assumed that a source agent has useful knowledge about the environment, and others would like to learn this information without being able to distinguish the source from non-source agents. Motivated by biological scenarios, our focus is on solutions that utilize passive communication. We identify an extremely simple algorithm, called FET (Protocol 1), which has a natural appeal: In each round, each (non-source) agent estimates the current tendency direction of the dynamics, and then adapts to the emerging trend. The correct operation of the algorithm does not require that the source actively cooperates with the algorithm, and instead, only assumes that it maintains its correct option throughout the execution.

Different performance parameters may be further optimized in future work. For example, our analysis uses $O(\log n)$ samples per round, and it would be interesting to see whether the problem can be solved in poly-logarithmic time w.h.p, by using only a constant number of samples per round. Also, we do not exclude the possibility that a tighter analysis of Algorithm FET would reduce our bound on the running time. In addition, our framework assumes the presence of a single source agent, but as mentioned, it can also allow for a constant number of sources, as long as it is guaranteed that all sources agree on the correct opinion. No attempt has been made to consider a larger regime of sources (beyond a constant), although we believe that such a framework is also manageable.

Finally, as a more philosophical remark, we note that early adapting to emerging trends is a common strategy in humans, which is, in some sense, encouraged by modern economic systems. For example, investing

in a successful company can yield large revenues, especially if such an investment is made before others notice its high potential. On a global scale, the collective benefits of this strategy are typically associated with economic growth. This paper shows that such a strategy can also have a collective benefit that traces back to basic aspects of collective decision-making, suggesting the possibility that it may have evolved via group-selection. With this in mind, it would be interesting to empirically check whether such a strategy exists also in other animal groups, e.g., fish schools Sumpter et al. (2008) or ants Rajendran et al. (2022).

Acknowledgements

The authors would particularly like to thank Pierre Fraigniaud and Ofer Feinerman for helpful and exciting discussions. In addition, we would like to thank Uriel Feige, Uri Zwick, Bernard Haeupler, and Emanuele Natale for preliminary discussions.

References

- Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. 2016. Time-Space Trade-offs in Population Protocols. *CoRR* abs/1602.08032 (2016). arXiv:1602.08032 <http://arxiv.org/abs/1602.08032>
- Dan Alistarh and Rati Gelashvili. 2015. Polylogarithmic-Time Leader Election in Population Protocols. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 9135)*, Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann (Eds.). Springer, 479–491. https://doi.org/10.1007/978-3-662-47666-6_38
- Dan Alistarh and Rati Gelashvili. 2018. Recent Algorithmic Advances in Population Protocols. *SIGACT News* 49, 3 (2018), 63–73. <https://doi.org/10.1145/3289137.3289150>
- Dana Angluin, James Aspnes, Zoë Diamadi, Michael J Fischer, and René Peralta. 2006. Computation in networks of passively mobile finite-state sensors. *Distributed computing* 18, 4 (2006), 235–253.
- Dana Angluin, James Aspnes, and David Eisenstat. 2008. A simple population protocol for fast robust approximate majority. *Distributed Comput.* 21, 2 (2008), 87–102. <https://doi.org/10.1007/s00446-008-0059-z>
- James Aspnes and Eric Ruppert. 2007. An Introduction to Population Protocols. *Bull. EATCS* 93 (2007), 98–117.
- Oran Ayalon, Yigal Sternklar, Ehud Fonio, Amos Korman, Nir S. Gov, and Ofer Feinerman. 2021. Sequential Decision-Making in Ants and Implications to the Evidence Accumulation Decision Model. *Frontiers Appl. Math. Stat.* 7 (2021), 672773. <https://doi.org/10.3389/fams.2021.672773>
- Robert MR Barclay. 1982. Interindividual use of echolocation calls: eavesdropping by bats. *Behavioral Ecology and Sociobiology* 10, 4 (1982), 271–275.
- Paul Bastide, George Giakkoupis, and Hayk Saribekyan. 2021. Self-Stabilizing Clock Synchronization with 1-bit Messages. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2154–2173.
- Luca Becchetti, Andrea Clementi, and Emanuele Natale. 2020. Consensus dynamics: An overview. *ACM SIGACT News* 51, 1 (2020), 58–104.
- Lucas Boczkowski, Ofer Feinerman, Amos Korman, and Emanuele Natale. 2018a. Limits for Rumor Spreading in Stochastic Populations. In *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA (LIPIcs, Vol. 94)*, Anna R. Karlin (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 49:1–49:21. <https://doi.org/10.4230/LIPIcs.ITCS.2018.49>

- Lucas Boczkowski, Amos Korman, and Emanuele Natale. 2019. Minimizing message size in stochastic communication patterns: fast self-stabilizing protocols with 3 bits. *Distributed Comput.* 32, 3 (2019), 173–191. <https://doi.org/10.1007/s00446-018-0330-x>
- Lucas Boczkowski, Emanuele Natale, Ofer Feinerman, and Amos Korman. 2018b. Limits on reliable information flows through stochastic populations. *PLoS Comput. Biol.* 14, 6 (2018). <https://doi.org/10.1371/journal.pcbi.1006195>
- Keren Censor-Hillel, Bernhard Haeupler, Jonathan Kelner, and Petar Maymounkov. 2012. Global computation in a poorly connected world: fast rumor spreading with no dependence on conductance. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing.* 961–970.
- Ho-Lin Chen, Rachel Cummings, David Doty, and David Soloveichik. 2014. Speed faults in computation by chemical reaction networks. In *International Symposium on Distributed Computing.* Springer, 16–30.
- Flavio Chierichetti, George Giakkoupis, Silvio Lattanzi, and Alessandro Panconesi. 2018. Rumor spreading and conductance. *Journal of the ACM (JACM)* 65, 4 (2018), 1–21.
- Noam Cvikel, Katya Egert Berg, Eran Levin, Edward Hurme, Ivailo Borissov, Arjan Boonman, Eran Amichai, and Yossi Yovel. 2015. Bats aggregate to improve prey search but might be impaired when their density becomes too high. *Current Biology* 25, 2 (2015), 206–211.
- Etienne Danchin, Luc-Alain Giraldeau, Thomas J Valone, and Richard H Wagner. 2004. Public information: from nosy neighbors to cultural evolution. *Science* 305, 5683 (2004), 487–491.
- Alan Demers, Dan Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. 1987. Epidemic algorithms for replicated database maintenance. In *Proceedings of the sixth annual ACM Symposium on Principles of distributed computing.* 1–12.
- Edsger W. Dijkstra. 1974. Self-stabilizing Systems in Spite of Distributed Control. *Commun. ACM* 17, 11 (1974), 643–644. <https://doi.org/10.1145/361179.361202>
- Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. 2011. Stabilizing consensus with the power of two choices. In *Proceedings of the twenty-third annual ACM symposium on Parallelism in algorithms and architectures.* 149–158.
- Chinmoy Dutta, Gopal Pandurangan, Rajmohan Rajaraman, Zhifeng Sun, and Emanuele Viola. 2013. On the complexity of information spreading in dynamic networks. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms.* SIAM, 717–736.
- Ofer Feinerman, Bernhard Haeupler, and Amos Korman. 2017. Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication. *Distributed Comput.* 30, 5 (2017), 339–355. <https://doi.org/10.1007/s00446-015-0249-4>
- Ofer Feinerman and Amos Korman. 2017. Individual versus collective cognition in social insects. *Journal of Experimental Biology* 220, 1 (2017), 73–82.
- Chryssis Georgiou, Seth Gilbert, Rachid Guerraoui, and Dariusz R. Kowalski. 2013. Asynchronous gossip. *J. ACM* 60, 2 (2013), 11:1–11:42. <https://doi.org/10.1145/2450142.2450147>
- Chryssis Georgiou, Seth Gilbert, and Dariusz R. Kowalski. 2011. Meeting the deadline: on the complexity of fault-tolerant continuous gossip. *Distributed Comput.* 24, 5 (2011), 223–244. <https://doi.org/10.1007/s00446-011-0144-6>
- George Giakkoupis. 2014. Tight bounds for rumor spreading with vertex expansion. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms.* SIAM, 801–815.
- George Giakkoupis and Philipp Woelfel. 2011. On the randomness requirements of rumor spreading. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms.* SIAM, 449–461.

- Luc-Alain Giraldeau and Thomas Caraco. 2018. *Social foraging theory*. Princeton University Press.
- Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vocking. 2000. Randomized rumor spreading. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*. IEEE, 565–574.
- Amos Korman, Efrat Greenwald, and Ofer Feinerman. 2014. Confidence Sharing: An Economic Strategy for Efficient Information Flows in Animal Groups. *PLoS Comput. Biol.* 10, 10 (2014). <https://doi.org/10.1371/journal.pcbi.1003862>
- Thomas Milton Liggett and Thomas M Liggett. 1985. *Interacting particle systems*. Vol. 2. Springer.
- Harikrishnan Rajendran, Amir Haluts, Nir S Gov, and Ofer Feinerman. 2022. Ants resort to majority concession to reach democratic consensus in the presence of a persistent minority. *Current Biology* (2022).
- Nitzan Razin, Jean-Pierre Eckmann, and Ofer Feinerman. 2013. Desert ants achieve reliable recruitment across noisy interactions. *Journal of the Royal Society Interface* 10, 82 (2013), 20130079.
- David JT Sumpter, Jens Krause, Richard James, Iain D Couzin, and Ashley JW Ward. 2008. Consensus decision making by fish. *Current Biology* 18, 22 (2008), 1773–1777.
- Gerald S Wilkinson. 1992. Information transfer at evening bat colonies. *Animal Behaviour* 44 (1992), 501–518.
- Jennifer Yick, Biswanath Mukherjee, and Dipak Ghosal. 2008. Wireless sensor network survey. *Computer networks* 52, 12 (2008), 2292–2330.

Appendix

A Probabilistic tools

A.1 Some well-known theorems

Theorem 2. [Multiplicative Chernoff's Bound] Let X_1, \dots, X_n be independent binary random variables, let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}(X)$. Then it holds for all $\delta > 0$ that

$$\mathbb{P}(X \geq (1 + \delta)\mu) \leq \exp\left(-\min\{\delta, \delta^2\} \cdot \frac{\mu}{3}\right),$$

and for all $0 < \epsilon < 1$,

$$\mathbb{P}(X \leq (1 - \epsilon)\mu) \leq \exp\left(-\epsilon^2 \cdot \frac{\mu}{2}\right).$$

Theorem 3. [Hoeffding's bound] Let X_1, \dots, X_n be independent random variables such that for every $1 \leq i \leq n$, $a_i \leq X_i \leq b_i$ almost surely. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}(X)$. Then it holds for all $\delta > 0$ that

$$\mathbb{P}(X - \mu \geq \delta) \leq \exp\left(-\frac{2\delta^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

Theorem 4. [Central Limit] Let X_1, \dots, X_n be i.i.d. random variables with $\mathbb{E}(X_1) = \mu$ and $\text{Var}(X_1) = \sigma^2 < +\infty$. Then as n tends to infinity, the random variables $\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n X_i - \mu\right)$ converges in distribution to $\mathcal{N}(0, \sigma^2)$.

Let Φ be the cumulative distribution function (c.d.f.) of the standard normal distribution:

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt.$$

Theorem 5. [Berry-Esseen] Let X_1, \dots, X_n be i.i.d. random variables, with $\mathbb{E}(X_1) = 0$, $\text{Var}(X_1) = \mathbb{E}(X_1^2) > 0$, and $\mathbb{E}(|X_1|^3) = \rho < +\infty$. Let $X = \sum_{i=1}^n X_i$ and F be the c.d.f. of $X/(\sigma\sqrt{n})$. Then it holds that

$$|F(x) - \Phi(x)| \leq \frac{C\rho}{\sigma^3\sqrt{n}},$$

where, e.g., $C = 0.4748$.

A.2 Competition between coins

Consider two coins such that one coin has a greater probability of yielding “heads”, and toss them k times each.

A.2.1 Lower bounds on the probability that the best coin wins

In Lemmas 13 and 14 we aim to lower bound the probability that the more likely coin yields more “heads”, or in other words, we lower bound the probability that the favorite coin wins. Lemma 13 is particularly effective when the difference between p and q is sufficiently large. Its proof is based on a simple application of Hoeffding's inequality.

Lemma 13. For every $p, q \in [0, 1]$ s.t. $p < q$ and every integer k , we have

$$\mathbb{P}(B_k(p) < B_k(q)) \geq 1 - \exp\left(-\frac{1}{2}k(q-p)^2\right).$$

Proof. Let $Y_i, i \in \{1, \dots, k\}$ be i.i.d. random variables with

$$Y_i = \begin{cases} 1 & \text{w.p. } p(1-q), \\ 0 & \text{w.p. } pq + (1-p)(1-q), \\ -1 & \text{w.p. } (1-p)q. \end{cases}$$

Then

$$\mathbb{P}(B_k(p) \geq B_k(q)) = \mathbb{P}\left(\sum_{i=1}^k Y_i \geq 0\right) = \mathbb{P}\left(\frac{1}{k} \sum_{i=1}^k (Y_i - (p-q)) \geq (q-p)\right).$$

Since each Y_i is bounded, and $\mathbb{E}(Y_i) = (p-q)$, we can apply Hoeffding's inequality (Theorem 3) to get

$$\mathbb{P}(B_k(p) \geq B_k(q)) \leq \exp\left(-\frac{2k^2(q-p)^2}{4k}\right) = \exp\left(-\frac{1}{2}k(q-p)^2\right).$$

□

Lemma 13 is not particularly effective when p and q are close to each other. For such cases, we shall use the following lemma.

Lemma 14. *Let $\lambda > 0$. There exist $\epsilon = \epsilon(\lambda)$ and $K = K(\lambda)$, s.t. for every $p, q \in [1/2 - \epsilon, 1/2 + \epsilon]$ with $p < q$, and every $k > K$,*

$$\mathbb{P}(B_k(p) < B_k(q)) > \frac{1}{2} + \lambda \cdot (q-p) - \frac{1}{2} \mathbb{P}(B_k(p) = B_k(q)).$$

Proof. For every $p, q \in [0, 1]$, we have

$$\begin{aligned} \mathbb{P}(B_k(q) < B_k(p)) &= \sum_{d=1}^k \mathbb{P}(B_k(q) = B_k(p) - d) \\ &= \sum_{d=1}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{\mathbb{P}(B_k(q) = B_k(p) - d)}{\mathbb{P}(|B_k(q) - B_k(p)| = d)} \\ &= \sum_{d=1}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{\mathbb{P}(B_k(q) = B_k(p) - d)}{\mathbb{P}(B_k(q) = B_k(p) - d) + \mathbb{P}(B_k(p) = B_k(q) - d)}, \end{aligned}$$

so

$$\begin{aligned} &\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) \\ &= \sum_{d=1}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{\mathbb{P}(B_k(p) = B_k(q) - d) - \mathbb{P}(B_k(q) = B_k(p) - d)}{\mathbb{P}(B_k(p) = B_k(q) - d) + \mathbb{P}(B_k(q) = B_k(p) - d)}. \end{aligned} \quad (11)$$

Let us compute $\mathbb{P}(B_k(q) = B_k(p) - d)$:

$$\begin{aligned} \mathbb{P}(B_k(q) = B_k(p) - d) &= \sum_{i=0}^{k-d} \mathbb{P}(B_k(q) = i) \cdot \mathbb{P}(B_k(p) = i + d) \\ &= \sum_{i=0}^{k-d} \binom{k}{i} \binom{k}{i+d} q^i (1-q)^{k-i} p^{i+d} (1-p)^{k-i-d} \\ &= (p(1-q))^d \sum_{i=0}^{k-d} \binom{k}{i} \binom{k}{i+d} (qp)^i ((1-q)(1-p))^{k-i-d} \\ &:= (p(1-q))^d A_{k,d,p,q}, \end{aligned}$$

where

$$A_{k,d,p,q} := \sum_{i=0}^{k-d} \binom{k}{i} \binom{k}{i+d} (qp)^i ((1-q)(1-p))^{k-i-d}.$$

Since $A_{k,d,p,q}$ is symmetric w.r.t. p, q , i.e., $A_{k,d,p,q} = A_{k,d,q,p}$, we can simplify Eq. (11) as

$$\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) = \sum_{d=1}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d}. \quad (12)$$

Intuitively, the quantity

$$\frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d}$$

can be seen as the “advantage” given by playing with the better coin (q) in a k -coin-tossing contest, knowing that one coin hit “head” d times more than the other. Before we continue, we need the following simple claim.

Claim 7. *For every $a, b \in [0, 1]$ with $a > b$, the sequence*

$$\left(\frac{a^n - b^n}{a^n + b^n} \right), n \in \mathbb{N}$$

is increasing in n .

Proof. Rewrite

$$\frac{a^n - b^n}{a^n + b^n} = \frac{2a^n}{a^n + b^n} - 1 = 2 \cdot \frac{1}{1 + (b/a)^n} - 1,$$

and notice that, since $a > b$, $((b/a)^n)$, $n \in \mathbb{N}$ is a decreasing sequence. \square

Claim 8. *Let $0 < \gamma < 1$ and $d \in \mathbb{N}$. There exists $\epsilon = \epsilon(\gamma, d)$, such that for every $p, q \in [1/2 - \epsilon, 1/2 + \epsilon]$ with $p < q$,*

$$\frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} > (q-p) \cdot 2d\gamma.$$

Proof. First, by a telescopic argument:

$$\begin{aligned} (q(1-p))^d - (p(1-q))^d &= (q(1-p) - p(1-q)) \sum_{i=0}^{d-1} (q(1-p))^{d-1-i} \cdot (p(1-q))^i \\ &= (q-p) \sum_{i=0}^{d-1} (q(1-p))^{d-1-i} \cdot (p(1-q))^i. \end{aligned}$$

Note that

$$\lim_{p,q \rightarrow 1/2} \sum_{i=0}^{d-1} (q(1-p))^{d-1-i} \cdot (p(1-q))^i = \sum_{i=0}^{d-1} \left(\frac{1}{4}\right)^{d-1-i} \cdot \left(\frac{1}{4}\right)^i = \sum_{i=0}^{d-1} \left(\frac{1}{2}\right)^{2d-2} = d \cdot \left(\frac{1}{2}\right)^{2d-2},$$

and that

$$\lim_{p,q \rightarrow 1/2} (q(1-p))^d + (p(1-q))^d = \left(\frac{1}{4}\right)^d + \left(\frac{1}{4}\right)^d = \left(\frac{1}{2}\right)^{2d-1}.$$

Hence, since $\gamma < 1$, and provided that p, q are close enough to $1/2$, we obtain

$$\frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} > (q-p) \cdot 2d\gamma,$$

which completes the proof of Claim 8. \square

Next, let $\lambda > 0$ as in the Lemma's statement, and let $\lambda' = \lambda + 1$. Denote $D = \lceil \lambda' \rceil + 1 > \lambda'$ and $\gamma = \lambda'/D < 1$. By Claim 8, there exists $\epsilon = \epsilon(\gamma, D) = \epsilon(\lambda)$, s.t. for $p, q \in [1/2 - \epsilon, 1/2 + \epsilon]$,

$$\frac{(q(1-p))^D - (p(1-q))^D}{(q(1-p))^D + (p(1-q))^D} > (q-p) \cdot 2\lambda'. \quad (13)$$

Now we derive a lower bound on Eq. (12):

$$\begin{aligned} \mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) &= \sum_{d=1}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} \quad (\text{Eq. (12)}) \\ &\geq \sum_{d=D}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} \\ &\geq \sum_{d=D}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{(q(1-p))^D - (p(1-q))^D}{(q(1-p))^D + (p(1-q))^D} \quad (\text{by Claim 7}) \\ &> (q-p) \cdot 2\lambda' \sum_{d=D}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \quad (\text{by Eq. (13)}) \\ &= (q-p) \cdot 2\lambda' \cdot (1 - \mathbb{P}(|B_k(q) - B_k(p)| < D)). \end{aligned}$$

Since $\lambda' > \lambda$, and since $\mathbb{P}(|B_k(q) - B_k(p)| < D)$ tends to 0 as k tends to $+\infty$, there exists $K = K(\lambda)$ s.t. for all $k > K$,

$$\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) > (q-p) \cdot 2\lambda. \quad (14)$$

Eventually, we write

$$\begin{aligned} \mathbb{P}(B_k(p) < B_k(q)) &= 1 - \mathbb{P}(B_k(q) < B_k(p)) - \mathbb{P}(B_k(p) = B_k(q)) \\ &> 1 - \mathbb{P}(B_k(p) < B_k(q)) + 2\lambda(q-p) - \mathbb{P}(B_k(p) = B_k(q)). \quad (\text{by Eq. (14)}) \end{aligned}$$

Hence,

$$\mathbb{P}(B_k(p) < B_k(q)) > \frac{1}{2} + \lambda(q-p) - \frac{1}{2}\mathbb{P}(B_k(p) = B_k(q)),$$

which concludes the proof of Lemma 14. \square

A.2.2 Lower bounds on the probability that the worse coin wins

We now deal with the opposite problem, that is, to lower bound the probability that the underdog coin wins. Formally,

Lemma 15. *For every $p, q \in [0, 1]$ s.t. $p < q$ and every integer k , we have*

$$\mathbb{P}(B_k(p) > B_k(q)) \geq 1 - \Phi\left(\frac{\sqrt{k}(q-p)}{\sigma}\right) - \frac{C}{\sigma\sqrt{k}},$$

where $C = 0.4748$ and $\sigma = \sqrt{p(1-p) + q(1-q)}$.

Proof. Let $Y_i, i \in \{1, \dots, k\}$ be i.i.d. random variables with

$$Y_i = \begin{cases} 1 & \text{w.p. } p(1-q), \\ 0 & \text{w.p. } pq + (1-p)(1-q), \\ -1 & \text{w.p. } (1-p)q. \end{cases}$$

Let $\mu = \mathbb{E}(Y_1) = (p - q)$, $\sigma = \sqrt{\text{Var}(Y_1)} = \sqrt{p(1-p) + q(1-q)}$, and $\rho = \mathbb{E}(|Y_1 - \mu|^3)$. We have

$$\begin{aligned} \mathbb{P}(B_k(p) > B_k(q)) &= \mathbb{P}\left(\sum_{i=1}^k Y_i > 0\right) = \mathbb{P}\left(\frac{1}{\sqrt{k}} \sum_{i=1}^k (Y_i - (p - q)) > \sqrt{k}(q - p)\right) \\ &= \mathbb{P}\left(\frac{1}{\sigma\sqrt{k}} \sum_{i=1}^k (Y_i - (p - q)) > \frac{\sqrt{k}(q - p)}{\sigma}\right) \\ &= \mathbb{P}\left(Z > \frac{\sqrt{k}(q - p)}{\sigma}\right), \end{aligned}$$

where

$$Z = \frac{1}{\sigma\sqrt{k}} \sum_{i=1}^k (Y_i - (p - q)).$$

By the Berry-Esseen theorem (Theorem 5),

$$\left| \mathbb{P}\left(Z < \frac{\sqrt{k}(q - p)}{\sigma}\right) - \Phi\left(\frac{\sqrt{k}(q - p)}{\sigma}\right) \right| < \frac{C\rho}{\sigma^3\sqrt{k}},$$

implying that

$$\left| \left(1 - \Phi\left(\frac{\sqrt{k}(q - p)}{\sigma}\right)\right) - \mathbb{P}\left(Z > \frac{\sqrt{k}(q - p)}{\sigma}\right) \right| < \frac{C\rho}{\sigma^3\sqrt{k}},$$

and so

$$\mathbb{P}\left(Z > \frac{\sqrt{k}(q - p)}{\sigma}\right) > 1 - \Phi\left(\frac{\sqrt{k}(q - p)}{\sigma}\right) - \frac{C\rho}{\sigma^3\sqrt{k}},$$

where, e.g., $C = 0.4748$. Since $\rho < \sigma^2$, we end up with

$$\mathbb{P}(B_k(p) > B_k(q)) \geq 1 - \Phi\left(\frac{\sqrt{k}(q - p)}{\sigma}\right) - \frac{C}{\sigma\sqrt{k}},$$

which concludes the proof. \square

Just as Lemma 14 was a version of Lemma 13 optimized for cases where p and q are close to each other, Lemma 12 (stated in Section 3.2) complements Lemma 15 in such situations.

Proof of Lemma 12. Recall that (see Eq. (12) in the proof of Lemma 14):

$$\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) = \sum_{d=1}^k \mathbb{P}(|B_k(q) - B_k(p)| = d) \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d}. \quad (15)$$

The following claim is analogous to Claim 8, but this time we are looking for an upper bound (instead of a lower bound) on the same quantity.

Claim 9. *There exists a constant $\alpha > 1$, s.t. for every integer k , every $p, q \in [1/3, 2/3]$ with $p < q$, and all $d \in \mathbb{N}$,*

$$\frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} < \alpha d \cdot (q - p).$$

Proof. First, we use a well-known identity:

$$\begin{aligned}
(q(1-p))^d - (p(1-q))^d &= (q(1-p) - p(1-q)) \sum_{i=0}^{d-1} (q(1-p))^{d-1-i} \cdot (p(1-q))^i \\
&= (q-p) \sum_{i=0}^{d-1} (q(1-p))^{d-1-i} \cdot (p(1-q))^i \\
&\leq d \cdot (q-p) (q(1-p))^{d-1} \\
&\leq \alpha d \cdot (q-p) (q(1-p))^d,
\end{aligned}$$

where α is any upper bound on $1/(q(1-p))$, e.g., $\alpha = 9$. Hence,

$$\frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} \leq \alpha d \cdot (q-p) \cdot \frac{(q(1-p))^d}{(q(1-p))^d + (p(1-q))^d} \leq \alpha d \cdot (q-p),$$

which concludes the proof of Claim 9. \square

Using Claim 9 on Eq. (15), we obtain

$$\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) \leq \alpha \cdot (q-p) \sum_{d=1}^k d \cdot \mathbb{P}(|B_k(q) - B_k(p)| = d). \quad (16)$$

Claim 10. For every $p, q \in [1/3, 2/3]$ with $p < q$, and every integer k ,

$$\mathbb{E}(|B_k(p) - B_k(q)|) \leq \sqrt{2kq(1-q)} + k \cdot (q-p).$$

Proof. For $i \in \{1, \dots, k\}$, let $X_i^{(1)}, X_i^{(2)} \sim \mathcal{B}(q)$ and $Y_i \sim \mathcal{B}(1-p/q)$ be independent random variables. Let

$$X^{(1)} = \sum_{i=1}^k X_i^{(1)}, \quad X^{(2)} = \sum_{i=1}^k X_i^{(2)}, \quad Z = \sum_{i=1}^k X_i^{(2)} \cdot Y_i, \quad \text{and} \quad \tilde{X}^{(2)} = \sum_{i=1}^k X_i^{(2)} \cdot (1 - Y_i) = X^{(2)} - Z.$$

Clearly, $X^{(1)} \sim \mathcal{B}_k(q)$ and $X^{(2)} \sim \mathcal{B}_k(q)$. Since for every i ,

$$X_i^{(2)} \cdot (1 - Y_i) = \begin{cases} 1 & \text{if } X_i^{(2)} = 1 \text{ and } Y_i = 0, \\ 0 & \text{otherwise,} \end{cases}$$

we obtain that $\tilde{X}^{(2)} \sim \mathcal{B}_k(q \cdot (1 - (1-p/q))) = \mathcal{B}_k(p)$. Similarly, for every i ,

$$X_i^{(2)} \cdot Y_i = \begin{cases} 1 & \text{if } X_i^{(2)} = 1 \text{ and } Y_i = 1, \\ 0 & \text{otherwise,} \end{cases}$$

hence, we obtain that $Z \sim \mathcal{B}_k(q \cdot (1 - p/q)) = \mathcal{B}_k(q-p)$. We notice that $(X^{(1)}, X^{(2)})$ are independent, as well as $(X^{(1)}, \tilde{X}^{(2)})$. Hence

$$\begin{aligned}
\mathbb{E}(|B_k(q) - B_k(p)|) &= \mathbb{E}\left(|X^{(1)} - \tilde{X}^{(2)}|\right) \\
&= \mathbb{E}\left(|X^{(1)} - X^{(2)} + Z|\right) \\
&\leq \mathbb{E}\left(|X^{(1)} - X^{(2)}| + Z\right) \\
&= \mathbb{E}\left(|X^{(1)} - X^{(2)}|\right) + \mathbb{E}(Z).
\end{aligned}$$

We have $\mathbb{E}(Z) = k(q - p)$, and

$$\begin{aligned} \mathbb{E}(|X^{(1)} - X^{(2)}|) &= \mathbb{E}\left(\sqrt{(X^{(1)} - X^{(2)})^2}\right) \\ &\leq \sqrt{\mathbb{E}\left((X^{(1)} - X^{(2)})^2\right)} \quad (\text{by Jensen inequality, and the fact that } g(x) = x^2 \text{ is convex}) \\ &= \sqrt{\text{Var}(X^{(1)} - X^{(2)})} \quad (\text{since } \mathbb{E}(X^{(1)} - X^{(2)}) = 0) \\ &= \sqrt{2kq(1 - q)}, \quad (X^{(1)}, X^{(2)} \sim \mathcal{B}_k(q) \text{ and are independent}). \end{aligned}$$

which concludes the proof of Claim 10. \square

We note that

$$\begin{aligned} \sum_{d=1}^k d \cdot \mathbb{P}(|B_k(q) - B_k(p)| = d) &= \mathbb{E}(|B_k(q) - B_k(p)|) \\ &\leq \sqrt{2kq(1 - q)} + k \cdot (q - p) \quad (\text{by Claim 10}) \\ &\leq \sqrt{2kq(1 - q)} + \sqrt{k} \quad (\text{since } q - p \leq 1/\sqrt{k}) \\ &\leq 2\sqrt{k}. \end{aligned}$$

Eventually, Eq. (16) becomes

$$\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) \leq 2\alpha \cdot (q - p)\sqrt{k}. \quad (17)$$

To conclude, we write

$$\begin{aligned} \mathbb{P}(B_k(p) < B_k(q)) &= 1 - \mathbb{P}(B_k(q) < B_k(p)) - \mathbb{P}(B_k(p) = B_k(q)) \\ &< 1 - \mathbb{P}(B_k(p) < B_k(q)) + 2\alpha \cdot (q - p)\sqrt{k} - \mathbb{P}(B_k(p) = B_k(q)). \quad (\text{by Eq. (17)}) \end{aligned}$$

Hence,

$$\mathbb{P}(B_k(p) < B_k(q)) < \frac{1}{2} + \alpha \cdot (q - p)\sqrt{k} - \frac{1}{2}\mathbb{P}(B_k(p) = B_k(q)),$$

which concludes the proof of Lemma 12. \square

A.3 Proof of Claim 1

The goal of this section is to prove Claim 1, whose statement can be found in Section 3.2. For this purpose, the key observation is that the derivative, w.r.t. x , of $\mathbb{P}(B_k(x) > B_k(p))$ in the neighborhood of p is relatively high. The following claim formalizes this idea.

Claim 11. *There exists a constant $\beta' > 0$ such that for every k large enough, and every $p, x \in [1/3, 2/3]$ satisfying $p \leq x \leq p + 1/\sqrt{k}$,*

$$\frac{d}{dx} \mathbb{P}(B_k(x) > B_k(p)) \geq \beta' \cdot \sqrt{k}.$$

Proof. Let $h > 0$. We will proceed by using a coupling argument. Let $X_i, i \in \{1, \dots, k\}$, be i.i.d. random variables uniformly distributed over the interval $[0, 1]$. Let $Y_1 = |\{i \text{ s.t. } X_i \leq x\}|$ and $Y_2 = |\{i \text{ s.t. } X_i \leq x + h\}|$. By construction, $Y_1 \sim \mathcal{B}_k(x)$ and $Y_2 \sim \mathcal{B}_k(x + h)$. Next, let $H = |\{i \text{ s.t. } x < X_i \leq x + h\}|$. By construction, $Y_2 = Y_1 + H \geq Y_1$. Let $Z \sim \mathcal{B}_k(p)$ be a binomially distributed random variable, independent from Y_1 and Y_2 . Now, we have:

$$\begin{aligned} \mathbb{P}(B_k(x + h) > B_k(p)) - \mathbb{P}(B_k(x) > B_k(p)) &= \mathbb{P}(Y_2 > Z) - \mathbb{P}(Y_1 > Z) \quad (\text{by definition of } Y_1, Y_2 \text{ and } Z) \\ &= \mathbb{P}(Y_1 \leq Z \cap Y_2 > Z) \quad (\text{because } Y_1 > Z \Rightarrow Y_2 > Z) \\ &= \sum_{j=0}^k \mathbb{P}(Z = j) \cdot \mathbb{P}(Y_1 \leq j \cap Y_2 > j). \end{aligned}$$

Let $J = \{j \in \mathbb{N} \text{ s.t. } kp \leq j \leq kp + \sqrt{k}\}$. We can rewrite the last equation as

$$\mathbb{P}(B_k(x+h) > B_k(p)) - \mathbb{P}(B_k(x) > B_k(p)) \geq \sum_{j \in J} \mathbb{P}(Z = j) \cdot \mathbb{P}(Y_1 \leq j \cap Y_2 > j). \quad (18)$$

The following result is a well-known fact.

Observation 2. *There exists a constant $\beta > 0$ such that for every k large enough, every $p \in [1/3, 2/3]$, and every i satisfying $|i - kp| \leq \sqrt{k}$, we have $\mathbb{P}(B_k(p) = i) \geq \frac{\beta}{\sqrt{k}}$.*

Proof. By the De Moivre-Laplace theorem, for any $i \in \{0, \dots, k\}$,

$$\mathbb{P}(B_k(p) = i) = \binom{k}{i} p^i (1-p)^{k-i} \approx \frac{1}{\sqrt{2kp(1-p)}} \exp\left(-\frac{(i-kp)^2}{2kp(1-p)}\right), \quad (19)$$

where we used \approx in the sense that the ratio between the left-hand side and the right-hand side tends to 1 as k tends to infinity. Since $|i - kp| \leq \sqrt{k}$,

$$\frac{1}{\sqrt{2kp(1-p)}} \exp\left(-\frac{(i-kp)^2}{2kp(1-p)}\right) \geq \frac{1}{\sqrt{2kp(1-p)}} \exp\left(-\frac{1}{2p(1-p)}\right) := \frac{f(p)}{\sqrt{k}}.$$

By Eq. (19), we can conclude the proof of Observation 2 for k large enough by taking, e.g.,

$$\beta = \frac{1}{2} \cdot \min_{p \in [1/3, 2/3]} f(p).$$

□

For $j \in J$, by Observation 2, $\mathbb{P}(Z = j) \geq \beta/\sqrt{k}$, for some constant $\beta > 0$. Moreover,

$$\begin{aligned} \mathbb{P}(Y_1 \leq j \cap Y_2 > j) &\geq \mathbb{P}(Y_1 = j \cap Y_2 > j) \\ &= \mathbb{P}(Y_1 = j \cap H \geq 1) && \text{(because } Y_2 = Y_1 + H\text{)} \\ &= \mathbb{P}(Y_1 = j) \cdot \mathbb{P}(H \geq 1 \mid Y_1 = j). \end{aligned}$$

By the assumption in the lemma, $p \leq x \leq p + 1/\sqrt{k}$, and so $kp \leq kx \leq kp + \sqrt{k}$. Therefore, for $j \in J$, $|j - kx| \leq \sqrt{k}$, and by Observation 2, we get that $\mathbb{P}(Y_1 = j) \geq \beta/\sqrt{k}$. Hence, we can rewrite Eq. (18) as

$$\mathbb{P}(B_k(x+h) > B_k(p)) - \mathbb{P}(B_k(x) > B_k(p)) \geq \frac{\beta^2}{k} \sum_{j \in J} \mathbb{P}(H \geq 1 \mid Y_1 = j). \quad (20)$$

Now, let us find a lower bound on $\mathbb{P}(H \geq 1 \mid Y_1 = j)$, for $j \in J$. Note that, by definition, $Y_1 = j$ if and only if $|\{i \text{ s.t. } X_i > x\}| = k - j$. Since $X_i, 1 \leq i \leq k$, is uniformly distributed over $[0, 1]$,

$$\mathbb{P}(x < X_i \leq x+h \mid X_i > x) = \frac{h}{1-x}.$$

Therefore, for every $j \in J$

$$\mathbb{P}(H = 0 \mid Y_1 = j) = \left(1 - \frac{h}{1-x}\right)^{k-j} \leq \left(1 - \frac{h}{1-x}\right)^{k-kp-\sqrt{k}}.$$

This implies that

$$\sum_{j \in J} \mathbb{P}(H \geq 1 \mid Y_1 = j) \geq \sqrt{k} \cdot \left(1 - \left(1 - \frac{h}{1-x}\right)^{k-kp-\sqrt{k}}\right).$$

We have

$$\lim_{h \rightarrow 0} \frac{1}{h} \cdot \sum_{j \in J} \mathbb{P}(H \geq 1 \mid Y_1 = j) \geq \lim_{h \rightarrow 0} \frac{\sqrt{k}}{h} \cdot \left(1 - \left(1 - \frac{h}{1-x}\right)^{k-kp-\sqrt{k}}\right) = \frac{\sqrt{k}(k-kp-\sqrt{k})}{1-x}.$$

Eventually, we get from Eq. (20)

$$\frac{d}{dx} \mathbb{P}(B_k(x) > B_k(p)) \geq \frac{\beta^2(1-p)}{1-x} \cdot \sqrt{k} + \underset{k \rightarrow \infty}{o}(\sqrt{k}).$$

We can conclude the proof of Claim 11 for k large enough by taking, e.g., $\beta' = \frac{\beta^2(1-p)}{2(1-x)}$. \square

Now, we are ready to write the proof of Claim 1.

Proof of Claim 1. We can rewrite Eq. (7) as

$$g(x, y) = \left(y - \frac{1}{n}\right) \mathbb{P}(B_\ell(y) \geq B_\ell(x)) + (1-y) \cdot \mathbb{P}(B_\ell(y) > B_\ell(x)) + \frac{1}{n}.$$

Hence,

$$\begin{aligned} \frac{d}{dy} g(x, y) &= \left[\mathbb{P}(B_\ell(y) \geq B_\ell(x)) - \mathbb{P}(B_\ell(y) > B_\ell(x)) \right] \\ &\quad + \left(y - \frac{1}{n}\right) \cdot \frac{d}{dy} \mathbb{P}(B_\ell(y) \geq B_\ell(x)) \\ &\quad + (1-y) \cdot \frac{d}{dy} \mathbb{P}(B_\ell(y) > B_\ell(x)). \end{aligned} \quad (21)$$

The first term in Eq. (21) is equal to $\mathbb{P}(B_\ell(y) = B_\ell(x))$, which is positive. Moreover, $\mathbb{P}(B_\ell(y) \geq B_\ell(x))$ is obviously increasing in y , so the second term is also non-negative. By Claim 11, the third term in Eq. (21) satisfies

$$(1-y) \cdot \frac{d}{dy} \mathbb{P}(B_\ell(y) > B_\ell(x)) \geq (1-y) \cdot \beta' \cdot \sqrt{\ell} \geq \frac{\beta'}{4} \cdot \sqrt{\ell},$$

where the last inequality comes from the fact that $x \in [1/3, 2/3]$ and $y \in [x, x + 1/\sqrt{\ell}] \subseteq [1/4, 3/4]$. For ℓ large enough, this implies that

$$\frac{d}{dy} g(x, y) \geq \frac{\beta'}{4} \cdot \sqrt{\ell} > 1,$$

which concludes the proof of Claim 1. \square

B Basic Observations regarding Algorithm FET

B.1 Proof of Observation 1

Let I be the set of agents (including the source). Let $I_t^1 \subset I$ be the set of all *non-source* agents with opinion 1 at round t . Recall that we condition on $x_t = \mathbf{x}_t$ and $x_{t+1} = \mathbf{x}_{t+1}$ (although we avoid writing this conditioning). In addition, the proof will proceed by conditioning on $I_{t+1}^1 = \mathbf{I}_{t+1}^1$. Since we shall show that the statements are true for every \mathbf{I}_{t+1}^1 , the lemma will hold without this latter conditioning.

By definition of the protocol, and because it operates under the \mathcal{PULL} model, $\text{count}^{(i)}$ and $\text{count}''^{(i)}$ are obtained by sampling ℓ agents uniformly at random in the population (with replacement) and counting how many have opinion 1. Therefore, conditioning on (x_t, x_{t+1}) and I_{t+1}^1 ,

- (i) variables $(\text{count}'_{t+1}{}^{(i)})_{i \in I}$ and $(\text{count}''_t{}^{(i)})_{i \in I}$ are mutually independent, thus variables $(Y_{t+2}^{(i)})_{i \in I}$ are mutually independent.
- (ii) for every $i \in I$, $\text{count}'_{t+1}{}^{(i)} \sim \mathcal{B}_\ell(x_{t+1})$, and $\text{count}''_t{}^{(i)} \sim \mathcal{B}_\ell(x_t)$, so we can write for every non-source agent $i \in I_{t+1}^1$,

$$\mathbb{P}\left(Y_{t+2}^{(i)} = 1\right) = \mathbb{P}\left(B_\ell(x_{t+1}) \geq B_\ell(x_t)\right),$$

and for every non-source agent $i \notin I_{t+1}^1$,

$$\mathbb{P}\left(Y_{t+2}^{(i)} = 1\right) = \mathbb{P}\left(B_\ell(x_{t+1}) > B_\ell(x_t)\right).$$

This establishes Eq. 1. Now, let us define independent binary random variables $(X_j)_{1 \leq j \leq n}$, taking values in $\{0, 1\}$, as follows;

- $X_1 = 1$,
- for every j s.t. $1 < j \leq n \cdot x_{t+1}$, $\mathbb{P}(X_j = 1) = \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t))$,
- for every j s.t. $n \cdot x_{t+1} < j \leq n$, $\mathbb{P}(X_j = 1) = \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t))$.

We assumed the source agent to have opinion 1, so there are $nx_t - 1$ non-source agents with opinion 1 and $n(1 - x_t)$ non-source agents with opinion 0. Therefore, by (i) and (ii) and by construction of the $(X_j)_{1 \leq j \leq n}$, $x_{t+2} = \frac{1}{n} \sum_{i \in I} Y_{t+2}^{(i)}$ is distributed as $\frac{1}{n} \sum_{j=1}^n X_j$, which establishes the second statement in Observation 1. Computing the expectation (still conditioning on x_t, x_{t+1}) is straightforward and does not depend on I_{t+1}^1 :

$$\begin{aligned} \mathbb{E}(x_{t+2}) &= \left(x_{t+1} - \frac{1}{n}\right) \cdot \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t)) + (1 - x_{t+1}) \cdot \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + \frac{1}{n} \\ &= x_{t+1} \cdot \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t)) + (1 - x_{t+1}) \cdot \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + \frac{1}{n}(1 - \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t))) \\ &= \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + x_{t+1} \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) + \frac{1}{n}(1 - \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t))). \end{aligned}$$

This establishes Eq. (2), and concludes the proof of Observation 1. \square

Remark 2. From Observation 1, we obtain the following straightforward bounds: for every non-source agent i ,

$$\mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) \leq \mathbb{P}(Y_{t+2}^{(i)} = 1) \leq \mathbb{P}(B_\ell(x_{t+1}) \leq B_\ell(x_t)), \quad (22)$$

and

$$\mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) - \frac{1}{n} \leq \mathbb{E}(x_{t+2}) \leq \mathbb{P}(B_\ell(x_{t+1}) \leq B_\ell(x_t)) + \frac{1}{n}. \quad (23)$$

Because of the source agent having opinion 1, the left hand side in Eq. (23) is loose (specifically, $-1/n$ is not necessary). Nevertheless, we will use this equation in the proofs, because it has a symmetric equivalent (w.r.t. to the center of \mathcal{G} , $(\frac{1}{2}, \frac{1}{2})$) which will allow our statements about x_{t+2} to hold symmetrically for $1 - x_{t+2}$, despite the asymmetry induced by the source.

Remark 3. Eq. (2) in Observation 1 implies the following convenient bounds:

$$\begin{aligned} \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + x_{t+1} \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) - \frac{1}{n} \\ < \mathbb{E}(x_{t+2}) < \\ \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + x_{t+1} \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) + \frac{1}{n}. \end{aligned} \quad (24)$$

B.2 Affects of noise

When it comes to the central area, YELLOW, we will need the following result to break ties.

Lemma 16. There exists a constant $\beta > 0$ s.t. for n large enough, and if $\mathbb{E}(x_{t+2}) \in [1/3, 2/3]$, then

$$\mathbb{P}(x_{t+2} \leq \mathbb{E}(x_{t+2}) - 1/\sqrt{n}), \mathbb{P}(x_{t+2} \geq \mathbb{E}(x_{t+2}) + 1/\sqrt{n}) \geq \beta.$$

Proof. Consider X_1, \dots, X_n from the statement of Observation 1. We have (see the proof of Observation 1)

- $X_1 = 1$,
- for every j s.t. $1 < j \leq n \cdot x_{t+1}$, $\mathbb{P}(X_j = 1) = \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t))$,
- for every j s.t. $n \cdot x_{t+1} < j \leq n$, $\mathbb{P}(X_j = 1) = \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t))$.

Let $p = \mathbb{P}(B_\ell(x_{t+1}) \geq B_\ell(x_t))$ and $q = \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t))$. By Observation 1,

$$\mathbb{E}(x_{t+2}) = \mathbb{E}\left(\frac{1}{n} \sum_{i=1}^n X_i\right) = x_{t+1} \cdot p + (1 - x_{t+1}) \cdot q + \frac{1}{n}(1 - p).$$

By assumption on $\mathbb{E}(x_{t+2})$, this implies that

$$x_{t+1} \cdot p + (1 - x_{t+1}) \cdot q \in \left[\frac{1}{3} - \frac{1}{n}, \frac{2}{3}\right].$$

Moreover, we have that $p - q = \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t))$ which tends to 0 as n tends to infinity, i.e., p and q are arbitrarily close. Hence, for n large enough, the last equation implies that $p \in [1/4, 3/4]$ and $q \in [1/4, 3/4]$. Let $Y_p = \sum_{i=2}^{n \cdot x_{t+1}} X_i$ and $Y_q = \sum_{i=n \cdot x_{t+1} + 1}^n X_i$. These two variables are binomially distributed, and since $p, q \in [1/4, 3/4]$, there is a constant probability that $Y_p \geq \mathbb{E}(Y_p)$, and there is a constant probability that $Y_q \geq \mathbb{E}(Y_q)$ as well. Without loss of generality, we assume that $x_{t+1} \geq 1/2$ and focus on Y_p (if $x_{t+1} < 1/2$, then we could consider Y_q instead). Let $m = n \cdot x_{t+1} - 1$ be the number of samples of Y_p . In this case, $m \geq n/2 - 1$ tends to $+\infty$ as n tends to $+\infty$. Let $\sigma_p = \sqrt{p(1-p)}$. By the central limit theorem (Theorem 4), the random variable

$$\frac{\sqrt{m}}{\sigma_p} \left(\frac{1}{m} Y_p - p \right) = \frac{Y_p - \mathbb{E}(Y_p)}{\sigma_p \sqrt{m}}$$

converges in distribution to $\mathcal{N}(0, 1)$. Moreover, $\text{Var}(Y_p) = m\sigma_p^2 = (n \cdot x_{t+1} - 1)p(1-p) \geq (n/2 - 1)p(1-p) \geq np(1-p)/3$, so for any $\epsilon > 0$ and n large enough,

$$\begin{aligned} \mathbb{P}(Y_p \geq \mathbb{E}(Y_p) + \sqrt{n}) &= \mathbb{P}\left(\frac{Y_p - \mathbb{E}(Y_p)}{\sigma_p \sqrt{m}} \geq \frac{\sqrt{n}}{\sigma_p \sqrt{m}}\right) \\ &\geq \mathbb{P}\left(\frac{Y_p - \mathbb{E}(Y_p)}{\sigma_p \sqrt{m}} \geq \sqrt{\frac{3}{p(1-p)}}\right) \\ &\geq 1 - \Phi\left(\sqrt{\frac{3}{p(1-p)}}\right) - \epsilon. \end{aligned}$$

For ϵ small enough, and because p is bounded, this probability is bounded away from zero. This concludes the proof of Lemma 16 (the other inequality can be obtained symmetrically). \square

We can use the previous result to show that the Markov process (x_t, x_{t+1}) is sufficiently noisy so that it is never too likely to be at any given point (x, y) .

Lemma 17. *There is a constant $c_1 = c_1(c) > 0$ (recall that the sample size is $\ell = c \cdot \log n$), such that for any $a \in [1/2 - 4\delta, 1/2 + 4\delta]$, and any round t s.t. $(x_t, x_{t+1}) \in \text{YELLOW}'$, we have*

$$\mathbb{P}\left(|x_{t+2} - a| > \frac{1}{\sqrt{n}}\right) > c_1.$$

Proof. Follows directly from Lemma 16. \square

C Analyzing Domains

C.1 Green area

Proof of Lemma 1. Let us prove the first part and assume $(x_t, x_{t+1}) \in \text{GREEN}_1$ (the proof of the second part is analogous). By Eq. (22) in Remark 2, we have for every agent i

$$\mathbb{P}\left(Y_i^{(t+2)} = 0\right) \leq \mathbb{P}(B_\ell(x_{t+1}) \leq B_\ell(x_t)).$$

By Lemma 13, we have

$$\mathbb{P}(B_\ell(x_{t+1}) \leq B_\ell(x_t)) \leq \exp\left(-\frac{1}{2}\ell(x_{t+1} - x_t)^2\right) \leq \exp\left(-\frac{1}{2}\ell\delta^2\right) = \exp\left(-\frac{c\delta^2}{2}\log n\right).$$

Then, by the union bound,

$$\mathbb{P}\left(\bigcup_{i \in I \setminus \{\text{source}\}} (Y_i^{(t+2)} = 0)\right) \leq (n-1) \cdot \exp\left(-\frac{c\delta^2}{2}\log n\right),$$

which tends to 0 as n goes to $+\infty$, provided that $c > 2/\delta^2$. \square

C.2 Purple area

Proof of Lemma 2. Let us prove the first part and assume $(x_t, x_{t+1}) \in \text{PURPLE}_1$ (the proof of the second part is analogous). By Eq. (23) in Remark 2,

$$\mathbb{E}(x_{t+2}) \geq \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) - \frac{1}{n}.$$

Since $(x_t, x_{t+1}) \in \text{PURPLE}_1$, and since in this area x_{t+1} is at least a factor $(1 - \lambda_n)$ greater than x_t , we have

$$\mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) \geq \mathbb{P}(B_\ell((1 - \lambda_n)x_t) > B_\ell(x_t)).$$

Let

$$\sigma = \sqrt{x_t(1 - x_t) + (1 - \lambda_n)x_t(1 - (1 - \lambda_n)x_t)} > \sqrt{x_t(1 - x_t)} > \sqrt{\frac{x_t}{2}}, \quad (25)$$

where the last inequality is by the fact that $x_t < 1/2$ which follows from the definition of PURPLE_1 . By Lemma 15,

$$\mathbb{P}(B_\ell((1 - \lambda_n)x_t) > B_\ell(x_t)) > 1 - \Phi\left(\frac{\sqrt{\ell}\lambda_n x_t}{\sigma}\right) - \frac{C}{\sigma\sqrt{\ell}}.$$

We have (Eq. (25) and definition of PURPLE_1)

$$\sigma > \sqrt{\frac{x_t}{2}} > \sqrt{\frac{1}{2\log n}}$$

so

$$\frac{C}{\sigma\sqrt{\ell}} < \frac{\sqrt{2}C}{\sqrt{c}}.$$

If c is large enough (specifically, if $c > 32C^2/\delta^2$), we obtain

$$\mathbb{P}(B_\ell((1 - \lambda_n)x_t) > B_\ell(x_t)) > 1 - \Phi\left(\frac{\sqrt{\ell}\lambda_n x_t}{\sigma}\right) - \frac{\delta}{4}.$$

We have

$$0 \leq \frac{\sqrt{\ell}\lambda_n x_t}{\sigma} \leq \sqrt{\ell}\lambda_n \sqrt{2x_t} \leq \sqrt{\ell}\lambda_n = \frac{\sqrt{c}}{\log^\delta n} \xrightarrow{n \rightarrow +\infty} 0.$$

where the second inequality is by Eq. (25), and the third is because $x_t < 1/2$. So, for n large enough

$$1 - \Phi\left(\frac{\sqrt{\ell}\lambda_n}{\sigma}\right) - \frac{\delta}{4} > 1 - \Phi(0) - \frac{\delta}{2} = \frac{1 - \delta}{2}.$$

Overall, we have proved that if n is large enough, then $\mathbb{E}(x_{t+2}) > (1 - \delta)/2$. By Observation 1, we can apply Chernoff's inequality (Theorem 2) to get that $x_{t+2} > 1/2 - \delta$ w.h.p. Since by definition of PURPLE_1 we have $1/2 - \delta > x_{t+1} + \delta$, we obtain $x_{t+2} > x_{t+1} + \delta$ w.h.p., which concludes the proof of the lemma. \square

C.3 Red area

Proof of Lemma 3. Without loss of generality, we assume that $t_0 = 0$. We assume that $(x_0, x_1) \in \text{RED}_1$ (the proof in the case that $(x_0, x_1) \in \text{RED}_0$ is the same). First we note that for every round t , by definition, if $(x_t, x_{t+1}) \in \text{RED}_1$ then $x_{t+1} < (1 - \lambda_n)x_t$. So, we can prove by induction on t that for every $1 \leq t \leq t_1$,

$$x_t < x_0(1 - \lambda_n)^t.$$

In particular, we have that $x_{t_1} < x_{t_0} < 1/2 - 3\delta$, and so $(x_{t_1}, x_{t_1+1}) \notin \text{YELLOW}$ by definition of **YELLOW**.

Also by definition, $x_0 < 1/2$ and $x_t > 1/\log(n)$ for every $0 \leq t \leq t_1$, hence, we obtain from the last equation that

$$\frac{1}{\log n} < \frac{1}{2}(1 - \lambda_n)^t.$$

Taking the logarithm and rearranging, we get

$$\log\left(\frac{1}{2}\right) + \log(\log n) > t \cdot \log\left(\frac{1}{1 - \lambda_n}\right).$$

We know that $\log(1 - \lambda_n) < -\lambda_n$ and thus $t \cdot \log(1/(1 - \lambda_n)) > t\lambda_n$. Together with the above equation, this gives

$$t < \frac{1}{\lambda_n} \left(\log\left(\frac{1}{2}\right) + \log(\log n) \right) = o\left(\log^{1/2+2\delta} n\right),$$

which concludes the proof. \square

C.4 Cyan area

Proof of Claim 4. We note that, since $x_t < 1/\log(n)$, the probability that an agent does not see a 1 in round t is

$$\mathbb{P}(B_\ell(x_t) = 0) = (1 - x_t)^\ell > \left(1 - \frac{1}{\log n}\right)^\ell = \exp\left(c \log(n) \log\left(1 - \frac{1}{\log n}\right)\right) > e^{-2c},$$

for n large enough. Moreover,

$$(1 - x_{t+1})^\ell < 1 - \ell x_{t+1} + \frac{1}{2}\ell^2 x_{t+1}^2,$$

so the probability that an agent sees at least a 1 in round $t + 1$ is

$$\mathbb{P}(B_\ell(x_{t+1}) \geq 1) = 1 - (1 - x_{t+1})^\ell > \ell x_{t+1} \left(1 - \frac{1}{2}\ell x_{t+1}\right) > \frac{1}{2}\ell x_{t+1},$$

where the last inequality comes from the assumption that $x_{t+1} \leq 1/\ell$. Eventually, we can write

$$\mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) \geq \mathbb{P}(B_\ell(x_t) = 0) \cdot \mathbb{P}(B_\ell(x_{t+1}) \geq 1) \geq \frac{c}{2} \cdot e^{-2c} \cdot x_{t+1} \log n = K x_{t+1} \log n.$$

Hence, by Eq. (23) in Remark 2, $\mathbb{E}(x_{t+2}) \geq K x_{t+1} \log n - 1/n$. By Observation 1, we can apply Chernoff's inequality (Theorem 2) to conclude the proof of Claim 4. \square

Proof of Claim 5. The proof follows along similar lines as the proof of Claim 4. We note that, since $x_t < 1/\log(n)$, the probability that an agent does not see a 1 in round t is

$$\mathbb{P}(B_\ell(x_t) = 0) = (1 - x_t)^\ell > \left(1 - \frac{1}{\log n}\right)^\ell = \exp\left(c \log n \log\left(1 - \frac{1}{\log n}\right)\right) > e^{-2c},$$

for n large enough. Moreover, the probability that an agent sees at least a 1 in round $t + 1$ is

$$\mathbb{P}(B_\ell(x_{t+1}) \geq 1) = 1 - (1 - x_{t+1})^\ell \geq 1 - \left(1 - \frac{1}{\ell}\right)^\ell > 1 - \frac{1}{e}.$$

Eventually, we can write

$$\mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) \geq \mathbb{P}(B_\ell(x_{t+1}) \geq 1) \cdot \mathbb{P}(B_\ell(x_t) = 0) \geq e^{-2c} \cdot \left(1 - \frac{1}{e}\right) = 2\gamma.$$

Hence, by Eq. (23) in Remark 2, $\mathbb{E}(x_{t+2}) \geq 2\gamma - 1/n$. By Observation 1, we can apply Chernoff's inequality (Theorem 2) to conclude the proof of Claim 5. \square

Proof of Claim 6. By assumption, $x_{t+1} - x_t \geq \gamma - 1/\log(n)$, and so by Lemma 13

$$\mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) \geq 1 - \exp\left(-\frac{1}{2}\ell\left(\gamma - \frac{1}{\log n}\right)^2\right) > \frac{3}{4}$$

for n large enough. Hence, by Eq. (23) in Remark 2, $\mathbb{E}(x_{t+2}) \geq 3/4 - 1/n$. By Observation 1, we can apply Chernoff's inequality (Theorem 2) to conclude the proof of Claim 6. \square

C.5 Yellow area

Recall the partitioning of the Yellow domain as illustrated in Figure 2. We analyze each of the resulting areas separately.

C.5.1 Area A

Proof of Lemma 7. Without loss of generality, we assume that $(x_t, x_{t+1}) \in \mathbf{A}_1$ (the same arguments apply to \mathbf{A}_0 symmetrically). We have, provided that δ is small enough and n is large enough,

$$\begin{aligned} \mathbb{E}(x_{t+2}) &> \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + x_{t+1} \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) - \frac{1}{n} && \text{(by Remark 3)} \\ &> \frac{1}{2} + 6(x_{t+1} - x_t) + \left(x_{t+1} - \frac{1}{2}\right) \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) && \text{(by Lemma 14, taking } \lambda > 6) \\ &> \frac{1}{2} + 6(x_{t+1} - x_t). && \text{(by the definition of } \mathbf{A}_1) \end{aligned}$$

Hence,

$$\mathbb{E}(x_{t+2}) - x_{t+1} > \frac{1}{2} - x_t + 5(x_{t+1} - x_t) = (x_{t+1} - (2x_t - \frac{1}{2})) + 4(x_{t+1} - x_t),$$

and by definition of \mathbf{A}_1 , $(x_{t+1} - (2x_t - 1/2)) \geq 0$ and so

$$\mathbb{E}(x_{t+2}) > 4(x_{t+1} - x_t) + x_{t+1}. \tag{26}$$

By Observation 1, we can apply Chernoff's inequality (Theorem 2). Taking $\epsilon = 2(x_{t+1} - x_t)/(4(x_{t+1} - x_t) + x_{t+1})$, we have

$$\begin{aligned} \mathbb{P}(x_{t+2} - x_{t+1} \leq 2(x_{t+1} - x_t)) &= \mathbb{P}(x_{t+2} \leq (1 - \epsilon)(4(x_{t+1} - x_t) + x_{t+1})) \\ &\leq \mathbb{P}(nx_{t+2} \leq (1 - \epsilon)\mathbb{E}(nx_{t+2})) && \text{(by Eq. (26))} \\ &\leq \exp\left(-\frac{\epsilon^2}{2}\mathbb{E}(nx_{t+2})\right) && \text{(by Theorem 2)} \\ &\leq \exp\left(-\frac{2x_{t+1}}{(4(x_{t+1} - x_t) + x_{t+1})^2}(x_{t+1} - x_t)^2n\right). && \text{(by Eq. (26) and definition of } \epsilon) \end{aligned}$$

Since x_t and x_{t+1} are close to $1/2$, we have for δ small enough

$$\mathbb{P}(x_{t+2} - x_{t+1} > 2(x_{t+1} - x_t)) \geq 1 - \exp(-3(x_{t+1} - x_t)^2n). \tag{27}$$

Now, we show that the event " $x_{t+2} - x_{t+1} > 2(x_{t+1} - x_t)$ " suffices for (x_{t+1}, x_{t+2}) to remain in \mathbf{A}_1 or leave YELLOW'.

Claim 12. *If $(x_t, x_{t+1}) \in \mathbf{A}_1$ and $x_{t+2} - x_{t+1} > 2(x_{t+1} - x_t)$, then $(x_{t+1}, x_{t+2}) \in \mathbf{A}_1$ or $(x_{t+1}, x_{t+2}) \notin \text{YELLOW}'$.*

Proof. If $(x_{t+1}, x_{t+2}) \notin \text{YELLOW}'$, the result holds. Otherwise, $(x_{t+1}, x_{t+2}) \in \text{YELLOW}'$ and we have to prove that (x_{t+1}, x_{t+2}) satisfies \mathbf{A}_1 .(i) and \mathbf{A}_1 .(ii). First we prove that (x_{t+1}, x_{t+2}) satisfies \mathbf{A}_1 .(i):

$$\begin{aligned} x_{t+2} &> x_{t+1} + 2(x_{t+1} - x_t) && \text{(by assumption in the claim)} \\ &\geq x_{t+1} && \text{(because } (x_t, x_{t+1}) \in \mathbf{A}_1 \Rightarrow x_{t+1} \geq x_t) \\ &\geq 1/2. && \text{(because } (x_t, x_{t+1}) \in \mathbf{A}_1 \text{ and by } \mathbf{A}_1\text{.(i))} \end{aligned}$$

Then we prove that (x_{t+1}, x_{t+2}) satisfies \mathbf{A}_1 .(ii):

$$\begin{aligned} x_{t+2} - x_{t+1} &> 2(x_{t+1} - x_t) && \text{(by assumption in the claim)} \\ &> (x_{t+1} - x_t) + (x_t - 1/2) && \text{(because } (x_t, x_{t+1}) \in \mathbf{A}_1 \text{ and by } \mathbf{A}_1\text{.(ii))} \\ &= x_{t+1} - 1/2, \end{aligned}$$

which concludes the proof of Claim 12. \square

Next, we apply Claim 12 to Eq. (27) to establish (a). Eventually, $x_{t+2} > x_{t+1} + 4(x_{t+1} - x_t) + 1/\sqrt{n}$ implies $x_{t+2} - x_{t+1} > +2(x_{t+1} - x_t)$ so we can use Claim 12,

$$\begin{aligned} &\mathbb{P}((x_{t+1}, x_{t+2}) \notin \text{YELLOW}' \setminus \mathbf{A}_1 \cap x_{t+2} > x_{t+1} + 4(x_{t+1} - x_t) + 1/\sqrt{n}) \\ &= \mathbb{P}(x_{t+2} > x_{t+1} + 4(x_{t+1} - x_t) + 1/\sqrt{n}) && \text{(by Claim 12)} \\ &> \mathbb{P}(x_{t+2} > \mathbb{E}(x_{t+2}) + 1/\sqrt{n}) && \text{(by Eq. (26))} \\ &> c_2 > 0, \end{aligned}$$

where the existence of c_2 is guaranteed by Lemma 16. This establishes (b). \square

Proof of Lemma 8. Without loss of generality, we assume that $(x_t, x_{t+1}) \in \mathbf{A}_1$ (the same arguments apply to \mathbf{A}_0 symmetrically). Let us define event H_{t_0+1} , that the system is either in A_1 or out of YELLOW' in round $t_0 + 1$, and that the ‘‘gap’’ $(x_{t_0+2} - x_{t_0+1})$ is not too small. Formally,

$$H_{t_0+1} : (x_{t_0+1}, x_{t_0+2}) \notin \text{YELLOW}' \setminus \mathbf{A}_1 \cap x_{t_0+2} - x_{t_0+1} > 1/\sqrt{n}.$$

For $t > t_0 + 1$, we define event H_t , that the system is either in A_1 or out of YELLOW' in round t , and that the gap $(x_{t+1} - x_t)$ doubles. Formally,

$$H_t : (x_t, x_{t+1}) \notin \text{YELLOW}' \setminus \mathbf{A}_1 \cap x_{t+1} - x_t > 2(x_t - x_{t-1}).$$

We start with the following observation, which results directly from the definition of H_t for $t \geq t_0 + 1$:

$$\bigcap_{s=t_0+1}^{t-1} H_s \Rightarrow (x_t - x_{t-1}) > 2^{(t-t_0-2)}(x_{t_0+2} - x_{t_0+1}) \Rightarrow (x_t - x_{t-1}) > 2^{(t-t_0-2)}/\sqrt{n}. \quad (28)$$

For every $t > t_0 + 1$,

$$\begin{aligned} \mathbb{P}\left(H_t \mid \bigcap_{s=t_0+1}^{t-1} H_s\right) &> 1 - \exp(-3n \cdot (x_t - x_{t-1})^2) && \text{(By Lemma 7)} \\ &> 1 - \exp\left(-\frac{3}{4} \cdot 4^{(t-t_0-1)}\right). && \text{(by Eq. (28))} \end{aligned}$$

By Lemma 7 (b), $(x_{t_0+1}, x_{t_0+2}) \in \mathbf{A}_1$ and $x_{t_0+2} - x_{t_0+1} > 1/\sqrt{n}$ w.p. $c_2 > 0$. Together with the last equation and using the union bound, we get

$$\mathbb{P}\left(\bigcap_{t=t_0+1}^{t_1} H_t\right) > c_2 \cdot \left(1 - \sum_{t=t_0+2}^{t_1} \exp\left(-\frac{3}{4} \cdot 4^{(t-t_0-1)}\right)\right).$$

We have the following very rough upper bounds

$$\sum_{t=t_0+2}^{t_1} \exp\left(-\frac{3}{4} \cdot 4^{(t-t_0-1)}\right) < \sum_{t=t_0+2}^{t_1} \exp\left(-\frac{3}{4} \cdot 4 \cdot (t-t_0-1)\right) < 2 \cdot e^{-3}.$$

Hence, we have proved that for every $t_1 > t_0 + 1$,

$$\mathbb{P}\left(\bigcap_{t=t_0+1}^{t_1} H_t\right) > c_2 \cdot (1 - 2 \cdot e^{-3}) := c_3 > 0.$$

By Eq. (28), it implies that for every $t_1 > t_0 + 1$,

$$\mathbb{P}\left((x_{t_1} - x_{t_1-1}) > 2^{(t_1-t_0-2)}/\sqrt{n}\right) > c_3.$$

For t_1 large enough (e.g., $t_1 = t_0 + \log n$), this implies that $(x_{t_1-1}, x_{t_1}) \notin \text{YELLOW}'$, otherwise the gap $(x_{t_1} - x_{t_1-1})$ would be greater than 8δ which is the diameter of YELLOW' . This concludes the proof of Lemma 8. \square

C.5.2 Area B: Proof of Lemma 10

Without loss of generality, we assume that $(x_t, x_{t+1}) \in \mathbf{B}_1$ (the same arguments apply to \mathbf{B}_0 symmetrically). For any round t , let H_t the event that $(x_t, x_{t+1}) \in \mathbf{B}$ and (a) of Lemma 9 holds. Let $t_{\max} = t_0 + (\sqrt{c}/c_4) \cdot \log^{3/2} n$, and let X be the number of rounds between t_0 and t_{\max} for which H_t does not happen. Each time (a) in Lemma 9 doesn't hold, (b) of Lemma 9 holds so there is a constant probability to leave \mathbf{B} , so

$$\mathbb{P}(\text{for every } t \text{ such that } t_0 \leq t \leq t_{\max}, (x_t, x_{t+1}) \in \mathbf{B} \mid X = \mathbf{x}) \leq (1 - c_5)^{\mathbf{x}}. \quad (29)$$

Note that

$$(1 - c_5)^{(t_{\max}-t_0)/4} = \exp\left(\log(1 - c_5) \cdot \frac{\sqrt{c}}{4c_4} \cdot \log^{3/2} n\right).$$

This, together with Eq. (29), implies that either (i) $X < (t_{\max} - t_0)/4$, or w.h.p. (ii) there is a time $t_0 \leq t \leq t_{\max}$ such that $(x_t, x_{t+1}) \notin \mathbf{B}$ (in which case Lemma 10 holds).

Now, consider case (i). Let $u_t = x_t - 1/2$. By Lemma 17, $u_{t_0} > 1/\sqrt{n}$ with constant probability. Therefore, for the price of waiting up to an additional $\log n$ rounds, we can assume that $u_{t_0} > 1/\sqrt{n}$ w.h.p. In what follows, we condition on that event.

Note now, that whenever $(x_t, x_{t+1}) \in \mathbf{B}_1$, by definition $x_{t+1} \geq x_t > 1/2$ and so (x_{t+1}, x_{t+2}) cannot be in \mathbf{B}_0 . This implies that the system must remain in \mathbf{B}_1 until it leaves \mathbf{B} . Also by definition of \mathbf{B}_1 , if $(x_t, x_{t+1}) \in \mathbf{B}_1$ then $x_t \leq x_{t+1}$, and so

$$u_t \leq u_{t+1}. \quad (30)$$

Moreover, by the fact that we are in case (i), we have that the number of rounds $t_0 \leq t \leq t_{\max}$ such that H_t happens is at least

$$k := \frac{3(t_{\max} - t_0)}{4} = \frac{3\sqrt{c}}{4c_4} \cdot \log^{3/2} n.$$

Note that at each such round, by definition of H_t and (a) in Lemma 9,

$$u_{t+1} > u_t \cdot \left(1 + \frac{c_4}{\sqrt{\ell}}\right).$$

Hence, by Eq. (30),

$$\begin{aligned}
u_{t_{\max}} &> u_{t_0} \cdot \left(1 + \frac{c_4}{\sqrt{\ell}}\right)^k = u_{t_0} \cdot \exp\left(k \log\left(1 + \frac{c_4}{\sqrt{\ell}}\right)\right) \\
&> u_{t_0} \cdot \exp\left(\frac{k \cdot c_4}{2\sqrt{\ell}}\right) && \text{(for } n \text{ large enough)} \\
&= u_{t_0} \cdot \exp\left(\frac{3}{8} \log n\right) && \text{(by definition of } k \text{ and } \ell) \\
&= \frac{1}{2} \cdot u_{t_0} \cdot n^{4/3} \\
&> \frac{1}{2} n^{1/4}. && \text{(since we assumed } u_{t_0} > 1/\sqrt{n})
\end{aligned}$$

When n is large, this quantity is larger than 1, hence (i) is impossible unless $(x_t, x_{t+1}) \notin \mathbf{B}$. This concludes the proof of Lemma 10. \square

C.5.3 Area C: Proof of Lemma 11

Without loss of generality, we assume that $(x_t, x_{t+1}) \in \mathbf{C}_1$ (the same arguments apply to \mathbf{C}_0 symmetrically). By Observation 1, we have

$$\mathbb{E}(x_{t+2}) = \mathbb{P}(B_\ell(x_{t+1}) > B_\ell(x_t)) + x_{t+1} \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) - \frac{1}{n}.$$

By Lemma 14 (taking $\lambda > 2$), this becomes

$$\mathbb{E}(x_{t+2}) > \frac{1}{2} + 2 \cdot (x_{t+1} - x_t) - \left(\frac{1}{2} - x_{t+1}\right) \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)). \quad (31)$$

Case 1. If $(x_{t+1} - x_t) > 1/2 - x_{t+1}$, then Eq. (31) implies

$$\mathbb{E}(x_{t+2}) > \frac{1}{2} + 2 \cdot (x_{t+1} - x_t) - \left(\frac{1}{2} - x_{t+1}\right) > \frac{1}{2} + (x_{t+1} - x_t) > \frac{1}{2},$$

so with constant probability $x_{t+2} > 1/2$ and thus $(x_{t+1}, x_{t+2}) \in \mathbf{A}_1$ or is not in \mathbf{YELLOW}' .

Case 2. Else, if $(x_{t+1} - x_t) \leq 1/2 - x_{t+1}$, Eq. (31) rewrites

$$\begin{aligned}
\mathbb{E}(x_{t+2}) &> \frac{1}{2} \left(\frac{1}{2} + x_{t+1}\right) + \frac{1}{2} \left(\frac{1}{2} + 4 \cdot (x_{t+1} - x_t) - 2 \left(\frac{1}{2} - x_{t+1}\right) \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) - x_{t+1}\right) \\
&= \frac{1}{2} \left(\frac{1}{2} + x_{t+1}\right) + \frac{1}{2} \left(4 \cdot (x_{t+1} - x_t) + \left(\frac{1}{2} - x_{t+1}\right) (1 - 2 \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)))\right)
\end{aligned}$$

Since $(x_t, x_{t+1}) \in \mathbf{C}_1$, we have $x_{t+1} \geq x_t$ and $1/2 > x_{t+1}$. Moreover, for ℓ large enough, $1 - 2 \cdot \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t)) > 0$. Hence,

$$\mathbb{E}(x_{t+2}) > \frac{1}{2} \left(\frac{1}{2} + x_{t+1}\right),$$

so, by Lemma 16, with constant probability $x_{t+2} > (1/2 + x_{t+1})/2$, i.e., $x_{t+2} - x_{t+1} > 1/2 - x_{t+1}$. If so, Case 1 applies and with constant probability, $(x_{t+2}, x_{t+3}) \in \mathbf{A}_1$ or is not in \mathbf{YELLOW}' . This concludes the proof of Lemma 11. \square