Distributed Computing
https://doi.org/10.1007/s00446-024-00462-8

Early adapting to trends: self-stabilizing information spread using passive communication

Amos Korman[1,2] · Robin Vacus[3]

Received: 14 November 2022 / Accepted: 17 January 2024

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 $\mathcal{PULL}$ 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 [16, 18, 21, 32, 34], including in multiple natural systems [28, 38, 39]. This task becomes particularly challenging when the system is prone to faults [15, 25, 29, 30], or when the agents or their interactions are constrained [5, 14]. Among others, these issues find relevance in insect populations [38], chemical reaction networks [17], and mobile sensor networks [41]. In particular, in many biological systems, the internal computational abilities of individuals are impressively diverse, whereas the communication capacity is highly limited [10, 28, 38]. An extreme situation, often referred to as *passive communication* [40], is when information is gained by observing the behavior of other animals, which, in some cases, may not even intend to communicate [19, 33]. Such public information can reflect on the quality of possible behavioral options, hence allowing to improve fitness when used properly [20].

Consider, for example, the following scenario that serves as an inspiration for our model. A group of n animals is scat-

✉ Amos Korman
akorman@ds.haifa.ac.il

Robin Vacus
robin.vacus@irif.fr

[1] Department of Computer Science, University of Haifa, Haifa, Israel

[2] The French National Center for Scientific Research (CNRS), UMI FILOFOCS, Tel-Aviv, Israel

[3] The French National Centre for Scientific Research (CNRS), The Institute for Research in Fundamental Computer Science (IRIF), Paris, France

Published online: 22 February 2024

tered 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. Knowledgeable animals have a clear incentive to remain on the preferable side, promoting this choice passively. On the other hand, uninformed animals may choose a side solely for communication purposes, regardless of which they think is best. However, such communication has to be limited in time, since all animals need to eventually converge towards the desirable side.

1.2 Informal description of the problem

This paper studies the *self-stabilizing bit-dissemination* problem, introduced by Boczkowski, Korman, and Natale in [14], 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 $\{0, 1\}$ (in the motivating example above, the opinion corresponds to being either on the eastern side of an area, or the western side). 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 $\in \{0, 1\}$ (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. The goal of the remaining agents, i.e., the non-source agents, is to converge on the opinion held by the source. Because the source does not change its opinion, it may be thought of as representing the environment and thus, not participating in the protocol. Instead, the protocol is executed by the non-source agents.

We study the basic $\mathcal{PULL}$ model of communication [15, 21, 34], in which in each round, each agent sees the information held by ℓ other agents, chosen uniformly at random, where ℓ is small compared to n. Although the sampling is done by non-source agents, the ℓ random agents are chosen from the whole population of agents (including the source). Importantly, however, agents are unable to distinguish the case of sampling the source from the case of sampling a non-source agent. Specifically, we consider the *passive communication* model which assumes that the only information that can be obtained by sampling an agent is its opinion bit.

The *convergence time* of the protocol corresponds to the first round that the configuration of opinions reached a consensus on the correct opinion, and remained unchanged forever after.

We consider the *self-stabilization* framework [4, 24] to model the lack of global organization inherent to biological systems. We assume that the source has pertinent knowledge about which opinion is correct. Conversely, all other agents cannot be sure of their opinion, hence, we assume that their opinion may be corrupt, and even set by an adversary. In the self-stabilizing setting, the goal of the non-source agents is to quickly converge on the correct opinion, despite having an arbitrary initial configuration, set by an adversary.

Our framework and proofs could be extended to allow for a constant number of sources, however, in this case it must be guaranteed that all source agents agree on which opinion is the correct one. Indeed, as mentioned in Sect. 1.4, 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.

1.3 Previous works

The self-stabilizing bit-dissemination problem was introduced in [14], 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 $T = 4 \log n$. In the first half 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 half, 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 half, 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 half 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-syncronization* 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 [14], as well as the focus of the subsequent work [11], was solving the self-stabilizing clock-synchronization using as few as possible bits per message. In fact, the authors in [11] 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 [14], then allowed to further compress the 2 bits into 1-bit messages.

Importantly, however, the protocols in [11, 14] do not fit the framework of passive communication. Indeed, in the corresponding protocols the message revealed by an agent is not necessarily the same as its opinion bit, which is kept in the protocols of [11, 14] as an internal variable. 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.

1.4 On the difficulties resulting from using passive communication

At first glance, to adhere to the passive communication model, one may suggest that agents simply choose their opinion to be the 1-bit messages used in [11], just for the purpose of communication, until a consensus is reached, and then switch their opinions 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 [11], 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.

To further illustrate the difficulty of self-stabilizing information spread under passive communication, let us consider a more general problem called *majority bit-dissemination*. As explained in [14], this problem could be solved efficiently when separating the messages from the opinions but, as shown below, could not be solved efficiently under passive communication. 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 time. The authors of [14] showed that the self-stabilizing majority bit-dissemination problem could be solved in logarithmic time, using messages of size 3 bits, and the authors of [11] showed how to reduce the message size to 1. As mentioned, the messages in these protocols do not necessarily coincide with the opinions, which were stored as internal variables, and hence the protocols in [11, 14] 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 efficiently 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 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.

1.5 Problem definition

Let $\mathcal{Y} = \{0, 1\}$ be the *opinion space*. The population contains *n agents*, for some $n \in \mathbb{N}$, with one specific agent, called *source*, and $n - 1$ *non-source agents*. The source agent holds an opinion $z \in \mathcal{Y}$ which is called the *correct opinion*. Informally, the source agent may be observed by other agents but does not actively participate in the protocol. In particular, its opinion, which is assumed to be set by an adversary, remains fixed throughout the execution. Instead, the *protocol* is run by the set $I = \{1, \dots, n-1\}$ of non-source agents.

Execution proceeds in discrete, synchronous rounds $t = 0, 1, \dots$. Each non-source agent is seen as a state machine over $\mathcal{Y} \times \Sigma$, where Σ depends on the protocol. We write $\gamma_t^{(i)} = (Y_t^{(i)}, \sigma_t^{(i)}) \in \mathcal{Y} \times \Sigma$ to denote the *state* of Agent i in round t. Specifically, we refer to $Y_t^{(i)} \in \mathcal{Y}$ as the *opinion* of Agent i in round t, and to $\sigma_t^{(i)}$ as its *internal memory state*. A *configuration* $C \in \mathcal{Y} \times (\mathcal{Y} \times \Sigma)^I$ of the system consists of the correct opinion, together with the state of every non-source agent. We write C_t to denote the configuration of the system in round t.

Non-source agents do not "communicate" in the classical sense. Instead, in every round t, they consecutively perform the two following operations. Note that both operations occur within a single round. Hence, since we are interested in the measure of the number of rounds, we assume that these operations occur instantaneously.

1. **Sampling.** Every agent i receives an *opinion sample* $S_t^{(i)} \in \mathcal{Y}^\ell$, where ℓ is a parameter called *sample size*. Every element in $S_t^{(i)}$ is equal to the correct opinion z with probability $1/n$, and is equal to the opinion of a non-source agent chosen uniformly at random (u.a.r) in I otherwise. This is equivalent to assuming that in each sample the opinion of an agent is sampled uniformly at random, where we consider all agents, including the source, as having equal probability of being sampled. We assume that all elements in $S_t^{(i)}$ are drawn independently and that the opinion samples $\left\{S_t^{(i)}\right\}_{i \in I}$ are also independent of each other.
2. **Computation.** Every agent i adopts a new state $\gamma_{t+1}^{(i)}$, based on its previous state $\gamma_t^{(i)}$ and the opinion sample $S_t^{(i)}$, according to a *transition function*

   $$f : (\mathcal{Y} \times \Sigma) \times \mathcal{Y}^\ell \mapsto (\mathcal{Y} \times \Sigma),$$

   which depends on the protocol. The transition function might be randomized.

An *execution* of a protocol on an initial configuration C_0 is a random sequence of configurations $\{C_t\}_{t \in \mathbb{N}}$, where every C_t is obtained from C_{t-1} by applying the two aforementioned steps simultaneously on all (non-source) agents in parallel.

Using the terminology from [24], an execution is considered *legal* if, for every round t and every non-source agent i, the opinion of Agent i in round t is equal to the correct opinion, i.e., $Y_t^{(i)} = z$. A configuration C is *safe* if every execution starting from C is legal. We say that an event happens *with high probability* (w.h.p.) if it happens with probability at least $1 - 1/n^2$ as long as n is large enough. A protocol is said to *solve the self-stabilizing bit-dissemination problem in time T* if, for any initial configuration C_0, the T'th configuration in an execution, namely C_T, is safe w.h.p., where the probability is taken over the possible executions.

Important comments. Although our model is now already well-defined, let us make a few remarks to clarify it further.

- The indices are used for analysis purposes, and it should be clear that a non-source agent is neither aware of its index, nor of the round number. Moreover, upon receiving an opinion sample S_t, it is not aware of where the opinions in the sample come from (and whether they are observing the correct opinion directly).
- Note that we do not require agents to irrevocably commit to their final opinion, but rather that they eventually converge on the correct opinion without necessarily being aware that convergence has happened.
- As mentioned, we consider the *source* as an agent that holds the correct opinion z throughout the execution. Despite the word "agent", the source does not have an internal memory state, and does not run the protocol. This is justified by the fact that in our interpretation, the source represents an informed individual, unwilling to cooperate, and incentivized to always keep the correct opinion. Alternatively, the source can be thought of as representing the environment, and hence, sampling the source is equivalent to sampling the environment. Importantly, as made clear in our problem definition, self-stabilization is defined only with respect to the non-source agents. In other words, at the beginning of an execution, an adversary may choose the opinion $Y_0^{(i)}$ and internal memory state $\sigma_0^{(i)}$ of every non-source agent, as well as the correct opinion z; Hence, since in our protocols the internal memory state of non-source agents does not include the possibility of indicating that the agent is a source, the adversary cannot "trick" non-source agents into believing that they are the source. The source itself should be seen as a part of the environment that the non-source agents are facing.
- One may consider a system where non-source agents suffer from transient faults that corrupt their internal states. As is common in self-stabilization contexts, we think of the initial configuration C_0 as the last configuration for

which a transient fault has occurred. In this respect, the convergence time corresponds to the time until the system fully recovers from the last transient fault.

- Our model is closely related to the *message-passing* setting, where the communication topology would be the complete graph over all agents (including the source). However, agents do not choose whether and where to send messages. Instead, following the PULL model of communication, (non-source) agents observe ℓ neighbors chosen uniformly at random. Note that, following the *passive communication* assumption, the only information that is revealed upon observing an agent is its opinion.
- Since all opinion samples are random by definition (and not chosen by an adversary), there is no need for any fairness assumption. Instead, our results will typically hold w.h.p.

Additional notations. For any $k \in \mathbb{N}$ and $p \in [0,1]$, we write $\mathcal{B}(p)$ (resp. $\mathcal{B}_k(p)$) to denote the Bernoulli (resp. Binomial) distribution with parameter p (resp. (k, p)). Moreover, we write $\mathbb{P}(B_k(p) > B_k(q))$ to represent $\mathbb{P}(X > Y)$ where $X \sim \mathcal{B}_k(p)$, $Y \sim \mathcal{B}_k(q)$, and X and Y are independent. Finally, we write

$$x_t := \frac{z + \#\{i \in I, Y_t^{(i)} = 1\}}{n}$$

to denote the proportion of agents with opinion 1 in round t among all agents (including the source agent).

1.6 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, the protocol uses $\ell = c \log n$ samples, for a sufficiently large constant c. The internal state space is $\Sigma = \{0, \dots, \ell\}$. For any opinion sample $A \in \mathcal{Y}^\ell$, let COUNT(A) denote the number of 1-opinions in A. The transition function is as follows:

```
Input : $Y_t \in \mathcal{Y}$, $\sigma_t \in \Sigma = \{0, \dots, \ell\}$, $S_t \in \mathcal{Y}^\ell$
$\sigma_{t+1} \leftarrow$ COUNT($S_t$) ;
if $\sigma_{t+1} > \sigma_t$ then $Y_{t+1} \leftarrow 1$ else if $\sigma_{t+1} < \sigma_t$ then $Y_{t+1} \leftarrow 0$
    else $Y_{t+1} \leftarrow Y_t$ Output : $(Y_{t+1}, \sigma_{t+1})$
```

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 σ_t is used to compute both Y_t and Y_{t+1}. For example, if the sample S_{t-1} at round $t-1$ happens to contain more 1's, then σ_t 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 divide the sample of round t into 2 samples of equal size. One sample will be used to compare with one sample of round $t-1$, and the other sample will be used to compare with one sample of round $t+1$. Note that this implies that the sample size is twice as big, 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. Its transition function is specified below (Protocol 1).

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

```
Input : $Y_t \in \mathcal{Y}$, $\sigma_t \in \Sigma = \{0, \dots, \ell\}$, $S_t \in \mathcal{Y}^{2\ell}$
Divide $S_t$ into two vectors $S_t', S_t'' \in \mathcal{Y}^\ell$ of equal size ;
$\sigma_{t+1} \leftarrow$ COUNT($S_t'$) ;   $\mathsf{tmp}_{t+1} \leftarrow$ COUNT($S_t''$) ;
if $tmp_{t+1} > \sigma_t$ then $Y_{t+1} \leftarrow 1$ else if $tmp_{t+1} < \sigma_t$ then
    $Y_{t+1} \leftarrow 0$ else $Y_{t+1} \leftarrow Y_t$ Output : $(Y_{t+1}, \sigma_{t+1})$
```

Note that, although we used time indices for clarity, the protocol does not require the agents to know t. The following theorem consists of the main result in the paper. Its proof is deferred to Sects. 2, 3, 5 and 6.

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 $\ell = \Theta(\log n)$ samples in each round, and using $\Theta(\log \ell)$ bits of memory per agent.*

The task of discovering and analyzing an algorithm using less than a logarithmic number of samples (e.g., a constant) appears challenging and is, therefore, left for future work. We note, however, that on a practical level, we believe that demonstrating the existence of algorithms utilizing logarithmic and constant sample sizes offers similar insights into the capacity of biological systems to efficiently spread information reliably.

Provided that $\ell = \Theta(\log n)$, it is worth noting that all algorithms require $\Omega(\log n / \log\log n)$ rounds to solve the bit-dissemination problem (even with active communications). This is because this is the time needed for merely spreading information from the source to the whole population. More precisely, with fewer rounds, the configurations in which the source has opinion 0 and 1 are perfectly indistinguishable for a non-empty subset of the agents. Therefore, although our upper bound on the convergence time of Protocol 1 may not be tight, one can only hope to decrease it by approximately a factor of $\log^{3/2} n$.

Finally, we consider the more general case with k opinions, for an arbitrary $k \in \mathbb{N}$. Importantly, however, we restrict attention to the case that the agents agree on a labeling of the opinions. That is, we assume that the set of opinions is $\mathcal{Y} = \{0, \ldots, k-1\}$. The case in which there may be conflicts in the way agents view the labeling of opinions remains for future work, see Sect. 8. The following theorem is proved in Sect. 7.

Theorem 2 *Let k be a positive integer and let $m = \lceil \log_2 k \rceil$. When $\mathcal{Y} = \{0, \ldots, k-1\}$, there exists a protocol that solves the bit-dissemination problem in $O(\log^{5/2} n)$ rounds with high probability, while relying on $\ell = \Theta(\log n)$ samples in each round and using $\Theta(m \log \ell)$ bits of memory.*

Our analysis involves partitioning the configuration space and subsequently studying the dynamics' behavior in each subset, using classical concentration and anti-concentration results. This approach required us to identify a proper partitioning for which the dynamics are analytically tractable both on each part of the partitioning and in between parts. While we find this approach interesting and challenging at times, it is not novel [23]. Moreover, it is anticipated that, in a specific setting, the method may necessitate customization, presenting challenges for its reuse in other related problems.

1.7 Other related works

In recent years, the study of *population protocols* has attracted significant attention in the distributed computing community [1–3, 6, 8]. These models often consider agents that interact under random meeting patterns while being restricted in both their memory and communication capacities. While these model are inspired by biological scenarios, in many cases, the algorithms used appear unlikely to be employed by biological ensembles. By now, we understand the computational power of such systems rather well, but apart from a few exceptions [7], 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 [12]. 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 [36]. In theoretical computer science, extensive research has been devoted to analyzing the time to reach consensus, following different updating rules including the *3-majority* [23], *Undecided-State Dynamics* [6], 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 contrast, in many natural settings 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 [9, 35, 37, 39]. 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., [16, 18, 21, 31, 32, 34]. A classical algorithm in the $\mathcal{PULL}$ model spreads the opinion of the source to all others in $\Theta(\log n)$ rounds w.h.p., by letting each uninformed agent copy the opinion of an informed agent whenever seeing one for the first time, as observed in, e.g., [34]. 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 [13, 27]. Another model of faults is *self-stabilization* [22], where the system must eventually converge on the opinion of the source regardless of the initial configuration of states [22]. For example, the algorithm in [34] 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. For an introduction to self-stabilizing algorithms, see, e.g., [4, 24], and see [26] for another work solving self-stabilizing problems using weak communications.

Finally, it is worth noting that the term "bio-inspired", often used in the literature, typically refers to research focused on applications in artificial intelligence or swarm robotics. In contrast, our aim is to employ algorithmic tools to comprehend the behavior of biological ensembles. Consequently, many articles labeled as bio-inspired may be irrelevant to our context.

2 Proof of Theorem 1: an overview

The goal of this section is to prove Theorem 1. The $O(\log \ell)$ bits upper bound on the memory complexity clearly follows from the fact that the internal state space $\Sigma = \{0, \ldots, \ell\}$ of the FET algorithm (Protocol 1) is only of size $\ell + 1$, as required to memorize the number of 1's in a sample.

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. The closure property is easy to verify. Indeed, if there exists a round t such that the system has reached consensus in round t, then all future samples will be equal to the correct opinion, and by construction of Protocol 1, this implies that every agent will keep the correct opinion.

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 $O(T)$ 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 $(1/n^{\epsilon})^{2/\epsilon} = 1/n^2$.

Recall that x_t denotes the fraction of agents with opinion 1 at round t among the whole population (including the source). We shall extensively use the two dimensional grid $\mathcal{G} := \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, 1\}^2$. 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 t, conditioning on $(x_t, x_{t+1}) = (\mathbf{x_t}, \mathbf{x_{t+1}})$, a non-source agent i has opinion 1 on round $t+2$ w.p.*

$$\mathbb{P}\left(Y^{(i)}_{t+2} = 1 \,\middle|\, \begin{array}{c} x_t = \mathbf{x_t} \\ x_{t+1} = \mathbf{x_{t+1}} \\ Y^{(i)}_{t+1} = \mathbf{Y^{(i)}_{t+1}} \end{array} \right) = \mathbb{P}\left(B_\ell\left(\mathbf{x_{t+1}}\right) > B_\ell\left(\mathbf{x_t}\right)\right) + \mathbb{1}_{\{\mathbf{Y^{(i)}_{t+1}} = 1\}} \cdot \mathbb{P}\left(B_\ell\left(\mathbf{x_{t+1}}\right) = B_\ell\left(\mathbf{x_t}\right)\right). \qquad (1)$$

Moreover, there are independent binary random variables[1] *$X_1, \ldots, X_n$ such that x_{t+2} is distributed as $\frac{1}{n}\sum X_i$. Eventually,*

$$\mathbb{E}\left(x_{t+2} \,\middle|\, \begin{array}{c} x_t = \mathbf{x_t} \\ x_{t+1} = \mathbf{x_{t+1}} \end{array} \right) = \mathbb{P}\left(B_\ell\left(\mathbf{x_{t+1}}\right) > B_\ell\left(\mathbf{x_t}\right)\right) + \mathbf{x_{t+1}} \cdot \mathbb{P}\left(B_\ell\left(\mathbf{x_{t+1}}\right) = B_\ell\left(\mathbf{x_t}\right)\right) + \frac{1}{n}\left(1 - \mathbb{P}\left(B_\ell\left(\mathbf{x_{t+1}}\right) \geq B_\ell\left(\mathbf{x_t}\right)\right)\right). \qquad (2)$$

The proof of Observation 1 is deferred to Sect. 4. 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 $\delta > 0$ (δ should be though of as a very small constant) and $\lambda_n = \frac{1}{\log^{1/2+\delta} n}$. We partition $\mathcal{G}$ into domains as follows (see illustration on Fig. 1a).

$$\text{GREEN}_1 = \left\{ (x_t, x_{t+1}) \,\middle|\, x_{t+1} \geq x_t + \delta \right\},$$

$$\text{PURPLE}_1 = \left\{ (x_t, x_{t+1}) \,\middle|\, \begin{array}{l} 1/\log n \leq x_t < 1/2 - 3\delta, \text{ and} \\ (1 - \lambda_n) \cdot x_t \leq x_{t+1} < x_t + \delta \end{array} \right\},$$

$$\text{RED}_1 = \left\{ (x_t, x_{t+1}) \,\middle|\, \begin{array}{l} 1/\log n \leq x_{t+1}, \text{ and} \\ x_t < 1/2 - 3\delta, \text{ and} \\ x_t - \delta \leq x_{t+1} < (1 - \lambda_n) \cdot x_t \end{array} \right\},$$

$$\text{CYAN}_1 = \left\{ (x_t, x_{t+1}) \,\middle|\, \begin{array}{l} 0 \leq \min(x_t, x_{t+1}) < 1/\log n, \text{ and} \\ x_t - \delta < x_{t+1} < x_t + \delta \end{array} \right\},$$

$$\text{YELLOW} = \left\{ (x_t, x_{t+1}) \,\middle|\, \begin{array}{l} 1/2 - 3\delta \leq x_t < 1/2 + 3\delta, \text{ and} \\ 1/2 - 4\delta \leq x_{t+1} \leq 1/2 + 4\delta, \text{ and} \\ |x_{t+1} - x_t| < \delta \end{array} \right\}.$$

Similarly, for the former 4 domains, we define GREEN_0, PURPLE_0, 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.

[1] In fact, in general, $Y^{(i)}_{t+2}$, $i \in I$ are *not* independent conditioned on (x_t, x_{t+1}). This is not a problem since we mainly care about their sum, but it forces us to introduce variables $X_1, \ldots, X_n$ in order to use classical concentration results on their sum.

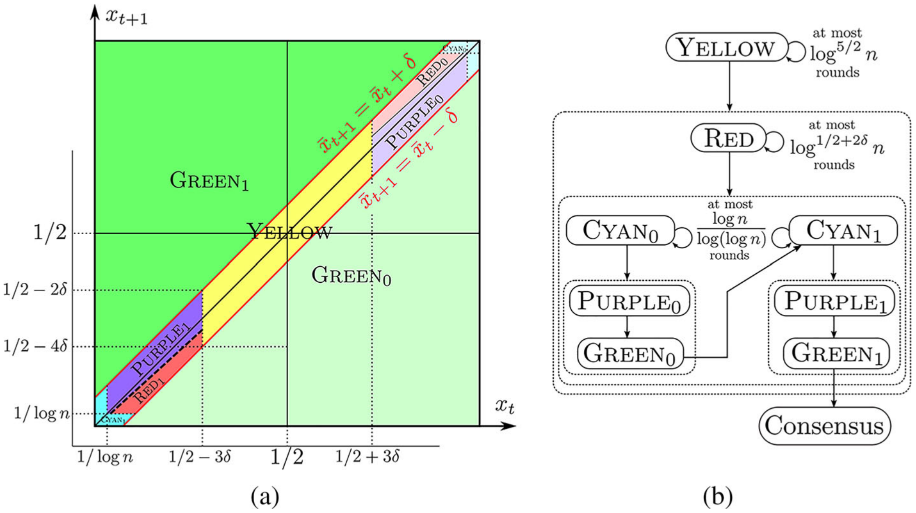

Fig. 1 **a** Partitioning the state space into domains. The x-axis (resp., y-axis) represents the proportion of agents with opinion 1 in round t (resp., $t+1$). The thick dashed line at the frontier between PURPLE$_1$ and RED$_1$ is defined by $x_{t+1} = (1 - \lambda_n)x_t$. **b** Sketch of the proof of Theorem 1. All transitions w.p. at least $1 - 1/n^{\Omega(1)}$. The process stays in a domain for as many rounds as indicated on the corresponding self-loop w.p. at least $1 - 1/n^{\Omega(1)}$, and at most a constant number of rounds when no self-loop is represented. The source is assumed to have opinion 1

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 process' "derivative" at time t.)

2.2 Dynamics at different domains

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 will be on 1, and we are done. If, on the other hand, the Markov chain is in GREEN$_0$, then the consensus of non-source agents would be on 0. As we show later, in that case the Markov chain would reach CYAN$_1$ in one round w.h.p.

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

The proof of Lemma 1 follows from a simple application of Hœffding's inequality, and is deferred to Sect. 6.1. 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 Sect. 6.2.

Lemma 2 *(Purple area) Assume that c is sufficiently large. If $(x_t, x_{t+1}) \in$ PURPLE$_1$, then $(x_{t+1}, x_{t+2}) \in$ GREEN$_1$ w.h.p. Similarly, if $(x_t, x_{t+1}) \in$ PURPLE$_0$, then $(x_{t+1}, x_{t+2}) \in$ 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 RED$_1$ (resp., RED$_0$), 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 Sect. 6.3.

Lemma 3 *(Red area) Consider the case that $(x_{t_0}, x_{t_0+1}) \in$ 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 $(x_{t_1}, x_{t_1+1}) \notin$ YELLOW ∪ RED.*

Next, we bound the time that can be spent in CYAN$_1$. (A similar result holds for CYAN$_0$.) 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 GREEN$_1$ ∪ 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 GREEN$_1$ ∪ PURPLE$_1$. The proof of the following lemma is given in Sect. 6.4.

Lemma 4 *(Cyan area) Consider the case that $(x_{t_0}, x_{t_0+1}) \in \textsc{Cyan}_1$ for some round t_0, and let $t_1 = \min\{t \geq t_0, (x_t, x_{t+1}) \notin \textsc{Cyan}_1\}$. Then with probability at least $1-1/n^{\Omega(1)}$ we have (1) $t_1 < t_0 + O(\log(n)/\log(\log n))$, and (2) $(x_{t_1}, x_{t_1+1}) \in \textsc{Green}_1 \cup \textsc{Purple}_1$. Moreover, the same holds symmetrically for $\textsc{Cyan}_0$.*

Eventually, we consider the central area, namely, $\textsc{Yellow}$, 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 Sect. 5.

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

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

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, and that we already checked the closure property. The reader is strongly encouraged to refer to Fig. 1b to follow the ensuing arguments more easily.

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

  $$\begin{cases} t_1 < O(\log^{5/2} n) \text{ w.h.p., and} \\ (x_{t_1}, x_{t_1+1}) \in \textsc{Red} \cup \textsc{Cyan} \cup \textsc{Purple} \cup \textsc{Green}. \end{cases} \tag{3}$$

  Else, $(x_0, x_1) \notin \textsc{Yellow}$ so $t_1 = 0$, and Eq. (3) also holds.
- Let $t_2 = \min\{t \geq t_1, (x_t, x_{t+1}) \notin \textsc{Red}\}$. If $(x_{t_1}, x_{t_1+1}) \in \textsc{Red}$, we apply Lemma 3 to get that

  $$\begin{cases} t_2 < t_1 + \log^{1/2+2\delta} n \text{ w.h.p., and} \\ (x_{t_2}, x_{t_2+1}) \in \textsc{Cyan} \cup \textsc{Purple} \cup \textsc{Green}. \end{cases} \tag{4}$$

  Else, $(x_{t_1}, x_{t_1+1}) \notin \textsc{Red}$ so $t_1 = t_2$, and by Eq. (3), it implies that Eq. (4) also holds.
- Let $t_3 = \min\{t \geq t_2, (x_t, x_{t+1}) \notin \textsc{Cyan}\}$. If $(x_{t_2}, x_{t_2+1}) \in \textsc{Cyan}$, we apply Lemma 4 to get that

  $$\begin{cases} t_3 < t_2 + \log(n)/\log(\log n), \quad \text{and} \\ (x_{t_3}, x_{t_3+1}) \in \textsc{Purple} \cup \textsc{Green} \text{ w.p. } \geq 1 - 1/n^{\Omega(1)}. \end{cases} \tag{5}$$

  Else, $(x_{t_2}, x_{t_2+1}) \notin \textsc{Cyan}$ so $t_2 = t_3$, and by Eq. (4), it implies that Eq. (5) also holds.
- Let $t_4 = \min\{t \geq t_3, (x_t, x_{t+1}) \in \textsc{Green}\}$. By Lemma 2, and by Eq. (5), we have that $t_4 = t_3$ or $t_4 = t_3 + 1$ w.h.p. If $(x_{t_4}, x_{t_4+1}) \in \textsc{Green}_1$, then by Lemma 1 the consensus is reached on round t_4+1. Otherwise, if $(x_{t_4}, x_{t_4+1}) \in \textsc{Green}_0$, 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 $(x_{t_4+1}, x_{t_4+2}) \in \textsc{Green}_0$ or $(x_{t_4+1}, x_{t_4+2}) \in \textsc{Cyan}_1$ (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 $(x_{t_4+2}, x_{t_4+3}) = (1/n, 1/n) \in \textsc{Cyan}_1$. 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 $\textsc{Purple}_1 \cup \textsc{Green}_1$, then to $\textsc{Green}_1$, 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 $1 - 1/n^c$. In other words, there exists $T = O(\log^{5/2} n)$ such that w.h.p., the configuration C_T of the system in round T is *safe*, which concludes the proof of Theorem 1. □

3 Probabilistic tools–competition between coins

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

3.1 Lower bounds on the probability that the best coin wins

In Lemmas 6 and 7 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 6 is particularly effective when the difference between p and q is sufficiently large. Its proof is based on a simple application of Hœffding's inequality.

Lemma 6 *For every $p, q \in [0, 1]$ s.t. $p < q$ and every integer k, we have*

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

Proof Let Y_i, $i \in \{1, \ldots, k\}$ 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

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

Since each Y_i is bounded, and $\mathbb{E}(Y_i) = (p-q)$, we can apply Hœffding's inequality (Theorem 5) to get

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

□

Lemma 6 is not particularly effective when p and q are close to each other. For such cases, we shall use the following lemma.

Lemma 7 *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 $k > K$,*

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

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

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

so

$$\mathbb{P}\left(B_k(p) < B_k(q)\right) - \mathbb{P}\left(B_k(q) < B_k(p)\right)$$
$$= \sum_{d=1}^{k} \mathbb{P}\left(|B_k(q) - B_k(p)| = d\right) \cdot \frac{\mathbb{P}\left(B_k(p) = B_k(q) - d\right) - \mathbb{P}\left(B_k(q) = B_k(p) - d\right)}{\mathbb{P}\left(B_k(p) = B_k(q) - d\right) + \mathbb{P}\left(B_k(q) = B_k(p) - d\right)}. \quad (6)$$

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

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

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. (6) as

$$\mathbb{P}\left(B_k(p) < B_k(q)\right) - \mathbb{P}\left(B_k(q) < B_k(p)\right)$$
$$= \sum_{d=1}^{k} \mathbb{P}\left(|B_k(q) - B_k(p)| = d\right) \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d}. \quad (7)$$

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 1 *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. □

Claim 2 *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:

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

Note that

$$\begin{aligned}&\lim_{p,q \to 1/2} \sum_{i=0}^{d-1} (q(1-p))^{d-1-i} \cdot (p(1-q))^i\\ &\quad = \sum_{i=0}^{d-1} \left(\frac{1}{4}\right)^{d-1-i} \cdot \left(\frac{1}{4}\right)^i = \sum_{i=0}^{d-1} \left(\frac{1}{2}\right)^{2d-2} = d \cdot \left(\frac{1}{2}\right)^{2d-2},\end{aligned}$$

and that

$$\begin{aligned}&\lim_{p,q \to 1/2} (q(1-p))^d + (p(1-q))^d\\ &\quad = \left(\frac{1}{4}\right)^d + \left(\frac{1}{4}\right)^d = \left(\frac{1}{2}\right)^{2d-1}.\end{aligned}$$

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 2. □

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 2, there exists $\epsilon = \epsilon(\gamma, D) = \epsilon(\lambda)$, s.t. for $p, q \in [1/2 - \epsilon, 1/2 + \epsilon]$,

$$\frac{(q(1-p))^D - (p(1-q))^D}{(q(1-p))^D + (p(1-q))^D} > (q-p) \cdot 2\lambda'. \tag{8}$$

Now we derive a lower bound on Eq. (7):

$$\begin{aligned}&\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p))\\ &= \sum_{d=1}^{k} \mathbb{P}(|B_k(q) - B_k(p)| = d)\\ &\quad \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} && \text{(Eq. (7))}\\ &\geq \sum_{d=D}^{k} \mathbb{P}(|B_k(q) - B_k(p)| = d)\\ &\quad \cdot \frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d}\\ &\geq \sum_{d=D}^{k} \mathbb{P}(|B_k(q) - B_k(p)| = d)\\ &\quad \cdot \frac{(q(1-p))^D - (p(1-q))^D}{(q(1-p))^D + (p(1-q))^D} && \text{(by Claim 1)}\\ &> (q-p)\\ &\quad \cdot 2\lambda' \sum_{d=D}^{k} \mathbb{P}(|B_k(q) - B_k(p)| = d) && \text{(by Eq. (8))}\\ &= (q-p) \cdot 2\lambda'\\ &\quad \cdot (1 - \mathbb{P}(|B_k(q) - B_k(p)| < D)) .\end{aligned}$$

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 $k > K$,

$$\mathbb{P}(B_k(p) < B_k(q)) - \mathbb{P}(B_k(q) < B_k(p)) > (q-p) \cdot 2\lambda. \tag{9}$$

Eventually, we write

$$\begin{aligned}&\mathbb{P}(B_k(p) < B_k(q))\\ &= 1 - \mathbb{P}(B_k(q) < B_k(p)) - \mathbb{P}(B_k(p) = B_k(q))\\ &> 1 - \mathbb{P}(B_k(p) < B_k(q))\\ &\quad + 2\lambda(q-p) - \mathbb{P}(B_k(p) = B_k(q)) . && \text{(by Eq.(9))}\end{aligned}$$

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 7.

3.2 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 8 *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 $Y_i, i \in \{1,\ldots,k\}$ 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 $\mu = \mathbb{E}(Y_1)$, $\sigma^2 = \mathrm{Var}(Y_1)$, and $\rho = \mathbb{E}(|Y_1 - \mu|^3)$. Writing the definitions and simplifying, we obtain

$$\mu = p - q, \quad \sigma^2 = p(1-p) + q(1-q),$$
$$\rho = (2p^3 - 3p^2 + p) - (2q^3 - 3q^2 + q). \tag{10}$$

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 7),

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

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., $C = 0.4748$. □

Claim 3 *We have that $\rho < \sigma^2$.*

Proof Let $f(p) = 2p^3 - p/2$ and $g(p) = 1/4 - p^2$. We start by proving that for every $p \in [-1/2, 1/2]$, $|f(p)| \leq g(p)$. Since f is anti-symmetric, $|f|$ is symmetric, and g is symmetric, we can restrict the analysis to $[0, 1/2]$. On this interval, $|f(p)| = p/2 - 2p^3$, and

$$g(p) - |f(p)| = \frac{1}{4} - \frac{p}{2} - p^2 + 2p^3 = \frac{1}{4}(1-2p)^2(1+2p) \geq 0.$$

We can rewrite Eq. (10) as

$$\sigma^2 = g(p + \tfrac{1}{2}) + g(q + \tfrac{1}{2}) \quad \text{and} \quad \rho = f(p + \tfrac{1}{2}) + f(q + \tfrac{1}{2}).$$

Therefore,

$$|\rho| \leq |f(p + \tfrac{1}{2})| + |f(q + \tfrac{1}{2})| \leq g(p + \tfrac{1}{2}) + g(q + \tfrac{1}{2}) = \sigma^2,$$

which concludes the proof of Claim 3. □

By Claim 3, 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 of Lemma 8. □

Just as Lemma 7 was a version of Lemma 6 optimized for cases where p and q are close to each other, Lemma 9 complements Lemma 8 in such situations.

Lemma 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 $q - p \leq 1/\sqrt{k}$, we have*

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

Proof Recall that (see Eq. (7) in the proof of Lemma 7):

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

The following claim is analogous to Claim 2, but this time we are looking for an upper bound (instead of a lower bound) on the same quantity. □

Claim 4 *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 $d \in \mathbb{N}$,*

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

As in the proof of Claim 2, we have

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

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

$$\begin{aligned}&\frac{(q(1-p))^d - (p(1-q))^d}{(q(1-p))^d + (p(1-q))^d} \le \alpha d \cdot (q-p) \\ &\cdot \frac{(q(1-p))^d}{(q(1-p))^d + (p(1-q))^d} \le \alpha d \cdot (q-p),\end{aligned}$$

which concludes the proof of Claim 4.

Using Claim 4 on Eq. (11), we obtain

$$\begin{aligned}&\mathbb{P}\left(B_k(p) < B_k(q)\right) - \mathbb{P}\left(B_k(q) < B_k(p)\right) \le \alpha \\ &\cdot (q-p) \sum_{d=1}^{k} d \cdot \mathbb{P}\left(|B_k(q) - B_k(p)| = d\right). \end{aligned} \tag{12}$$

Claim 5 *For every $p, q \in [1/3, 2/3]$ with $p < q$, and every integer k,*

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

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

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

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

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

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

$$\begin{aligned}&\mathbb{E}\left(|X^{(1)} - X^{(2)}|\right) = \mathbb{E}\left(\sqrt{\left(X^{(1)} - X^{(2)}\right)^2}\right) \\ &\le \sqrt{\mathbb{E}\left(\left(X^{(1)} - X^{(2)}\right)^2\right)} \quad \text{(Jensen inequality, } x \mapsto \sqrt{x} \\ &\text{being concave)} \\ &= \sqrt{\operatorname{Var}\left(X^{(1)} - X^{(2)}\right)} \quad \left(\text{since } \mathbb{E}\left(X^{(1)} - X^{(2)}\right) = 0\right) \\ &= \sqrt{2kq(1-q)}, \quad (X^{(1)}, X^{(2)} \sim \mathcal{B}_k(q) \text{and are independent}).\end{aligned}$$

which concludes the proof of Claim 5. □

We note that

$$\begin{aligned}&\sum_{d=1}^{k} d \cdot \mathbb{P}\left(|B_k(q) - B_k(p)| = d\right) = \mathbb{E}\left(|B_k(q) - B_k(p)|\right) \\ &\le \sqrt{2kq(1-q)} + k \cdot (q-p) (\text{by Claim 5}) \\ &\le \sqrt{2kq(1-q)} + \sqrt{k} (\text{since } q - p \le 1/\sqrt{k}\ ) \\ &\le 2\sqrt{k}.\end{aligned}$$

Eventually, Eq. (12) becomes

$$\mathbb{P}\left(B_k(p) < B_k(q)\right) - \mathbb{P}\left(B_k(q) < B_k(p)\right) \le 2\alpha \cdot (q-p)\sqrt{k}. \tag{13}$$

To conclude, we write

$$\begin{aligned}\mathbb{P}\left(B_k(p) < B_k(q)\right) &= 1 - \mathbb{P}\left(B_k(q) < B_k(p)\right) - \\ &\quad \mathbb{P}\left(B_k(p) = B_k(q)\right) \\ &< 1 - \mathbb{P}\left(B_k(p) < B_k(q)\right) + 2\alpha \cdot (q-p)\sqrt{k} - \\ &\quad \mathbb{P}\left(B_k(p) = B_k(q)\right). \qquad \text{(by Eq. (13))}\end{aligned}$$

Hence,

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

which concludes the proof of Lemma 9.

4 Preliminaries

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.

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, $\mathsf{tmp}^{(i)}$ and $\sigma^{(i)}$ 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 $({\mathsf{tmp}_{t+2}}^{(i)})_{i \in I}$ and $({\sigma_{t+1}}^{(i)})_{i \in I}$ are mutually independent, thus variables $(Y_{t+2}^{(i)})_{i \in I}$ are mutually independent.

(ii) for every $i \in I$, ${\mathsf{tmp}_{t+2}}^{(i)} \sim \mathcal{B}_\ell(x_{t+1})$, and ${\sigma_{t+1}}^{(i)} \sim \mathcal{B}_\ell(x_t)$, so we can write for every non-source agent $i \in I_{t+1}^1$,

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

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

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

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

We assumed the source 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:

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

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

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

$$\begin{aligned}&\mathbb{P}\left(B_\ell\left(x_{t+1}\right) > B_\ell\left(x_t\right)\right) \\ &\quad \leq \mathbb{P}\left(Y_{t+2}^{(i)} = 1\right) \leq \mathbb{P}\left(B_\ell\left(x_{t+1}\right) \geq B_\ell\left(x_t\right)\right), \quad (14)\end{aligned}$$

and

$$\begin{aligned}&\mathbb{P}\left(B_\ell\left(x_{t+1}\right) > B_\ell\left(x_t\right)\right) - \frac{1}{n} \\ &\quad \leq \mathbb{E}\left(x_{t+2}\right) \leq \mathbb{P}\left(B_\ell\left(x_{t+1}\right) \geq B_\ell\left(x_t\right)\right) + \frac{1}{n}. \quad (15)\end{aligned}$$

Because the source has opinion 1, the left-hand side in Eq. (15) 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:

$$\begin{aligned}\mathbb{P}\left(B_\ell\left(x_{t+1}\right) > B_\ell\left(x_t\right)\right) + x_{t+1}\cdot\mathbb{P}\left(B_\ell\left(x_{t+1}\right) = B_\ell\left(x_t\right)\right) - \frac{1}{n}&\\ < \mathbb{E}\left(x_{t+2}\right) &<\\ \mathbb{P}\left(B_\ell\left(x_{t+1}\right) > B_\ell\left(x_t\right)\right) + x_{t+1}\cdot\mathbb{P}\left(B_\ell\left(x_{t+1}\right) = B_\ell\left(x_t\right)\right) + \frac{1}{n}.&\end{aligned} \tag{16}$$

5 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:

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

Obviously, YELLOW $\subset$ YELLOW$'$, so in order to prove Lemma 5 it suffices to prove Lemma 10 below.

Lemma 10 *Consider that* $(x_{t_0}, x_{t_0+1}) \in$ YELLOW$'$*. Then, w.h.p.,* $\min\{t > t_0 \text{ s.t. } (x_t, x_{t+1}) \notin \text{YELLOW}'\} < t_0 + O(\log^{5/2} n)$.

5.1 Effects of noise

We will need the following result to break ties.

Lemma 11 *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*

$$\mathbb{P}\left(x_{t+2} \le \mathbb{E}(x_{t+2}) - 1/\sqrt{n}\right), \mathbb{P}\left(x_{t+2} \ge \mathbb{E}(x_{t+2}) + 1/\sqrt{n}\right) \ge \beta.$$

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

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

Let $p = \mathbb{P}\left(B_\ell\left(x_{t+1}\right) \ge B_\ell\left(x_t\right)\right)$ and $q = \mathbb{P}\left(B_\ell\left(x_{t+1}\right) > B_\ell\left(x_t\right)\right)$. 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, $p - q = \mathbb{P}(B_\ell(x_{t+1}) = B_\ell(x_t))$ 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 follow binomial distributions, and since $p, q \in [1/4, 3/4]$, there is a constant probability that $Y_p \ge \mathbb{E}(Y_p)$, and there is a constant probability that $Y_q \ge \mathbb{E}(Y_q)$ as well. Without loss of generality, we assume that $x_{t+1} \ge 1/2$ 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, $m \ge n/2 - 1$ tends to $+\infty$ as n tends to $+\infty$. Let $\sigma_p = \sqrt{p(1-p)}$. By the central limit theorem (Theorem 6), 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, $\text{Var}(Y_p) = m\sigma_p^2 = (n \cdot x_{t+1} - 1)p(1-p) \ge (n/2 - 1)p(1-p) \ge np(1-p)/3$, so for any $\epsilon > 0$ and n large enough,

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

By taking a small enough ϵ, and because p is bounded, we conclude the proof of Lemma 11 (the other inequality can be obtained symmetrically). □

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 12 *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.* $(x_t, x_{t+1}) \in \textsc{Yellow}'$, *we have either*

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

or $(x_{t+1}, x_{t+2}) \notin \textsc{Yellow}'$ *w.h.p.*

Proof If $\mathbb{E}(x_{t+2}) \notin [1/3, 2/3]$, then $(x_{t+1}, x_{t+2}) \notin \textsc{Yellow}'$ w.h.p. Otherwise, let $a \in [1/2 - 4\delta, 1/2 + 4\delta]$. If $a > \mathbb{E}(x_{t+2})$, then by Lemma 11,

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

Similarly, if $a \leq \mathbb{E}(x_{t+2})$, Lemma 11 gives

$$\mathbb{P}\left(|x_{t+2} - a| > \frac{1}{\sqrt{n}}\right) \geq \mathbb{P}\left(x_{t+2} \geq \mathbb{E}(x_{t+2}) + \frac{1}{\sqrt{n}}\right) \geq \beta,$$

which concludes the proof of Lemma 12. □

5.2 General structure of the proof

In order to prove Lemma 10, we first partition $\textsc{Yellow}'$, as follows (for an illustration, see Fig. 2):

$$\mathbf{A}_1 = \left\{(x_t, x_{t+1})\text{s.t.} \;\middle|\; \begin{array}{l}\text{(i) } x_{t+1} \geq 1/2, \text{ and} \\ \text{(ii) } x_{t+1} - x_t \geq x_t - 1/2.\end{array}\right\} \bigcap \textsc{Yellow}',$$

$$\mathbf{B}_1 = \left\{(x_t, x_{t+1})\text{s.t.} \;\middle|\; \begin{array}{l}\text{(i) } x_{t+1} \geq x_t, \text{ and} \\ \text{(ii) } x_{t+1} - x_t < x_t - 1/2.\end{array}\right\} \bigcap \textsc{Yellow}',$$

$$\mathbf{C}_1 = \left\{(x_t, x_{t+1})\text{s.t.} \;\middle|\; \begin{array}{l}\text{(i) } x_{t+1} < 1/2, \text{ and} \\ \text{(ii) } x_{t+1} \geq x_t.\end{array}\right\} \bigcap \textsc{Yellow}'.$$

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 $\mathbf{C} = \mathbf{C}_0 \cup \mathbf{C}_1$.

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 $\textsc{Yellow}'$, 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 $\textsc{Yellow}'$. 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 $\textsc{Yellow}'$.

Lemma 13 *If* $(x_t, x_{t+1}) \in \mathbf{A}$, *and provided that* δ *is small enough and* n *is large enough,*

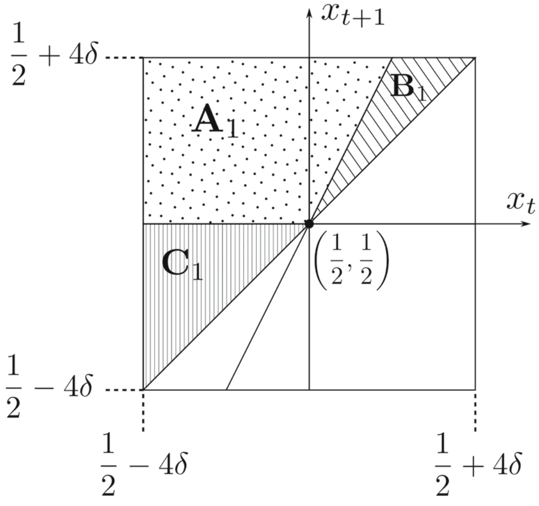

Fig. 2 Partitioning the $\textsc{Yellow}'$ domain

(a) $\mathbb{P}\left((x_{t+1}, x_{t+2}) \notin \textsc{Yellow}' \backslash \mathbf{A} \bigcap |x_{t+2} - x_{t+1}| >\right)$ $2|x_{t+1} - x_t| > 1 - \exp\left(-3n \cdot (x_{t+1} - x_t)^2\right)$.

(b) *There exists a constant* $c_2 = c_2(c) > 0$ *s.t.* $\mathbb{P}\left((x_{t+1}, \right)$ $xsps_{t+2}) \notin \textsc{Yellow}' \backslash \mathbf{A} \bigcap |x_{t+2} - x_{t+1}| > 1/\sqrt{n} > c_2$.

Proof 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,

$$\begin{aligned}
\mathbb{E}(x_{t+2}) &> \mathbb{P}\left(B_\ell\,(x_{t+1}) > B_\ell\,(x_t)\right) + x_{t+1} \\
&\cdot \mathbb{P}\left(B_\ell\,(x_{t+1}) = B_\ell\,(x_t)\right) - \frac{1}{n} \quad \text{(by Remark 3)} \\
&> \frac{1}{2} + 6(x_{t+1} - x_t) + \left(x_{t+1} - \frac{1}{2}\right) \\
&\cdot \mathbb{P}\left(B_\ell\,(x_{t+1}) = B_\ell\,(x_t)\right) \text{(by Lemma 7, taking } \lambda > 6) \\
&> \frac{1}{2} + 6(x_{t+1} - x_t). \quad \text{(by the definition of } \mathbf{A}_1)
\end{aligned}$$

More precisely, in the second inequality, the term $-1/n$ disappears within the $6(x_{t+1} - x_t)$ lower bound. Indeed, we can take $\lambda \gg 6$ (from Lemma 7) for this purpose. Moreover, we can assume $x_{t+1} - x_t \geq 1/n$, by definition of $\mathbf{A}_1$, and ruling out the case $x_t = x_{t+1}$ by Lemma 12.

Hence,

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

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

$$\mathbb{E}(x_{t+2}) > 4(x_{t+1} - x_t) + x_{t+1}. \tag{17}$$

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

$$\begin{aligned}
&\mathbb{P}\left(x_{t+2}-x_{t+1}\le 2(x_{t+1}-x_t)\right)\\
&\quad= \mathbb{P}\left(x_{t+2}\le (1-\epsilon)\left(4(x_{t+1}-x_t)+x_{t+1}\right)\right)\\
&\quad\le \mathbb{P}\left(nx_{t+2}\le (1-\epsilon)\mathbb{E}(nx_{t+2})\right) && \text{(by Eq. (17))}\\
&\quad\le \exp\left(-\frac{\epsilon^2}{2}\mathbb{E}(nx_{t+2})\right) && \text{(by Theorem 4)}\\
&\quad\le \exp\left(-\frac{2x_{t+1}}{(4(x_{t+1}-x_t)+x_{t+1})^2}(x_{t+1}-x_t)^2 n\right).\\
&\qquad\qquad\qquad\text{(by Eq. (17) and definition of } \epsilon)
\end{aligned}$$

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

$$\begin{aligned}
&\mathbb{P}\left(x_{t+2}-x_{t+1} > 2(x_{t+1}-x_t)\right)\\
&\quad\ge 1-\exp\left(-3(x_{t+1}-x_t)^2 n\right). \qquad (18)
\end{aligned}$$

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 $\textsc{Yellow}'$. □

Claim 6 *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 $(x_{t+1},x_{t+2})\notin\textsc{Yellow}'$.*

Proof If $(x_{t+1},x_{t+2})\notin\textsc{Yellow}'$, the result holds. Otherwise, $(x_{t+1},x_{t+2})\in\textsc{Yellow}'$ 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):

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

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

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

which concludes the proof of Claim 6. □

Next, we apply Claim 6 to Eq. (18) to establish (a). Finally, $x_{t+2} > x_{t+1}+4(x_{t+1}-x_t)+1/\sqrt{n}$ implies $x_{t+2}-x_{t+1} > 2(x_{t+1}-x_t)$ so we can use Claim 6,

$$\begin{aligned}
&\mathbb{P}\Big((x_{t+1},x_{t+2})\notin\textsc{Yellow}'\setminus\mathbf{A}_1 \bigcap x_{t+2}\\
&\quad > x_{t+1}+4(x_{t+1}-x_t)+1/\sqrt{n}\Big)\\
&\quad = \mathbb{P}\left(x_{t+2} > x_{t+1}+4(x_{t+1}-x_t)+1/\sqrt{n}\right) && \text{(by Claim6)}\\
&\quad > \mathbb{P}\left(x_{t+2} > \mathbb{E}(x_{t+2})+1/\sqrt{n}\right) && \text{(by Eq.(17))}\\
&\quad > c_2 > 0,
\end{aligned}$$

where the existence of c_2 is guaranteed by Lemma 11. This establishes (b).

Now, we can iteratively use the previous result to prove that any state in $\mathbf{A}$ has a reasonable chance to escape $\textsc{Yellow}'$.

Lemma 14 *There is a constant $c_3 = c_3(c)$ s.t. if $(x_{t_0},x_{t_0+1})\in\mathbf{A}$, then*

$$\mathbb{P}\left(\exists t_1 < t_0+\log n, (x_{t_1},x_{t_1+1})\notin\textsc{Yellow}'\right) > c_3.$$

Proof 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 $\textsc{Yellow}'$ in round t_0+1, and that the "gap" $(x_{t_0+2}-x_{t_0+1})$ is not too small. Formally,

$$\begin{aligned}
H_{t_0+1}:\ &(x_{t_0+1},x_{t_0+2})\\
&\notin\textsc{Yellow}'\setminus\mathbf{A}_1\bigcap x_{t_0+2}-x_{t_0+1} > 1/\sqrt{n}.
\end{aligned}$$

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

$$\begin{aligned}
H_t:\ &(x_t,x_{t+1})\notin\textsc{Yellow}'\setminus\mathbf{A}_1\bigcap x_{t+1}-x_t\\
&\quad > 2(x_t-x_{t-1}).
\end{aligned}$$

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

$$\begin{aligned}
&\bigcap_{s=t_0+1}^{t-1} H_s \Rightarrow (x_t-x_{t-1}) > 2^{(t-t_0-2)}(x_{t_0+2}-x_{t_0+1})\\
&\quad\Rightarrow (x_t-x_{t-1}) > 2^{(t-t_0-2)}/\sqrt{n}. \qquad (19)
\end{aligned}$$

For every $t > t_0+1$ such that $(x_t,x_{t+1})\in\mathbf{A}_1$,

$$\begin{aligned}
&\mathbb{P}\left(H_t \,\middle|\, \bigcap_{s=t_0+1}^{t-1} H_s\right)\\
&\quad > 1-\exp\left(-3n\cdot(x_t-x_{t-1})^2\right) && \text{(By Lemma 13)}\\
&\quad > 1-\exp\left(-\frac{3}{4}\cdot 4^{(t-t_0-1)}\right). && \text{(by Eq. (19))}
\end{aligned}$$

By Lemma 13 (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

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

Hence, we have proved that for every $t_1 > t_0 + 1$ such that $(x_{t_1}, x_{t_1+1}) \in \mathbf{A}_1$,

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

By Eq. (19), it implies that for every $t_1 > t_0 + 1$ such that $(x_{t_1}, x_{t_1+1}) \in \mathbf{A}_1$,

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

For t_1 large enough (e.g., $t_1 = t_0 + \log n$), this implies that $(x_{t_1-1}, x_{t_1}) \notin$ 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 14. □

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 rather long and is deferred to Sect. 5.3.

Lemma 15 *There are constants $c_4, c_5 > 0$ such that if $(x_t, x_{t+1}) \in \mathbf{B}$, then either*

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

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.

Lemma 16 *If $(x_{t_0}, x_{t_0+1}) \in \mathbf{B}$, then, w.h.p.,* $\min\{t \geq t_0, (x_t, x_{t+1}) \notin \mathbf{B}\} < t_0 + \frac{\sqrt{c}}{c_4} \cdot \log^{3/2} n$.

Proof 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 15 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 15 doesn't hold, (b) of Lemma 15 holds so there is a constant probability to leave $\mathbf{B}$, so

$$\begin{aligned}&\mathbb{P}\,((\text{ for every } t \text{ such that } t_0 \leq t \leq t_{\max},\\ &\quad (x_t, x_{t+1}) \in \mathbf{B}) \bigcap X = \mathbf{x}\big) \leq (1 - c_5)^{\mathbf{x}}.\end{aligned} \tag{20}$$

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. (20), 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 16 holds).

Now, consider case (i). Let $u_t = x_t - 1/2$. By Lemma 12, after $O(\log n)$ rounds, we'll have $u_t > 1/\sqrt{n}$ w.h.p., so up to waiting a logarithmic number of rounds, we assume that $u_{t_0} > 1/\sqrt{n}$.

Note 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

$$u_t \leq u_{t+1}. \tag{21}$$

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 15,

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

Hence, by Eq. (21),

$$\begin{aligned}u_{t_{\max}} &> u_{t_0} \cdot \left(1 + \frac{c_4}{\sqrt{\ell}}\right)^k\\ &= u_{t_0} \cdot \exp\left(k \log\left(1 + \frac{c_4}{\sqrt{\ell}}\right)\right)\\ &> u_{t_0} \cdot \exp\left(\frac{4}{5} \cdot \frac{k \cdot c_4}{\sqrt{\ell}}\right) \qquad \text{(for } n \text{ large enough)}\\ &= u_{t_0} \cdot \exp\left(\frac{4}{5} \cdot \frac{3}{4} \log n\right) \qquad \text{(by definition of } k \text{ and } \ell\text{)}\end{aligned}$$

$$= u_{t_0} \cdot n^{3/5}$$
$$> n^{1/10}. \qquad \text{(since we assumed } u_{t_0} > 1/\sqrt{n}\text{)}$$

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 16. □

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.

Lemma 17 *There is a constant $c_6 > 0$ such that if $(x_t, x_{t+1}) \in \mathbf{C}$, then*

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

Proof 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

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

By Lemma 7 (taking $\lambda > 2$), this becomes

$$\mathbb{E}(x_{t+2}) > \frac{1}{2} + 2 \cdot (x_{t+1} - x_t) - \left(\frac{1}{2} - x_{t+1}\right) \cdot \mathbb{P}\left(B_\ell\left(x_{t+1}\right) = B_\ell\left(x_t\right)\right). \tag{22}$$

Case 1. If $(x_{t+1} - x_t) > 1/2 - x_{t+1}$, then Eq. (22) 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 YELLOW$'$.

Case 2. Else, if $(x_{t+1} - x_t) \leq 1/2 - x_{t+1}$, Eq. (22) rewrites

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

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, $1 - 2 \cdot \mathbb{P}\left(B_\ell(x_{t+1}) = B_\ell(x_t)\right) > 0$. Hence,

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

If $\mathbb{E}(x_{t+2}) \notin [1/3, 2/3]$, then $(x_{t+1}, x_{t+2}) \notin \mathbf{A}_1$ w.h.p. Otherwise, we can apply Lemma 11: with constant probability $x_{t+2} > (1/2 + x_{t+1})/2$, i.e., $x_{t+2} - x_{t+1} > 1/2 - x_{t+2}$. If so, Case 1 applies and with constant probability, $(x_{t+2}, x_{t+3}) \in \mathbf{A}_1$ or is not in YELLOW$'$. This concludes the proof of Lemma 17. □

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

Proof of Lemma 10 By Lemma 14, if $(x_{t_0}, x_{t_0+1}) \in \mathbf{A}$, then

$$\mathbb{P}\left(\exists t_1 < t_0 + \log n, (x_{t_1}, x_{t_1+1}) \notin \text{YELLOW}'\right) > c_3 > 0.$$

By Lemma 17, this implies that if $(x_{t_0}, x_{t_0+1}) \in \mathbf{A} \cup \mathbf{C}$,

$$\mathbb{P}\left(\exists t_1 < t_0 + \log n + 2, (x_{t_1}, x_{t_1+1}) \notin \text{YELLOW}'\right) > \min(c_3, c_3 \cdot c_6) = c_3 \cdot c_6 > 0. \tag{23}$$

By Lemma 16, 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. (23), 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 10. □

5.3 Proof of Lemma 15—regarding Area B

The goal of this section is to prove Lemma 15, which concerns Area 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). Let us define

$$g(x, y) = \mathbb{P}\left(B_\ell\left(y\right) > B_\ell\left(x\right)\right) + y \cdot \mathbb{P}\left(B_\ell\left(y\right) = B_\ell\left(x\right)\right) + \frac{1}{n}\left(1 - \mathbb{P}\left(B_\ell\left(y\right) \geq B_\ell\left(x\right)\right)\right), \tag{24}$$

so that, conditioning on (x_t, x_{t+1}), $\mathbb{E}(x_{t+2}) = g(x_t, x_{t+1})$ by Observation 1. Informally, our plan is the following. For all points (x_t, x_{t+1}) such that $\mathbb{E}(x_{t+2}) = g(x_t, x_{t+1})$ is "small" compared to x_{t+1}, then the process will lose speed and get out of $\mathbf{B}$, corresponding to item (b) in the statement of Lemma 15. Otherwise, $\mathbb{E}(x_{t+2}) = g(x_t, x_{t+1})$ is sufficiently "large" compared to x_{t+1}, and the process makes a significant step towards escaping YELLOW, corresponding to item (a).

We will proceed in two steps: first, we analyze function g, and only then, we prove Lemma 15.

Analysis of *g*

We start with the following claim, which will be used to prove the subsequent claim.

Claim 7 *Let $x_0 \in [1/3, 2/3]$. On the interval $[x_0, x_0+1/\sqrt{\ell}]$, and for ℓ large enough, $y \mapsto g(x_0, y)-y$ is a strictly increasing function of y.*

Proof The key observation is that the derivative, w.r.t. x, of $\mathbb{P}\left(B_k(x) > B_k(p)\right)$ in the neighborhood of p is relatively high. The following claim formalizes this idea. □

Claim 8 *There exists a constant $\beta' > 0$ such that for every k large enough, and every $p, x \in [1/3, 2/3]$ satisfying $p \le x \le p + 1/\sqrt{k}$,*

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

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

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

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

$$\begin{aligned}
&\mathbb{P}\left(B_k(x+h) > B_k(p)\right) - \mathbb{P}\left(B_k(x) > B_k(p)\right) \\
&\quad \ge \sum_{j \in J} \mathbb{P}\left(Z = j\right) \cdot \mathbb{P}\left(Y_1 \le j \bigcap Y_2 > j\right).
\end{aligned} \tag{25}$$

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 $|i - kp| \le \sqrt{k}$, we have $\mathbb{P}\left(B_k(p) = i\right) \ge \frac{\beta}{\sqrt{k}}$.*

Proof By the De Moivre-Laplace theorem, for any i in $\{kp - \sqrt{k}, \ldots, kp + \sqrt{k}\}$,

$$\begin{aligned}
\mathbb{P}\left(B_k(p) = i\right) &= \binom{k}{i} p^i (1-p)^{k-i} \\
&\approx \frac{1}{\sqrt{2kp(1-p)}} \exp\left(-\frac{(i-kp)^2}{2kp(1-p)}\right),
\end{aligned} \tag{26}$$

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 $|i - kp| \le \sqrt{k}$,

$$\begin{aligned}
&\frac{1}{\sqrt{2kp(1-p)}} \exp\left(-\frac{(i-kp)^2}{2kp(1-p)}\right) \\
&\ge \frac{1}{\sqrt{2kp(1-p)}} \exp\left(-\frac{1}{2p(1-p)}\right) := \frac{f(p)}{\sqrt{k}}.
\end{aligned}$$

By Eq. (26), 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, $\mathbb{P}\left(Z = j\right) \ge \beta/\sqrt{k}$, for some constant $\beta > 0$. Moreover,

$$\begin{aligned}
&\mathbb{P}\left(Y_1 \le j \bigcap Y_2 > j\right) \\
&\ge \mathbb{P}\left(Y_1 = j \bigcap Y_2 > j\right) \\
&= \mathbb{P}\left(Y_1 = j \bigcap H \ge 1\right) \qquad \text{(because } Y_2 = Y_1 + H\text{)} \\
&= \mathbb{P}\left(Y_1 = j\right) \cdot \mathbb{P}\left(H \ge 1 \mid Y_1 = j\right).
\end{aligned}$$

By the assumption in the lemma, $p \le x \le p + 1/\sqrt{k}$, and so $kp \le kx \le kp + \sqrt{k}$. Therefore, for $j \in J$, $|j - kx| \le \sqrt{k}$, and by Observation 2, we get that $\mathbb{P}\left(Y_1 = j\right) \ge \beta/\sqrt{k}$. Hence, we can rewrite Eq. (25) as

$$\begin{aligned}
&\mathbb{P}\left(B_k(x+h) > B_k(p)\right) - \mathbb{P}\left(B_k(x) > B_k(p)\right) \\
&\quad \ge \frac{\beta^2}{k} \sum_{j \in J} \mathbb{P}\left(H \ge 1 \mid Y_1 = j\right).
\end{aligned} \tag{27}$$

Now, let us find a lower bound on $\mathbb{P}\left(H \ge 1 \mid Y_1 = j\right)$, 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, $1 \le i \le k$, is uniformly distributed over $[0, 1]$,

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

Therefore, for every $j \in J$

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

This implies that

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

We have

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

Eventually, we get from Eq. (27)

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

We can conclude the proof of Claim 8 for k large enough by taking, e.g., $\beta' = \frac{\beta^2(1-p)}{2(1-x)}$.

Now, we are ready to conclude the proof of Claim 7. We can rewrite Eq. (24) as

$$g(x_0, y) = \left(y - \frac{1}{n}\right)\mathbb{P}\left(B_\ell\left(y\right) \geq B_\ell\left(x_0\right)\right) + (1-y) \\ \cdot \mathbb{P}\left(B_\ell\left(y\right) > B_\ell\left(x_0\right)\right) + \frac{1}{n}.$$

Hence,

$$\frac{d}{dy}g(x_0, y) \\ = \left[\mathbb{P}\left(B_\ell\left(y\right) \geq B_\ell\left(x_0\right)\right) - \mathbb{P}\left(B_\ell\left(y\right) > B_\ell\left(x_0\right)\right)\right] \\ + \left(y - \frac{1}{n}\right) \cdot \frac{d}{dy}\mathbb{P}\left(B_\ell\left(y\right) \geq B_\ell\left(x_0\right)\right) + (1-y) \\ \cdot \frac{d}{dy}\mathbb{P}\left(B_\ell\left(y\right) > B_\ell\left(x_0\right)\right). \tag{28}$$

The first term in Eq. (28) is equal to $\mathbb{P}\left(B_\ell\left(y\right) = B_\ell\left(x_0\right)\right)$, which is positive. Moreover, $\mathbb{P}\left(B_\ell\left(y\right) \geq B_\ell\left(x_0\right)\right)$ is obviously increasing in y, so the second term is also non-negative. By Claim 8, the third term in Eq. (28) satisfies

$$(1-y) \cdot \frac{d}{dy}\mathbb{P}\left(B_\ell\left(y\right) > B_\ell\left(x_0\right)\right) \\ \geq (1-y) \cdot \beta' \cdot \sqrt{\ell} \geq \frac{\beta'}{4} \cdot \sqrt{\ell},$$

where the last inequality comes from the fact that $x_0 \in [1/3, 2/3]$ and $y \in [x_0, x_0 + 1/\sqrt{\ell}] \subseteq [1/4, 3/4]$. For ℓ large enough, this implies that

$$\frac{d}{dy}g(x_0, y) \geq \frac{\beta'}{4} \cdot \sqrt{\ell} > 1,$$

which concludes the proof of Claim 7.

Finishing the proof

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

Claim 9 *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 7, 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 $y^\star \in [x, x + 1/\sqrt{\ell}]$ such that $h(y^\star) = 0$; 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. □

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 9, with this definition we always have

$$g(x, f(x)) \leq f(x). \tag{29}$$

Claim 10 *For any $x \in [1/2 + 4/n, 1/2 + 4\delta]$, it holds that*

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

where $\alpha > 1$ is the constant stated in Lemma 9, stated in Sect. 3.

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 9, there exists $\alpha > 0$ s.t.

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

This can be plugged into the definition of f (Eq. (24)) to give

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

which we can rewrite,

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

This gives

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

where the last inequality comes from $\mathbb{P}\left(B_\ell(f(x)) = B_\ell(x)\right) < 1/2$ (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 10. □

Next, rewriting $f(x) - x = (f(x) - 1/2) - (x - 1/2)$, we get from Claim 10 that for every $x \in [1/2 + 4/n, 1/2 + 4\delta]$,

$$\left(f(x) - \frac{1}{2}\right) > \left(1 + \frac{1}{4\alpha\sqrt{\ell}}\right) \cdot \left(x - \frac{1}{2}\right). \tag{30}$$

We are now ready to conclude the proof of Lemma 15. Let $c_4 = 1/4\alpha$. We will use the fact that, within the Yellow area, $x_{t+2} \leq \mathbb{E}(x_{t+2}) - 1/\sqrt{n}$ and $x_{t+2} \geq \mathbb{E}(x_{t+2}) + 1/\sqrt{n}$ both happen with constant probability (Lemma 11) or the system leaves the area.

- 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 $x_{t+2} \in [1/2, 1/2 + 16/n]$. 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. (30),

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

  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:

  $$x_t + 1/\sqrt{\ell} \geq f(x_t) \geq x_{t+1} \geq x_t. \tag{31}$$

  By Eq. (29), $g(x_t, f(x_t)) \leq f(x_t)$, i.e., $g(x_t, f(x_t)) - f(x_t) \leq 0$. By Claim 7, for ℓ large enough, function $y \mapsto g(x_t, y) - y$ is strictly increasing on $[x_t, x_t + 1/\sqrt{\ell}]$. Equation (31) 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, $\mathbb{E}(x_{t+2}) \leq x_{t+1}$.
  Hence, 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 15. □

6 Other areas

6.1 Proof of Lemma 1—Green area

The goal of this section is to prove Lemma 1, which concerns Area Green. 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. (14) in Remark 2, we have for every agent i

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

By Lemma 6, we have

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

Then, by the union bound,

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

which is $o(n^{-\epsilon})$ for some $\epsilon > 0$ provided that $c > 2/\delta^2$. □

6.2 Proof of Lemma 2—Purple area

The goal of this section is to prove Lemma 2, which concerns Area Purple. 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. (15) 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 $x_{t+1} \geq (1-\lambda_n)x_t$, 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

$$\sigma = \sqrt{x_t(1-x_t)+(1-\lambda_n)x_t(1-(1-\lambda_n)x_t)}$$
$$> \sqrt{x_t(1-x_t)} > \sqrt{\frac{x_t}{2}}, \tag{32}$$

where the last inequality is by the fact that $x_t < 1/2$ which follows from the definition of Purple_1. By Lemma 8,

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

We have (Eq. (32) 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

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

where the second inequality is by Eq. (32), 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 4) 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. □

6.3 Proof of Lemma 3—Red area

The goal of this section is to prove Lemma 3, which concerns Area Red. 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. □

6.4 Proof of Lemma 4—Cyan area

The goal of this section is to prove Lemma 4, which concerns Area CYAN. We only prove the result for CYAN$_1$, but the same arguments apply to CYAN$_0$ symmetrically. 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 $(x_{t_0+1}, x_{t_0+2}) \in$ CYAN$_1$ (but this time Case 2 applies).
- or $x_{t_0+2} \geq x_{t_0+1} + \delta$, and so $(x_{t_0+1}, x_{t_0+2}) \in$ 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 11 below concerns small values of x_{t+1}, Claim 12 concerns intermediate values of x_{t+1}, and Claim 13 concerns large values of x_{t+1}.

Claim 11 *If $x_t < 1/\log(n)$, and if $0 < x_{t+1} \leq 1/\ell$, then*

$$\mathbb{P}\left(x_{t+2} > \frac{K}{2} x_{t+1} \log n\right) > 1 - \exp\left(-\frac{K}{8} \log n\right).$$

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

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

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 one 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

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

Hence, by Eq. (15) in Remark 2, $\mathbb{E}\left(x_{t+2}\right) \geq K x_{t+1} \log n - 1/n$. By Observation 1, we can apply Chernoff's inequality (Theorem 4) to conclude the proof of Claim 11. □

Claim 12 *If $x_t < 1/\log(n)$, and if $1/\ell < x_{t+1} \leq \gamma$, then*

$$\mathbb{P}\left(x_{t+2} > \gamma\right) > 1 - \exp\left(-\frac{\gamma n}{8}\right) > 1 - \exp\left(-\frac{K}{8} \log n\right).$$

Proof The proof follows along similar lines as the proof of Claim 11. We note that, since $x_t < 1/\log(n)$, the probability that an agent does not see a 1 in round t is

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

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

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

Hence, by Eq. (15) in Remark 2, $\mathbb{E}\left(x_{t+2}\right) \geq 2\gamma - 1/n$. By Observation 1, we can apply Chernoff's inequality (Theorem 4) to conclude the proof of Claim 12. □

Claim 13 *If $x_t < 1/\log(n)$, and if $x_{t+1} > \gamma$, then*

$$\mathbb{P}\left(x_{t+2} > \frac{1}{2}\right) > 1 - \exp\left(-\frac{n}{18}\right) > 1 - \exp\left(-\frac{K}{8} \log n\right).$$

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

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

for n large enough. Hence, by Eq. (15) in Remark 2, $\mathbb{E}(x_{t+2}) \geq 3/4 - 1/n$. By Observation 1, we can apply Chernoff's inequality (Theorem 4) to conclude the proof of Claim 13. □

We say that a round t is *successful* if $(x_t, x_{t+1}) \in \textsc{Cyan}_1$, and the event of either Claim 11, 12 or 13 happens. 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 $\textsc{Cyan}_1$), $\log(n/\ell)/\log(K \cdot \log(n)/2)$ rounds are always enough to get $x_{t+1} > 1/\ell$; one more round is enough to get $x_{t+1} > \gamma$; and one more round is enough to get $x_{t+1} > 1/2$, in which case $(x_t, x_{t+1}) \notin \textsc{Cyan}_1$. Therefore, by Claims 11, 12 and 13, the probability that, starting from t_0, all rounds are successful until the system is out of $\textsc{Cyan}_1$ is at least $\left(1 - \exp\left(-\frac{K}{8}\log n\right)\right)^{X_{\max}} \geq 1 - X_{\max} \cdot \exp\left(-\frac{K}{8}\log n\right) = 1 - 1/n^{\Omega(1)}$. 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 $\textsc{Cyan}_1$). Thus, by construction of the partition, we must have $(x_{t+1}, x_{t+2}) \in \textsc{Cyan}_1 \cup \textsc{Green}_1 \cup \textsc{Purple}_1$. This implies that $(x_{t_1}, x_{t_1+1}) \in \textsc{Green}_1 \cup \textsc{Purple}_1$, which concludes the proof of Lemma 4. □

7 Extension to more than two opinions

In this section, we consider the more general case with k opinions, for an arbitrary $k \in \mathbb{N}$, and we prove Theorem 2. We assume that agents can agree on a labeling of the opinions beforehand, and hence, on a total order over the opinion space. Therefore, without loss of generality, we may assume that the set of opinions is $\mathcal{Y} = \{0, \ldots, k-1\}$.[2] Investigating settings where agents are denied this common knowledge is left for future work.

Let $m = \lceil \log_2 k \rceil \in \mathbb{N}$. Given an opinion $y \in \mathcal{Y}$, we identify y with its binary representation (as an m-bit string). We write $y[j]$ to denote the j^{th} bit of y, with the convention that $y[1]$ is the most significant bit, and $y[m]$ the least significant. A solution to the bit-dissemination problem in this setting is to let the agents execute Protocol 1 independently and simultaneously on each bit of the opinions. The analysis can then be reduced to the one of Protocol 1.

However, before it reaches a consensus on every bit of the correct opinion, the protocol might face the following issue. Whenever the most significant bit is set to 1, the resulting integer may be too large, i.e., may belong to $\{k, \ldots, 2^m - 1\}$, and therefore not correspond to any opinion in $\mathcal{Y}$. Thus, to ensure that agents only output valid opinions, we need to set the most significant bit to 0 whenever it is necessary; that is, whenever the opinion's value exceeds k.[3] We will show that this modification does not prevent the protocol from solving the bit-dissemination problem.

The internal state space of our algorithm is $\Sigma = \{0, \ldots, \ell\}^m$. In the state $\sigma_t^{(i)} = (\sigma_t^{(i)}[1], \ldots, \sigma_t^{(i)}[m])$ of Agent i in round t, $\sigma_t^{(i)}[j] \in \{0, \ldots, \ell\}$ represents the number of samples received in the previous round, for which the j^{th} bit is equal to 1. Our protocol is described in details by Protocol 2 below.

Protocol 2: FET bit by bit

```
  Input  : $Y_t \in \mathcal{Y} = \{0, \ldots, k-1\}$, $\sigma_t \in \Sigma = \{0, \ldots, \ell\}^m$, $S_t \in \mathcal{Y}^{2\ell}$
1 for $j \in \{1, \ldots, m\}$ do
2     Let $S_t[j] \in \{0, 1\}^{2\ell}$ be a vector containing the $j^{\text{th}}$ bit of all elements in $S_t$;
3     $Y_{t+1}[j], \sigma_{t+1}[j] \leftarrow$ Protocol 1 $(Y_t[j], \sigma_t[j], S_t[j])$ ;
4 end
5 if $Y_{t+1} \notin \mathcal{Y}$ then
6     $Y_{t+1}[1] \leftarrow 0$ ;
7 end
  Output : $(Y_{t+1}, \sigma_{t+1})$
```

Theorem 2 is a corollary of the following result.

Theorem 3 *Let k be a positive integer and let $m = \lceil \log_2 k \rceil$. When $\mathcal{Y} = \{0, \ldots, k-1\}$, Protocol 2 solves the bit-dissemination problem in $O(\log^{5/2} n)$ rounds with high probability, while relying on $\ell = \Theta(\log n)$ samples in each round and using $\Theta(m \log \ell)$ bits of memory.*

Proof We start by making the following simple observation. □

Claim 14 *Any m-bit string a with $a[1] = 0$ corresponds to a valid opinion in $\mathcal{Y}$.*

Proof Since the most significant bit of a is equal to 0, we have $a \leq 2^{m-1} - 1$. Moreover, $m = \lceil \log_2 k \rceil < 1 + \log_2 k$, so $2^{m-1} - 1 < 2^{\log_2 k} - 1 = k - 1$. Hence, $a \in \{0, \ldots, k-2\} \subset \mathcal{Y}$. □

By Claim 14, and because of the "if" statement in line 5, Protocol 2 only outputs valid opinions.

Now, let $z \in \mathcal{Y}$ be the opinion of the source. By construction, Protocol 2 behaves exactly as Protocol 1 with respect to every bit j, for $j \in \{2, \ldots, m\}$. Therefore, by Theorem 1, and

[2] Indeed, whenever an agent sees the i'th opinion in the total order over $\mathcal{Y}$, he can treat it as if it is opinion i in the set $\{0, \ldots, k-1\}$.

[3] We would like to thank one of the anonymous reviewers for suggesting this implementation which is more elegant than our original idea, although closely related.

since k is fixed w.r.t. n, w.h.p. there exists $t_0 = O(\log^{5/2} n)$, such that:

$$\text{for every } t \geq t_0, \text{ every } j \in \{2, \dots, m\}, \\ \text{and every } i \in I, \quad Y_t^{(i)}[j] = z[j]. \tag{33}$$

Now, consider integers z_0, z_1, given by their binary representations: $z_0 = (0, z[2], \dots, z[m])$, and $z_1 = (1, z[2], \dots, z[m])$. By definition, either $z = z_0$ or $z = z_1$. By Claim 14, $z_0 \in \mathcal{Y}$.

- If $z_1 \notin \mathcal{Y}$, then necessarily the opinion of the source is z_0. Moreover, by the "if" statement on line 5 of Protocol 2, each agent adopts opinion z_0 from round $t_0 + 1$ onward.
- Otherwise, if $z_1 \in \mathcal{Y}$, then from round t_0 and by construction, Protocol 2 behaves exactly as Protocol 1 with respect to the most significant bit.Therefore, there exists a round $t_1 = t_0 + O(\log^{5/2} n)$ such that all agents agree on the most significant bit. Together with Eq. (33), this implies that from round t_1 onward, all agents adopt opinion z.

In both cases, Protocol 2 converges in $O(\log^{5/2} n)$ rounds w.h.p. Finally, we note that the $O(m \log \ell)$ bits upper bound on the memory complexity follows from the fact that the internal state space $\Sigma = \{0, \dots, \ell\}^m$ is only of size $(\ell + 1)^m$, which concludes the proof of Theorem 3.

8 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 $\Theta(\log n)$ 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.

In Sect. 7 we study the case of multiple opinions. We showed how to handle this case assuming that the agents agree on the ordering of the opinions. As mentioned, the case when there is a possible conflict in the ordering of the opinions remains for future work.

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 [39] or ants [37].

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. Finally, we would like to thank the anonymous reviewers for their meticulous proofreading of the manuscript and their insightful comments, which notably improved Sect. 7. 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).

Author Contributions All authors contributed equally.

Funding Open access funding provided by University of Haifa.

Declarations

Conflict of interest The authors declare no competing interests.

Appendix

A Probabilistic tools–Some well-known theorems

Theorem 4 *[Multiplicative Chernoff's Bound] Let $X_1, \ldots, X_n$ 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*

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

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

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

Theorem 5 *[Hœffding's bound] Let $X_1, \ldots, X_n$ 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*

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

Theorem 6 *[Central Limit] Let $X_1, \ldots, X_n$ be i.i.d. random variables with $\mathbb{E}(X_1) = \mu$ and $\mathrm{Var}(X_1) = \sigma^2 < +\infty$. Then as n tends to infinity, the random variables $\sqrt{n}\big(\frac{1}{n}\sum_{i=1}^n X_i - \mu\big)$ 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 7 *[Berry-Esseen] Let $X_1, \ldots, X_n$ be i.i.d. random variables, with $\mathbb{E}(X_1) = 0$, $\mathrm{Var}(X_1) = \mathbb{E}(X_1^2) = \sigma^2 > 0$, 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., $C = 0.4748$.

References

1. Alistarh, D., Aspnes, J., Eisenstat, D., Gelashvili, R., Rivest, R.L.: Time-space trade-offs in population protocols. In: Proceedings of the 2017 annual ACM-SIAM symposium on discrete algorithms (SODA), Proceedings, pp. 2560–2579. Society for Industrial and Applied Mathematics
2. Alistarh, D., Gelashvili, R.: Polylogarithmic-time leader election in population protocols. In: Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II, volume 9135 of Lecture Notes in Computer Science, pp. 479–491. Springer (2015)
3. Alistarh, D., Gelashvili, R.: Recent algorithmic advances in population protocols. SIGACT News **49**(3), 63–73 (2018)
4. Altisen, K., Devismes, S., Dubois, S., Petit, F.: Introduction to distributed self-stabilizing algorithms. Springer Nature
5. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M.J., Peralta, R.: Computation in networks of passively mobile finite-state sensors. Distrib. Comput. **18**(4), 235–253 (2006)
6. Angluin, D., Aspnes, J., Eisenstat, D.: A simple population protocol for fast robust approximate majority. Distrib. Comput. **21**(2), 87–102 (2008)
7. Angluin, D., Aspnes, J., Fischer, M.J., Jiang, H.: Self-stabilizing population protocols. ACM Trans. Auton. Adapt. Syst. (TAAS) **3**(4), 131–1328 (2008)
8. Aspnes, J., Ruppert, E.: An introduction to population protocols. Bull. EATCS **93**, 98–117 (2007)
9. Ayalon, O., Sternklar, Y., Fonio, E., Korman, A., Gov, N.S., Feinerman, O.: Sequential decision-making in ants and implications to the evidence accumulation decision model. Front. Appl. Math. Stat. **7**, 672773 (2021)
10. Barclay, R.M.R.: Interindividual use of echolocation calls: eavesdropping by bats. Behav. Ecol. Sociobiol. **10**(4), 271–275 (1982)
11. Bastide, P., Giakkoupis, G., Saribekyan, H.: Self-stabilizing clock synchronization with 1-bit messages. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2154–2173, Alexandria, Virginia, U.S. SIAM (2021)
12. Becchetti, L., Clementi, A., Natale, E.: Consensus dynamics: an overview. ACM SIGACT News **51**(1), 58–104 (2020)
13. Boczkowski, L., Feinerman, O., Korman, A., Natale, E.: Limits for rumor spreading in stochastic populations. In: Anna R. Karlin, editor, 9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA, volume 94 of LIPIcs, pp. 49:1–49:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018)
14. Boczkowski, L., Korman, A., Natale, E.: Minimizing message size in stochastic communication patterns: fast self-stabilizing protocols with 3 bits. Distrib. Comput. **32**(3), 173–191 (2019)
15. Boczkowski, L., Natale, E., Feinerman, O., Korman, A.: Limits on reliable information flows through stochastic populations. PLoS Comput. Biol. **14**(6), e1006195 (2018)
16. Censor-Hillel, K., Haeupler, B., Kelner, J., Maymounkov, P.: 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, pp. 961–970 (2012)
17. Chen, H.L., Cummings, R., Doty, D., Soloveichik, D.: Speed faults in computation by chemical reaction networks. In: International Symposium on Distributed Computing, pp. 16–30. Springer (2014)
18. Chierichetti, F., Giakkoupis, G., Lattanzi, S., Panconesi, A.: Rumor spreading and conductance. J. ACM (JACM) **65**(4), 1–21 (2018)
19. Cvikel, N., Berg, K.E., Levin, E., Hurme, E., Borissov, I., Boonman, A., Amichai, E., Yovel, Y.: Bats aggregate to improve prey search but might be impaired when their density becomes too high. Curr. Biol. **25**(2), 206–211 (2015)
20. Danchin, E., Giraldeau, L.-A., Valone, T.J., Wagner, R.H.: Public information: from nosy neighbors to cultural evolution. Science **305**(5683), 487–491 (2004)
21. Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H., Swinehart, D., Terry, D.: Epidemic algorithms for replicated database maintenance. In: Proceedings of the sixth annual ACM Symposium on Principles of distributed computing, pp. 1–12 (1987)

22. Dijkstra, E.W.: Self-stabilizing systems in spite of distributed control. Commun. ACM **17**(11), 643–644 (1974)
23. Doerr, B., Goldberg, L.A., Minder, L., Sauerwald, T., Scheideler.: Stabilizing consensus with the power of two choices. In: Proceedings of the twenty-third annual ACM symposium on Parallelism in algorithms and architectures, pp. 149–158 (2011)
24. Dolev, S.: Self-Stabilization. MIT Press, Cambridge (2000)
25. Dutta, C., Pandurangan, G., Rajaraman, R., Sun, Z., Viola, E.: On the complexity of information spreading in dynamic networks. In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 717–736. SIAM (2013)
26. Emek, Y., Keren, E.: A thin self-stabilizing asynchronous unison algorithm with applications to fault tolerant biological networks. In: Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing, PODC'21, pp. 93–102. Association for Computing Machinery
27. Feinerman, O., Haeupler, B., Korman, A.: Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication. Distrib. Comput. **30**(5), 339–355 (2017)
28. Feinerman, O., Korman, A.: Individual versus collective cognition in social insects. J. Exp. Biol. **220**(1), 73–82 (2017)
29. Georgiou, C., Gilbert, S., Guerraoui, R., Kowalski, D.R.: Asynchronous gossip. J. ACM **60**(2), 11:1-11:42 (2013)
30. Georgiou, C., Gilbert, S., Kowalski, D.R.: Meeting the deadline: on the complexity of fault-tolerant continuous gossip. Distrib. Comput. **24**(5), 223–244 (2011)
31. Giakkoupis, G.: Tight bounds for rumor spreading with vertex expansion. In: Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms, pp. 801–815. SIAM (2014)
32. Giakkoupis, G., Woelfel, P.: On the randomness requirements of rumor spreading. In: Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 449–461. SIAM (2011)
33. Giraldeau, L.A., Caraco, T.: Social foraging theory. Princeton University Press, Princeton (2018)
34. Karp R., Schindelhauer C., Shenker S., Vocking B.: Randomized rumor spreading. In: Proceedings 41st Annual Symposium on Foundations of Computer Science, pp. 565–574. IEEE (2000)
35. Korman, A., Greenwald, E., Feinerman, O.: Confidence sharing: an economic strategy for efficient information flows in animal groups. PLoS Comput. Biol. **10**(10), e1003862 (2014)
36. Liggett, T.M., Liggett, T.M.: Interacting particle systems, vol. 2. Springer, Berlin (1985)
37. Rajendran, H., Haluts, A., Gov, N.S., Feinerman, O.: Ants resort to majority concession to reach democratic consensus in the presence of a persistent minority. Current Biol. **32**(3), 645–53 (2022)
38. Razin, N., Eckmann, J.-P., Feinerman, O.: Desert ants achieve reliable recruitment across noisy interactions. J. R. Soc. Interface **10**(82), 20130079 (2013)
39. Sumpter, D.J.T., Krause, J., James, R., Couzin, I.D., Ward, A.J.W.: Consensus decision making by fish. Current Biol. **18**(22), 1773–1777 (2008)
40. Wilkinson Gerald, S.: Information transfer at evening bat colonies. Anim. Behav. **44**, 501–518 (1992)
41. Yick, J., Mukherjee, B., Ghosal, D.: Wireless sensor network survey. Comput. Netw. **52**(12), 2292–2330 (2008)