Asynchronous Majority Dynamics on Binomial Random Graphs*

Divyarthi Mohan† Paweł Prałat‡

Abstract

We study information aggregation in networks when agents interact to learn a binary state of the world. Initially each agent privately observes an independent signal which is *correct* with probability $\frac{1}{2}+\delta$ for some $\delta > 0$. At each round, a node is selected uniformly at random to update their public opinion to match the majority of their neighbours (breaking ties in favour of their initial private signal). Our main result shows that for sparse and connected binomial random graphs $\mathcal{G}(n,p)$ the process stabilizes in a *correct* consensus in $\mathcal{O}(n\log^2 n/\log\log n)$ steps with high probability. In fact, when $\log n/n \ll p = o(1)$ the process terminates at time $\hat{T} = (1+o(1))n\log n$, where $\hat{T}$ is the first time when all nodes have been selected at least once. However, in dense binomial random graphs with $p = \Omega(1)$, there is an information cascade where the process terminates in the *incorrect* consensus with probability bounded away from zero.

1 Introduction

Our opinions and actions we take as individuals are often influenced by both our private knowledge of the world and the information we obtain through our interactions with others. For example, a voter deciding which candidate's economics policies would decrease inflation, might have an initial belief based on her own past expenditure and later might be swayed by her friends' opinions. Now more than ever, with the advent of social media and online platforms, our interactions have increased many folds and our social networks are massive. Hence, an important research question is to understand if and how the structure of the social network and the dynamics of the interactions impact the (mis)information propagated [31]. Do our social networks enable successful information aggregation and lead to social learning, or do they amplify incorrect beliefs leading to an information cascade?

There has been extensive work modeling these opinion dynamics formally to study the network effects on information aggregation; see Section 1.4. In this paper, we focus on the model of *asynchronous majority dynamics*, where agents in a network (asynchronously) update their opinions to match the majority opinion amongst their neighbours. In particular, each agent initially has a private belief over a binary state of the world and no publicly announced opinion. At each time step, an agent is chosen uniformly at random to announce/update her opinion and she does so by simply copying the majority of the neighbours' current opinions, breaking ties with her initial belief. In our model, there is a *correct* opinion (i.e., the true state of the world) and each agent's initial private belief is independently drawn and is biased towards being correct (with probability $1/2+\delta$). So initially, in a large network, there is enough information so that an omniscient central

*This work was done while the authors were visiting the Simons Institute for the Theory of Computing.

†Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel, email: `divyarthim@tau.ac.il`.

‡Department of Mathematics, Toronto Metropolitan University, Toronto, ON, Canada, email: `pralat@torontomu.ca`.

planner can infer the true state (with very high probability). However, agents in the network are updating their opinions based on local heuristics, so the network structure can crucially alter the final outcome of the dynamics. For example, in a complete graph, with a constant probability all the nodes converge to the wrong opinion. On the other hand, in a star graph with high probability all the nodes converge to the correct opinion. This brings us to the main question of interest:

> *"What network structures enable efficient social learning, where the dynamics stabilizes with every agent in the network reaching the correct opinion?"*

Feldman et al. [18], who initiated the study of asynchronous majority dynamics, showed that when the network is sparse (has bounded degree) and expansive, a correct consensus is reached with high probability. More recently, Bahrani et al. [3] studied networks that have certain tree structures (like preferential attachment trees and balanced m-ary trees) and showed that the dynamics stabilizes in a correct majority. Both results heavily rely on these particular assumptions on the network. For example, to even establish that a *majority of the nodes* have the correct opinion *at some point* in the process, it is crucial that the network is either a bounded degree graph or is a tree. In this paper, our goal is to extend the guarantees of asynchronous majority dynamics beyond these assumptions and to develop techniques applicable to more general networks formed through random graph models.

1.1 The Model

Consider any undirected graph $G = (V, E)$ on $n = |V|$ nodes. Individuals initially have one of two private *beliefs* which we will refer to as "Correct" (or 1) and "Incorrect" (or 0). Formally, each $v \in V(G)$ receives an independent private signal $X(v) \in \{0, 1\}$, and $\Pr(X(v) = 1) = 1/2 + \delta$, for some universal constant $\delta \in (0, 1/2)$. Individuals also have a publicly announced *opinion* which we will simply refer to as an announcement. We define $C^t(v) \in \{\perp, 0, 1\}$ to be the public announcement of $v \in V$ at time t.

Initially, no announcement have been made, that is, $C^0(v) = \perp$ for all $v \in V$. In each subsequent step, a single node v^t is chosen uniformly at random from V, independently from the history of the process. In particular, as in the classical coupon collector problem, some nodes will be chosen many times before others will get lucky to get chosen for the first time. In step t, v^t updates her announcement using *majority dynamics*, while announcements of other nodes stay the same. To be specific, for any $i \in \{\perp, 0, 1\}$ and $v \in V$, let $N_i^t(v)$ denotes the number of neighbours of v that have opinion i at time t. Then,

$$C^t(v) = \begin{cases} 1 & \text{if } N_1^{t-1}(v) > N_0^{t-1}(v) \text{ and } v = v^t, \\ 0 & \text{if } N_1^{t-1}(v) < N_0^{t-1}(v) \text{ and } v = v^t, \\ X(v) & \text{if } N_1^{t-1}(v) = N_0^{t-1}(v) \text{ and } v = v^t, \\ C^{t-1}(v) & \text{if } v \neq v^t. \end{cases}$$

Finally, for any $i \in \{\perp, 0, 1\}$, let Y_i^t be the number of nodes that have opinion i at time t, that is, $Y_i^t = |\{v \in V : C^t(v) = i\}|$.

As shown in [18], it is easy to see that in any network this process stabilizes with high probability in $\mathcal{O}(n^2)$ steps. In fact, the process stabilizes in $O(n \log n + n \cdot d(G))$ where $d(G)$ is the diameter of the graph [3]. That is, the network reaches a state at some time T where no node will want to change its announcement and thus the process terminates. Our goal is to understand what fraction of nodes converges to the correct opinion, that is, what the value of Y_1^T/n is.

1.2 Our Results

The main contribution of this paper is the proof that the asynchronous majority dynamics on *binomial random graph* $\mathcal{G}(n,p)$ converges to the correct opinion, provided that the graph is sparse (that is, the average degree $np = o(n)$) and connected (that is, $np - \log n \gg 1$). If $np \gg \log n$, then the process converges to the correct opinion as quickly as it potentially could.

Theorem 1.1. *Let $\delta \in (0, 1/10]$. Let $\omega' = \omega'(n) = o(\log n)$ be any function that tends to infinity as $n \to \infty$. Suppose that $p = p(n) \ll 1$ and $p \gg \log n/n$, and consider the asynchronous majority dynamics on $\mathcal{G}(n,p)$.*

Then, asymptotically almost surely (a.a.s.) after $n(\log n + \omega') = (1 + o(1))n \log n$ rounds the process terminates with all nodes announcing the correct opinion. In fact, it happens exactly at time $\hat{T}$, where $\hat{T}$ is the first time when all nodes are selected at least once.

For sparser (but still connected) graphs, the process also converges to the correct opinion. In this case, we do not aim to show that it happens at time $\hat{T}$ and we only provide an upper bound for the number of rounds. It remains an open problem to determine if the process terminates at time $\hat{T}$ or it needs more time to converge.

Theorem 1.2. *Let $\delta \in (0, 1/10]$. Let $\omega' = \omega'(n) = o(\log n)$ be any function that tends to infinity as $n \to \infty$. Suppose that $p = p(n) \le \omega' \log n/n$ and $p \ge (\log n + \omega')/n$, and consider the asynchronous majority dynamics on $\mathcal{G}(n,p)$.*

Then, a.a.s. after $\mathcal{O}(n(\log n)^2/(\log\log n))$ rounds the process terminates with all nodes announcing the correct opinion.

These results are best possible in the following sense. If $p \le (\log n - \omega')/n$, then a.a.s. $\mathcal{G}(n,p)$ is disconnected. In fact, a.a.s. there are at least ω' isolated nodes which announce their own private believes. As a result, a.a.s. some nodes announce the correct opinion but some of them announce the incorrect one. Indeed, the probability that all isolated nodes converge to the same opinion is at most $o(1) + (1/2 + \delta/2)^{\omega'} + (1/2 - \delta/2)^{\omega'} = o(1)$. On the other hand, if $p \in (0, 1]$ is a constant separated from zero, then with positive probability the process converges to the correct opinion and with positive probability it converges to the incorrect opinion.

Theorem 1.3. *Let $\delta \in (0, 1/2)$. Let $\omega' = \omega'(n) = o(\log n)$ be any function that tends to infinity as $n \to \infty$. Suppose that $p \in (0, 1]$ is a constant, and consider the asynchronous majority dynamics on $\mathcal{G}(n,p)$.*

Then, the following is true for $i \in \{0, 1\}$: with probability at least p_i, after $n(\log n + \omega') = (1 + o(1))n \log n$ rounds the process terminates with all nodes announcing opinion i, where

$$\begin{aligned} p_1 &= (1/2 + \delta) \exp\Big(-\log(1/p)(1/p)\Big) > 0 \\ p_0 &= (1/2 - \delta) \exp\Big(-\log(1/p)(1/p)\Big) > 0. \end{aligned}$$

Finally, let us mention that for some technical reason, in Theorems 1.1 and 1.2 it is assumed that $\delta \le 1/10$. However, it is easy to couple the process with $\delta \le 1/10$ with the one with $\delta \in (1/10, 1/2)$ to show that the result holds for any $\delta \in (0, 1/2)$—see Subsection 2.3 for more details.

1.3 Future Directions

Let us highlight a few potential directions one might want to consider.

- As already mentioned above, for very sparse graphs ($np - \log n \to \infty$ and $np = \mathcal{O}(\log n)$), it would be interesting to determine if the process terminates at time $\hat{T}$ or it needs more time to converge to the correct opinion—see Theorem 1.2.
- Theorem 1.2 holds as long as $pn = \log n + \omega$ for some $\omega = \omega(n) \to \infty$ as $n \to \infty$. It is known that if $pn = \log n + c$ for some constant $c \in \mathbb{R}$, then with probability bounded away from one and from zero the graphs is disconnected. As a result, there is no hope to extend the result for this range of p. But it is plausible that a.a.s. it holds right at the time the random graph process creates a connected graph. This would be an optimal "hitting time" result.
- For disconnected graphs ($np - \log n \to -\infty$), it would be interesting to investigate the process run on the giant component of $\mathcal{G}(n, p)$.
- For dense graphs, it is not true that a.a.s. all nodes converge to the correct opinion—see Theorem 1.3. Having said that, it is reasonable to expect that a.a.s. all nodes converge to the same opinion (for example, [19] show that a consensus is reached in this case in a synchronous setting). Is is true in our asynchronous setting? In any case, what is the asymptotic value of the probability that all nodes converge to the correct opinion?

1.4 Related Work

In this section, we briefly discuss prior work on social learning mainly focusing on the setting with a binary state of the world and the agents initially have a correct opinion independently with probability $1/2 + \delta$. We refer to some recent surveys on social learning and opinion dynamics [33, 7, 10] for a more detailed literature review.

Majority dynamics falls under a wide class of naive or non-Bayesian models, where agents use a simple local heuristic to update their opinions, to capture simple behaviours exhibited by non-expert decision makers. Prior works have studied majority dynamics under a variety of modeling assumptions, to understand when a consensus is possible and when there is social learning—that is, the consensus (or the majority) is correct. These works study a variety of networks such as k-regular trees [25, 28, 3], bounded degree graphs [18], "symmetric" graphs and expanders [35]. In [40], a different perspective on social learning asks when is it possible to "recover the correct opinion" at the end of the dynamics through any function (not just a consensus or majority vote). Prior work has also considered models with different notions of bias towards correct opinion [2, 41].

Recently, there has been a series of work studying synchronous majority dynamics in binomial random graphs [9, 19, 14], with a focus to showing that 99% of the nodes converge to the same opinion (with high probability) for sparse random graphs, with $p = \Omega(\log n/n^{3/5})$ being the best known lower bound for the average degree. In contrast to these works, we focus on asynchronous dynamics and prove that a correct consensus is reached with high probability for $p = \Omega(\log n/n)$ and $p = o(1)$. Binomial random graphs are also studied under label propagation [30] which is a special case of synchronous majority dynamics with non-binary opinion in $[0, 1]$.

Many of the works mentioned above focus on synchronous updates, where all agents update their opinions synchronously in each round. Majority dynamics with synchronous updates leads to a correct consensus for all networks that are sufficiently connected [35], whereas with asynchronous updates the network structure can have a huge impact on social learning. This is best illustrated by the complete graph. With asynchronous updates, once the first agent announces their opinion (which can be wrong with probability $1/2 - \delta$) everyone will copy this. Hence, with a probability bounded away from zero all the nodes converge to the wrong opinion. In contrast, if all agents were to update synchronously, then the majority of the round one updates will be correct with high probability, so there will be a correct consensus in round two. Recent work [4], studies the DeGroot model with uninformed agents, to capture the different phases of information diffusion and social learning, which is a key phenomena that occurs in our asynchronous model.

Other non-Bayesian dynamics have also been extensively studied. In the Voter model, agents choose a random neighbour and copy their opinion [15, 24]. A similar dynamics called k-majority model are studied in the distributed computing literature, where agents choose k-neighbours at random and copy their majority [8, 22, 21, 16]. In the DeGroot Model, an agent's opinion lies in $[0, 1]$ (as opposed to binary $\{0, 1\}$) and agents update to the average of their neighbours [17, 23]. A key difference between these works and majority dynamics is that in these models a consensus is reached with probability 1 for any connected graphs. This is not the case in majority dynamics even with synchronous updates.

While our focus is in non-Bayesian dynamics, there has also been a long line of work studying Bayesian models, where agents update their beliefs rationally given their (local) observations exhibiting more sophisticated decision-making. Seminal works [6, 11] introduced the study of Bayesian dynamics and identified conditions that lead to information cascades. Here, the agents arrive sequentially and observe all the announcements (i.e., they form a complete graph), and many other subsequent works consider Bayesian dynamics under different assumptions and variations [39, 5, 13]. Bayesian dynamics in general social networks were first studied in [1]. There is also a long line of work studying Bayesian learning with repeated interactions [20, 38, 27, 37, 36, 34].

2 Preliminaries

2.1 Notation

Let us first precisely define the $\mathcal{G}(n, p)$ binomial random graph. $\mathcal{G}(n, p)$ is a distribution over the class of graphs with the set of nodes $[n] := \{1, \ldots, n\}$ in which every pair $\{i, j\} \in \binom{[n]}{2}$ appears independently as an edge in G with probability p. Note that $p = p(n)$ may (and usually does) tend to zero as n tends to infinity. We say that $\mathcal{G}(n, p)$ has some property *asymptotically almost surely* or a.a.s. if the probability that $\mathcal{G}(n, p)$ has this property tends to 1 as n goes to infinity. For more about this model see, for example, [12, 26, 29].

Given two functions $f = f(n)$ and $g = g(n)$, we will write $f(n) = \mathcal{O}(g(n))$ if there exists an absolute constant $c \in \mathbb{R}_+$ such that $|f(n)| \le c|g(n)|$ for all n, $f(n) = \Omega(g(n))$ if $g(n) = \mathcal{O}(f(n))$, $f(n) = \Theta(g(n))$ if $f(n) = \mathcal{O}(g(n))$ and $f(n) = \Omega(g(n))$, and we write $f(n) = o(g(n))$ or $f(n) \ll g(n)$ if $\lim_{n\to\infty} f(n)/g(n) = 0$. In addition, we write $f(n) \gg g(n)$ if $g(n) = o(f(n))$ and we write $f(n) \sim g(n)$ if $f(n) = (1 + o(1))g(n)$, that is, $\lim_{n\to\infty} f(n)/g(n) = 1$.

2.2 Concentration Tools

In this section, we state a few specific instances of Chernoff's bound that we will find useful. Let $(Z_1, \ldots, Z_n)$ be a sequence of independent Bernoulli(p) random variables. For each $j \in [n]$, let $X_j = \sum_{i=1}^{j} Z_i$. In particular, $X_n \in \mathrm{Bin}(n, p)$ is a random variable distributed according to a

Binomial distribution with parameters n and p. Then, a consequence of *Chernoff's bound* (see e.g. [26, Theorem 2.1]) is that for any $t \geq 0$ we have

$$\mathbb{P}(X_n - \mathbb{E}[X_n] \geq t) \leq \exp\left(-\frac{t^2}{2(\mathbb{E}[X_n] + t/3)}\right) \quad (1)$$

$$\mathbb{P}(\mathbb{E}[X_n] - X_n \geq t) \leq \exp\left(-\frac{t^2}{2\mathbb{E}[X_n]}\right). \quad (2)$$

Moreover, let us mention that the above bounds hold in a more general setting as well, that is, for any sequence $(Z_j)_{1\leq j\leq n}$ of independent random variables such that for every $j \in [n]$ we have $Z_j \in \text{Bernoulli}(p_j)$ with (possibly) different p_j-s (again, see e.g. [26] for more details).

Finally, we note that $X_n - \mathbb{E}[X_n]$ in (1) can be replaced with $\max_{1\leq j\leq n}(X_j - \mathbb{E}[X_j])$ and $\mathbb{E}[X_n] - X_n$ in (2) can be replaced with $\max_{1\leq j\leq n}(\mathbb{E}[X_j] - X_j)$. That is, we have

$$\mathbb{P}(\max_{1\leq j\leq n}(X_j - \mathbb{E}[X_j]) \geq t) \leq \exp\left(-\frac{t^2}{2(\mathbb{E}[X_n] + t/3)}\right) \quad (3)$$

$$\mathbb{P}(\max_{1\leq j\leq n}(\mathbb{E}[X_j] - X_j) \geq t) \leq \exp\left(-\frac{t^2}{2\mathbb{E}[X_n]}\right). \quad (4)$$

This is a consequence of a standard martingale bound (see e.g. [32] for more details).

2.3 Coupling

Suppose that at some point of the process, the public announcement is captured by $C^t(v)$, $v \in V$. Let $\hat{C}^t(v)$ be any sequence of opinions such that the following properties hold: (a) if $\hat{C}^t(v) = 1$, then $C^t(v) = 1$, (b) if $\hat{C}^t(v) = 0$, then $C^t(v) \in \{0, 1, \perp\}$, (c) if $\hat{C}^t(v) = \perp$, then $C^t(v) = \perp$. In other words, we get the auxiliary sequence $C^t(v)$ by modifying some of the opinions 1 and $\perp$ in $C^t(v)$ to 0. Hence, the process starting from $C^t(v)$ can be coupled with the auxiliary process starting from $\hat{C}^t(v)$ such that all the properties (a)–(c) are satisfied in every step of the process. In particular, if the auxiliary process converges to all nodes having opinion 1, then so does the original process. This easy observation will turn out to be useful in analyzing the process.

Similarly, suppose private beliefs in the auxiliary process are dominated by private beliefs in the original process: for any $v \in V$, $\hat{X}(v) \leq X(v)$. If the two processes are coupled, then properties (a)–(c) hold again. As before, if the auxiliary process converges to all nodes having opinion 1, then so does the original process. In particular, as mentioned above, the assumption that $\delta \in (0, 1/10]$ in Theorems 1.1 and 1.2 can be relaxed to $\delta \in (0, 1/2)$.

3 Sparse Random Graphs

In this section, we consider sparse random graphs, that is, we will assume that $p = o(1)$. Let $\omega = \omega(n)$ be a function that tends to infinity as $n \to \infty$, arbitrarily slowly. In particular, each time we refer to ω, we will assume that $\omega \ll pn$ and $\omega \ll (1/p)^{1/2}$ so that $1/p \gg 1/(p\omega) \gg 1/(p\omega^2) \gg 1$.

We will consider a few phases. During the first phase (Subsection 3.1), most of the nodes that are chosen have not yet announced their opinions ($C^{t-1}(v^t) = \perp$) and none of their neighbours have announced ($N_1^{t-1}(v^t) = N_0^{t-1}(v^t) = 0$). Hence, the announcement of v^t will typically coincide with its private belief. Moreover, most of the nodes selected will not be chosen again during this phase.

During the second phase (Subsection 3.2), it is still the case that most selected nodes are selected for the first time but this time they might have neighbours that announced their opinions. As a result, the argument is more involved but the conclusion is that at the end of the second phase more nodes have correct opinion than not.

The analysis of the first two phases can be applied for all sparse graphs, even below the threshold for connectivity. The analysis of the final steps of the process is slightly more involved. We first present an easy argument for not very sparse graphs (Subsection 3.3), that is, when the asymptotic expected degree degree satisfies $pn \gg \log n$. Very sparse graphs for which $pn = \Theta(\log n)$ (but, of course, above the connectivity threshold) are considered in Subsection 3.4.

Overview. A key phenomena in asynchronous dynamics is that the process involves both information diffusion and conventional social learning. Intuitively, the process initially produces some independent beliefs/opinions pop up sporadically throughout the network. These opinions then diffuses in the network during the process as more nodes are selected to announce/update their opinion by learning from their neighbours. With this in mind, our analysis considers multiple phases of the process. We provide a brief description of the different phases below.

- **Phase 1.** In the first few time steps, most nodes that are selected to announce have not been selected earlier and, more importantly, do not have neighbours who have been selected before. So almost all of the opinions in the network at the end of phase one are just the independent private beliefs of the selected nodes. Since the private signals are biased towards being correct, a strict majority of the opinions are correct at the end of the first phase. In particular, we show that at time $T_1 = \delta/2p$, the number of nodes with opinion 1 is at least $(1/2 + 3\delta/5)T_1$ and opinion 0 is represented at most $(1/2 - 3\delta/5)T_1$ times. Moreover, $T_1(1 - o(1))$ many nodes have made some announcement in this phase, that is, very few nodes were selected more than once.

- **Phase 2.** In the second phase, again most nodes that are selected to announce have not been selected earlier. In particular, we show that at any time t during the second phase (i.e., after time T_1 but before time T_2), the number of nodes that were selected twice before time t is $o(t)$. Moreover, since a *super majority* of the opinions at the end of the previous phase were correct, we prove that nodes that are selected to announce for the first time are more likely to learn the correct opinion even if we pretend that the few nodes that are selected *again* were to change their opinion to 0.

- **Phase 3 (a).** For not very sparse graphs, we are able to show all nodes which were not selected in the first two phases have more neighbours with opinion 1 than not. Again, very few nodes who were selected before are selected again before time T_3, so even if all of them announce 0 all nodes who make their first announcement between time T_2 and time T_3 announce the correct opinion. Finally, even if all the nodes that were selected before time T_2 are to have opinion 0 and all nodes that were selected for the first time between time T_2 and time T_3 have opinion 1, we show that a.a.s. all announcements after time T_3 are always correct.

- **Phase 3 (b).** For very sparse graphs, the proof of the last phase is more involved as there might be nodes whose degree is too small to guarantee that a majority of their neighbours have opinion 1, even though there is a super majority of opinion 1 in the graphs. However, we may bound the number of nodes with small degrees and show that no large degree node has more than one small degree neighbour. With this in hand, we show that after every batch of $O(n \log n)$ many time steps the number of large degree nodes with opinion 0 shrinks by at least $(\log \log n)^{1/4}$ factor. Hence, after $o(\log n)$ many such batches all large nodes have opinion 1. Finally, we show that no two small degree nodes are adjacent to each other, and hence all the small degree nodes will also switch to opinion 1 by copying the opinions of their large degree neighbours.

We highlight a few simple techniques that help us in the analysis. Firstly, separating the randomness of the graph, the node selection process and the opinion formation. For example, we wait to reveal the edges adjacent to a node only when she is selected to announce for the first time. Second, considering an auxiliary dynamics that is coupled with the actual dynamics in order to ignore problematic but rare events such as the repeated nodes in the first two phases. Finally, finding independent sequences of random variables that stochastically dominate the opinion dynamics sequence in order to compute probability bounds more easily.

3.1 Phase 1: $T_1 = \delta/(2p)$

In the analysis of the process, it will be convenient to ignore opinions of a small fraction of nodes, and consider the following *auxiliary dynamics*. We will use $D^t(v) \in \{\perp, ?, 0, 1\}$ to denote the *auxiliary announcement* of $v \in V$ at time t. For any $i \in \{\perp, ?, 0, 1\}$, let Z_i^t be the number of nodes that have auxiliary opinion i at time t, that is, $Z_i^t = |\{v \in V : D^t(v) = i\}|$. We will explain how the values of $D^t(v)$ are determined soon but the auxiliary dynamics will be coupled with the original one and, in particular, we will make sure that the following property holds.

Property 3.1. *If $D^t(v) = i$ for some $i \in \{\perp, 0, 1\}$ and time t, then $C^t(v) = D^t(v)$. On the other hand, if $D^t(v) = ?$, then $C^t(v) \in \{0, 1\}$. As a result, for $i \in \{0, 1\}$ and any time t during the first phase, we have*

$$Z_i^t \le Y_i^t \le Z_i^t + Z_?^t. \tag{5}$$

The first phase takes $T_1 = \delta/(2p) = \Theta(1/p) \gg \omega^2 \gg 1$ rounds. In order to keep the analysis easy, we postpone exposing edges of $\mathcal{G}(n,p)$ for as long as possible, and keep the following useful property.

Property 3.2. *At any time t, only edges of $\mathcal{G}(n,p)$ with both endpoints in the set $\{v : D^t(v) \neq \perp\}$ are exposed.*

The auxiliary dynamics, coupled with the original one, that we aim to understand is defined as follows. Consider a node v^t chosen at time t. For all other nodes $v \neq v^t$ we have $D^t(v) = D^{t-1}(v)$. For v^t we have,

$$D^t(v^t) = \begin{cases} ? & \text{if } D^{t-1}(v^t) \neq \perp, \\ ? & \text{if } \exists \text{ node } v \text{ such that } v \in N(v^t) \text{ and } D^{t-1}(v) \neq \perp, \\ X(v) & \text{otherwise.} \end{cases}$$

That is, if v^t had announced her opinion at least once before time t ($D^{t-1}(v^t), C^{t-1}(v^t) \neq \perp$), then we fix $D^t(v^t) = ?$. On the other had, if v^t has not announced her opinion yet (that is, $D^{t-1}(v^t) = C^{t-1}(v^t) = \perp$), then we expose edges of $\mathcal{G}(n,p)$ between v^t and the set $\{v : D^{t-1}(v) \neq \perp\}$. If no edge between v^t and the set $\{v : D^{t-1}(v) \neq \perp\}$ is present, then no neighbour of v^t has an announced opinion and so $D^t(v^t) = C^t(v^t) = X(v^t)$ is fixed to the private belief of v^t. Otherwise (that is, at least one edge is present), then we simply fix $D^t(v^t) = ?$. Let us note that, alternatively, we could try to investigate the value of $C^t(v^t)$ and then fix $D^t(v^t) = C^t(v^t)$ but we expect at most $pt \le pT_1 = \delta/2$ edges between v^t and $\{v : D^{t-1}(v) \neq \perp\}$, and so there will not be very many nodes v^t of this type. As a result, we may simply ignore the announcements of such nodes.

Moreover, a useful implication of this approach is that in order to estimate the values of $Z^t_\perp$ and $Z^t_?$ in this process, we do not need to uncover nodes' private believes ($X(v)$'s). Hence, we may postpone exposing private beliefs of nodes with $D^t(v) \notin \{\perp, ?\}$ to the very end of this phase, and only then expose this information to determine how many nodes satisfy $D^{T_1}(v) = 1$ and how many of them satisfy $D^{T_1}(v) = 0$. Finally, it is easy to see that Property 3.1 is satisfied at time T_1 and Property 3.2 is satisfied in any point of the process.

Here is the main result of this subsection.

Theorem 3.3. *Suppose that $p = p(n) \ll 1$ and $p \gg 1/n$. Set $T_1 = \delta/(2p)$. Let $\omega = \omega(n) \ll \min\{pn, (1/p)^{1/2}\}$ be any function that tends to infinity as $n \to \infty$. Then, a.a.s. the following holds:*

$$Z^{T_1}_? \le \frac{\delta T_1}{4}\left(1 + \mathcal{O}(1/\omega)\right) \tag{6}$$

$$Z^{T_1}_1 \ge (1/2 + 3\delta/5)\ T_1 \tag{7}$$

$$Z^{T_1}_? + Z^{T_1}_1 + Z^{T_1}_0 = T_1\left(1 - \mathcal{O}(1/\omega)\right). \tag{8}$$

As a result, by Property 3.1,

$$\begin{aligned} Y^{T_1}_1 &\ge (1/2 + 3\delta/5)\ T_1 \\ Y^{T_1}_0 &\le (1/2 - 3\delta/5)\ T_1 \\ Y^{T_1}_1 + Y^{T_1}_0 &= T_1\left(1 + \mathcal{O}(1/\omega)\right). \end{aligned}$$

Proof. Let us start with investigating $Z^{T_1}_?$. Recall that in our auxiliary dynamics, there are two ways node v^t could change its state to $D^t(v^t) = ?$ at time t. Let I_t be the indicator random variable that this happens because $D^{t-1}(v^t) \neq \perp$, and let $I = \sum_{t=1}^{T_1} I_t$. Similarly, let J_t be the indicator random variable that $D^{t-1}(v^t) = \perp$ but there is an edge between v^t and the set $\{v : D^{t-1}(v) \neq \perp\}$. Let $J = \sum_{t=1}^{T_1} J_t$.

Note that, at most $t-1$ distinct nodes have made an announcement before round t. In particular, at most one node can change its state from $D^{t-1}(v) = \perp$ to $D^t(v) \neq \perp$, deterministically, at any round t of the process. So, the number of nodes with $D^{t-1}(v) \neq \perp$ is $n - Z^{t-1}_\perp \le t - 1$. We get that

$$\Pr(I_t = 1) = \frac{n - Z^{t-1}_\perp}{n} \le \frac{t-1}{n},$$

and so I can be stochastically upper bound by $\hat{I} = \sum_{t=1}^{T_1} \hat{I}_t$ where $(\hat{I}_t)_{1 \le t \le T_1}$ are independent variables and for every $t \in [T_1]$ we have $\hat{I}_t \in \text{Bernoulli}((t-1)/n)$. Note that, since $pn \gg \omega$,

$$\mathbb{E}[\hat{I}] = \sum_{t=1}^{T_1} \frac{t-1}{n} = \frac{(T_1 - 1)T_1}{2n} \sim \frac{\delta T_1}{4pn} \ll \frac{T_1}{\omega}. \tag{9}$$

It follows from Chernoff's bound (1) (and the comment right after it) applied with $t = T_1/\omega = \Theta(1/(p\omega)) \gg \omega \gg 1$ that

$$\Pr(\hat{I} \ge \mathbb{E}[\hat{I}] + t) \le \exp\left(-\frac{t^2}{(2/3 + o(1))t}\right) = \exp\left(-\Theta(t)\right) = o(1).$$

So a.a.s. $I \le \hat{I} = \mathcal{O}(T_1/\omega)$. Similarly, since $pt \le pT_1 = \delta/2 < 1/4$,

$$\Pr(J_t = 1) = \frac{Z_{\perp}^{t-1}}{n}\left(1 - (1-p)^{n - Z_{\perp}^{t-1}}\right) \le 1 - (1-p)^t = 1 - \left(1 - pt + p^2\binom{t}{2} - \dots\right) \le pt.$$

As before, we stochastically upper bound J by $\hat{J} = \sum_{t=1}^{T_1} \hat{J}_t$, where $\hat{J}_t \in \text{Bernoulli}(pt)$. We get that

$$\mathbb{E}[\hat{J}] = \sum_{t=1}^{T_1} pt = \frac{p(T_1+1)T_1}{2} = \frac{pT_1^2}{2}\left(1 + \mathcal{O}(1/T_1)\right) = \frac{\delta T_1}{4}\left(1 + \mathcal{O}(1/\omega)\right),$$

and Chernoff's bound (1) (applied with $t = \mathbb{E}[\hat{J}]/\omega$) implies that

$$\Pr(\hat{J} \ge \mathbb{E}[\hat{J}] + t) \le \exp\left(-\frac{\mathbb{E}[\hat{J}]}{(2+o(1))\omega^2}\right) = \exp\left(-\Theta(T_1/\omega^2)\right) = \exp\left(-\Theta(1/(p\omega^2))\right) = o(1).$$

Hence, a.a.s. $J \le \hat{J} \le \frac{\delta T_1}{4}\left(1 + \mathcal{O}(1/\omega)\right)$ and so a.a.s. $Z_?^{T_1} \le I + J \le \frac{\delta T_1}{4}\left(1 + \mathcal{O}(1/\omega)\right)$. This proves (6).

It remains to investigate $Z_0^{T_1}$ and $Z_1^{T_1}$. Let us summarize the situation at time T_1. The number of rounds when nodes were not chosen for the first time is at most $I = \mathcal{O}(T_1/\omega)$ a.a.s. Hence, a.a.s. the number of nodes that were chosen at least once is equal to $T_1 - \mathcal{O}(T_1/\omega)$. This proves (8). Moreover, it implies that a.a.s. the number of nodes with $D^{T_1}(v) \notin \{\perp, ?\}$ is equal to

$$Z_1^{T_1} + Z_0^{T_1} = T_1 - \mathcal{O}(T_1/\omega) - Z_?^{T_1} \ge (1 - \delta/4)T_1\left(1 + \mathcal{O}(1/\omega)\right).$$

More importantly, as mentioned above, in the analysis so far we did not use their opinions which are consistent with their private beliefs. We conveniently deferred this information up to now. After exposing this information, we get that $Z_1^{T_1}$ is stochastically lower bounded by the random variable $\hat{Z}_1 \in \text{Bin}((1-\delta/4)T_1 - cT_1/\omega, 1/2 + \delta)$, where $c > 0$ is a large enough constant. After applying Chernoff's bound (2) (with $t = T_1/\omega$) we get that

$$\begin{aligned} Z_1^{T_1} \ge \hat{Z}_1 &= (1/2+\delta)(1-\delta/4)T_1(1+\mathcal{O}(1/\omega)) \\ &\ge (1/2+\delta-\delta/4)T_1(1+\mathcal{O}(1/\omega)) \\ &\ge (1/2+3\delta/5)T_1 \end{aligned}$$

with probability at least

$$1 - \exp(-\Theta(T_1/\omega^2)) = 1 - \exp(-\Theta(1/(p\omega^2))) = 1 - o(1).$$

This proves (7).

The conclusion for $Y_1^{T_1}$ follows immediately from Property 3.1, and the bound for $Y_0^{T_1}$ is a trivial implication of the fact that $Y_1^{T_1} + Y_0^{T_1} \le T_1$. The proof of the theorem is finished. □

3.2 Phase 2: $T_2 = T_2(n)$ such that $\omega/p \le T_2 \le n/\omega$

By Theorem 3.3, since we aim for a statement that holds a.a.s., we may assume that at the beginning of Phase 2,

$$\begin{aligned} Y_1^{T_1} &\ge (1/2+3\delta/5)\,T_1 \\ Y_0^{T_1} &\le (1/2-3\delta/5)\,T_1 \\ Y_1^{T_1} + Y_0^{T_1} &= T_1\left(1+\mathcal{O}(1/\omega)\right). \end{aligned}$$

As in the previous phase, it will be convenient to ignore opinions of some problematic nodes and assign auxiliary announcements $D^t(v) =?$ to such nodes. We will continue using Z_i^t to denote the number of nodes that have auxiliary opinion i at time t. We fix $D^{T_1}(v) = C^{T_1}(v)$ for all v so, initially, auxiliary announcements coincide with the truth announcements. However, this time we assign $D^t(v^t) =?$ only if $D^{t-1}(v^t) \neq \perp$ (that is, the node chosen at time t has made an announcement in the past); otherwise, the auxiliary announcement $D^t(v^t)$ is determined immediately pretending that all neighbours v of v^t with $D^{t-1}(v) =?$ announced 0. More formally, for each node v and $i \in \{0, 1, \perp, ?\}$ let $\hat{N}_i^t(v)$ denote the number of neighbours v' of v with auxiliary opinion $D^t(v') = i$ at time t. Then we have,

$$D^t(v^t) = \begin{cases} ? & \text{if } D^{t-1}(v^t) \neq \perp, \\ 1 & \text{if } \hat{N}_1^{t-1}(v^t) > \hat{N}_0^{t-1}(v^t) + \hat{N}_?^{t-1}(v^t), \\ 0 & \text{if } \hat{N}_1^{t-1}(v^t) < \hat{N}_0^{t-1}(v^t) + \hat{N}_?^{t-1}(v^t), \\ X(v) & \text{if } \hat{N}_1^{t-1}(v^t) = \hat{N}_0^{t-1}(v^t) + \hat{N}_?^{t-1}(v^t). \end{cases}$$

As a consequence, $D^t(v)$ and $C^t(v)$ are coupled so that the following property is satisfied.

Property 3.4. *If $D^t(v) = i$ for some $i \in \{\perp, 1\}$ and time t, then $C^t(v) = D^t(v)$. On the other hand, if $D^t(v) = i$ for some $i \in \{0, ?\}$, then $C^t(v) \in \{0, 1\}$. As a result, for any time t during the second phase, we have $Y_1^t \geq Z_1^t$.*

As before, it is easy to see that Property 3.2 is also satisfied during this phase. Here is the main result of this subsection.

Theorem 3.5. *Suppose that $p = p(n) \ll 1$ and $p \gg 1/n$. Let $\omega = \omega(n) \ll \min\{pn, (1/p)^{1/2}\}$ be any function that tends to infinity as $n \to \infty$. Set $T_2 = T_2(n)$ such that $\omega/p \leq T_2 \leq n/\omega$. Then, a.a.s. the following holds:*

$$\begin{aligned} Z_?^{T_2} &= \mathcal{O}(T_2/\omega) \\ Z_1^{T_2} &\geq (1/2 + \delta/2)\ T_2 \\ Z_?^{T_2} + Z_1^{T_2} + Z_0^{T_2} &= T_2\,(1 + \mathcal{O}(1/\omega))\,. \end{aligned}$$

As a result, by Property 3.4,

$$\begin{aligned} Y_1^{T_2} &\geq (1/2 + \delta/2)\ T_2 \\ Y_0^{T_2} &\leq (1/2 - \delta/2)\ T_2 \\ Y_1^{T_2} + Y_0^{T_2} &= T_2\,(1 + \mathcal{O}(1/\omega))\,. \end{aligned}$$

Before we move to the proof of this theorem, let us make some simple but useful observations. First, note that only a negligible fraction of the nodes have opinion that we do not control.

Lemma 3.6. *Suppose that $p = p(n) \ll 1$ and $p \gg 1/n$. Let $\omega = \omega(n) \ll \min\{pn, (1/p)^{1/2}\}$ be any function that tends to infinity as $n \to \infty$. Set $T_2 = T_2(n)$ such that $\omega/p \leq T_2 \leq n/\omega$. Then, a.a.s., for any t such that $T_1 \leq t \leq T_2$, $Z_?^t \leq 2t/\omega$.*

Proof. In fact, we will prove something stronger. Let X_t be the number of nodes that were selected at least two times up to time t (which could happen before time T_1). We will prove that a.a.s. for any $1 \leq t \leq n/\omega$, $X_t \leq 2t/\omega$.

Case 1: $1 \le t \le n^{2/5}$. As argued in the proof of Theorem 3.3 (see (9)), one can bound the expected value of $X_{n^{2/5}}$ as follows:

$$\mathbb{E}[X_{n^{2/5}}] \le \sum_{t \le n^{2/5}} \frac{t-1}{n} \sim \frac{n^{4/5}}{2n} = o(1).$$

Since X_t is non-decreasing, it follows immediately from Markov's inequality that a.a.s. $X_t \le X_{n^{2/5}} = 0$ for all t such that $1 \le t \le n^{2/5}$.

Case 2: $n^{2/5} \le t \le n^{3/5}$. As before, we observe that $\mathbb{E}[X_{n^{3/5}}] = \mathcal{O}(n^{6/5}/n) = \mathcal{O}(n^{1/5})$ and so using Markov's inequality again we get that a.a.s. $X_t \le X_{n^{3/5}} \le \mathbb{E}[X_{n^{3/5}}] \log n = \mathcal{O}(n^{1/5} \log n) \le 2t/\omega$ for all t such that $n^{1/3} \le t \le n^{2/3}$.

Case 3: $n^{3/5} \le t \le n/\omega$. Fix any t in this range. This time we stochastically upper bound X_t by $X'_t \in \mathrm{Bin}(t, t/n)$ with $\mathbb{E}[X'_t] = t^2/n \ge n^{1/5} \gg \log n$. Chernoff's bound (1) implies that $X_t \le X'_t \le 2\mathbb{E}[X'_t] = 2t^2/n \le 2t/\omega$ with probability $1 - \mathcal{O}(n^{-2})$. The desired result holds by the union bound over all t in this range. □

Let us fix $k \in \mathbb{N}$ and consider random variable $X_k \in \mathrm{Bin}(k, 1/2+\delta/2)$. We will need to understand the following sequence of constants (the connection to our problem will become clear soon):

$$q_k := \mathbb{P}(X_k > k/2) + \mathbb{P}(X_k = k/2) \cdot (1/2 + \delta). \tag{10}$$

Clearly, $q_0 = 1/2 + \delta$ and $q_1 = 1/2 + \delta/2$. For any other value of $k \ge 2$, $q_k \ge 1/2 + 51\delta/100$ as we show in the next technical lemma. The proof of this fact can be found in the appendix.

Lemma 3.7. *Fix $k \in \mathbb{N}$ such that $k \ge 2$, and $\delta \in (0, 1/10]$. Then,*

$$q_k \ge \frac{1}{2} + \frac{51}{100}\delta.$$

Now, we are ready to go back to analyzing the behaviour of the process during the second phase.

Proof of Theorem 3.5. Our goal is to show that a.a.s. the following inequalities hold for any t such that $T_1 \le t \le T_2$:

$$\frac{Z_1^t}{Z_0^t + Z_?^t} \ \ge\ \frac{1/2+\delta/2}{1/2-\delta/2} \tag{11}$$

$$Z_?^t \ \le\ 2t/\omega. \tag{12}$$

Formally, we define the stopping time S to be the minimum value of $t \ge T_1$ such that either (11) fails, (12) fails or $t = T_2$. (A stopping time is any random variable S with values in $\{T_1, T_1+1, \ldots, T_2\}$ such that, for any time $\hat{t}$, it is determined whether $S = \hat{t}$ from knowledge of the process up to and including time $\hat{t}$.)

Property (12) is trivially satisfied at the beginning of the second phase as $Z_?^{T_1} = 0$. By Theorem 3.3, since we aim for a statement that holds a.a.s., we may assume that (11) is satisfied at the beginning of the second phase. In fact,

$$\frac{Z_1^{T_1}}{Z_0^{T_1} + Z_?^{T_1}} = \frac{Y_1^{T_1}}{Y_0^{T_1} + 0} \ge \frac{1/2 + 101\delta/200}{1/2 - 101\delta/200} \ge \frac{1/2+\delta/2}{1/2-\delta/2}.$$

It will be convenient to define $Z^t = Z_1^t + Z_0^t + Z_?^t$; that is, Z^t is the number of nodes that announced their opinions by time t. If (12) is satisfied, then only a negligible fraction of nodes were selected more than once and we get that $Z^t = t(1 - \mathcal{O}(1/\omega)) \sim t$.

Let us first show that if (11) and (12) are satisfied at time t and the node selected at time $t+1$ was not selected before (that is, $D^t(v^{t+1}) = C^t(v^{t+1}) = \perp$), then the probability that v^{t+1} announces an auxiliary opinion 1 is at least $1/2 + 101/200\delta$.

We first expose edges from v^{t+1} to the set $\{v : D^t(v) \neq \perp\}$ (see Property 3.2) and let us define p_k to be the probability that v^{t+1} has precisely k neighbours in that set. In particular, we have

$$\begin{aligned} p_1 \quad &= Z^t p(1-p)^{Z^t - 1} \quad = \lambda(1-p)^{\lambda/p - 1} \\ &\qquad\qquad\qquad\quad \le \lambda e^{-\lambda}/(1-p) \\ &\qquad\qquad\qquad\quad \le 1/e + o(1) \ < \ 1/2, \end{aligned} \tag{13}$$

where $\lambda = pZ^t = pt(1 - \mathcal{O}(1/\omega))$ and the second inequality follows because $xe^{-x} \le e^{-1}$ and $p = o(1)$.

Now, condition on v^t having exactly k neighbours that already announced their opinion. Note that we did not expose the neighbours yet (only the number of them) so neighbours form a random set of cardinality k from the set $\{v : D^{t-1}(v) \neq \perp\}$. Let r_k to be the probability that v^t announces auxiliary opinion 1 in this conditional probability space. It happens if more than $k/2$ neighbours of v^t have $D^{t-1}(v) = 1$. Moreover, if exactly $k/2$ neighbours have this property, then v^t announces opinion 1 with probability $1/2 + \delta$, which is the probability that its private belief is 1. Since (11) holds, r_k can be lower bounded by q_k which we defined in (10). It follows that the probability that v^t announces 1 is asymptotic to

$$\begin{aligned} \sum_{k \ge 0} r_k \cdot p_k \quad &\ge \quad \sum_{k \ge 0} q_k \cdot p_k = q_1 p_1 + \sum_{k \ge 0, k \neq 1} q_k \cdot p_k \\ &\ge \quad \left(\frac{1}{2} + \frac{\delta}{2}\right) p_1 + \left(\frac{1}{2} + \frac{51}{100}\delta\right)(1 - p_1) \\ &= \quad \left(\frac{1}{2} + \frac{51}{100}\delta\right) - p_1 \left(\frac{1}{100}\delta\right) \\ &\ge \quad \frac{1}{2} + \frac{101}{200}\delta, \end{aligned} \tag{14}$$

where the second inequality follows from Lemma 3.7 and the last one from (13).

Let s be the number of rounds t in the second phase in which v^t was not selected before, i.e., $D^{t-1}(v^t) = \perp$, and let $t_1, t_2, \ldots, t_s$ denote such round. Clearly, $s \le T_2 - T_1 = T_2(1 - \mathcal{O}(1/\omega))$ but, in fact, a.a.s. we have $s = T_2(1 - \mathcal{O}(1/\omega))$ by Lemma 3.6. For $i \in [s]$, let L_i be the indicator random variable for the event that v^{t_i} announced an auxiliary opinion 1, that is, $L_i = Z_1^{t_i} - Z_1^{t_i - 1}$. If both (11) and (12) hold at time $t_i - 1$, then $\mathbb{P}(L_i = 1) \ge 1/2 + 101\delta/200$ but, of course, we cannot condition on these two properties to hold. Instead, we will use a small trick and consider an auxiliary sequence of random variables after the stopping time S when one of the properties fails.

Fix $\hat{p} = 1/2 + 101\delta/200$ and let $M_1, \ldots, M_s$ be a sequence of independent Bernoulli variables with parameter $\hat{p}$. For each $i \in [s]$, we define $L'_i = L_i$ if both (11) and (12) hold at times $t < t_i$ and otherwise $L'_i = M_i$. That is, the process "stops" at our stopping time S which, in our context, means that it simply follows part of the sequence $(M_i)_{i=1}^s$ (namely, $(M_i)_{i=S+1}^s$) from that point on, ignoring the behaviour of the original process. Thus, defining $L'_{\le j} = \sum_{i=1}^j L'_i$ and $M_{\le j} = \sum_{i=1}^j M_i$, (14) implies that one can couple $L'_{\le j}$ and $M_{\le j}$ such that $L'_{\le j} \ge M_{\le j}$ for all $j \in [s]$.

Note that $\mathbb{E}[M_{\le j}] = \hat{p}j = (1/2 + 101\delta/200)j$ for any $j \in [s]$. If follows from Chernoff's bound (4),

$$\begin{aligned}
&\mathbb{P}\left(\exists_{1\le j\le s}\ \mathbb{E}[M_{\le j}] - M_{\le j} \ge \frac{\delta}{400}(T_1 + j)\right)\\
&\le \sum_{a\ge 1} \mathbb{P}\left(\max_{(2^{a-1}-1)T_1 < j \le (2^a-1)T_1} \left(\mathbb{E}[M_{\le j}] - M_{\le j}\right) \ge 2^{a-1}\frac{\delta}{400}T_1\right)\\
&\le \sum_{a\ge 1} \exp\left(-\Theta(2^a T_1)\right) = \exp\left(-\Theta(T_1)\right) = o(1),
\end{aligned}$$

since $T_1 = \Theta(1/p) \to \infty$. In other words, a.a.s. for any $j \in [s]$,

$$L'_{\le j} \ge M_{\le j} \ge \left(\frac{1}{2} + \frac{101}{200}\delta\right) j - \frac{\delta}{400}(T_1 + j).$$

Since $L_{\le j} = L'_{\le j}$ for any $j \in [s]$ such that $t_j < S$, and Z_1^t can decrease by at most one in a single round, a.a.s.

$$\begin{aligned}
Z_1^S &\ge Z_1^{S-1} - 1\\
&\ge Z_1^{T_1} + L_{\le S - T_1 - \mathcal{O}(S/\omega)} - \mathcal{O}(S/\omega)\\
&\ge \left(\frac{1}{2} + \frac{101}{200}\delta\right) T_1 + \left(\frac{1}{2} + \frac{101}{200}\delta\right)(S - T_1 - \mathcal{O}(S/\omega)) - \frac{\delta}{400}(S - \mathcal{O}(S/\omega)) - \mathcal{O}(S/\omega)\\
&\ge \left(\frac{1}{2} + \frac{201}{400}\delta\right) S - \mathcal{O}(S/\omega)\\
&\ge \left(\frac{1}{2} + \frac{\delta}{2}\right) S,
\end{aligned}$$

implying that (11) holds at time S. Indeed, there were $Z_1^{T_1}$ nodes with auxiliary opinion 1 at the beginning of the second phase. By Lemma 3.6, $S - T_1 - \mathcal{O}(S/\omega)$ nodes were selected for the first time before the stopping time and $L_{\le S - T_1 - \mathcal{O}(S/\omega)}$ of them announced 1 at that time. Finally, at most $\mathcal{O}(S/\omega)$ nodes that already announced their opinion were selected again. It implies that a.a.s. the process does not "stop" because of (11) failing. By Lemma 3.6, a.a.s. it also does not stop because of (12). Hence, a.a.s. $S = T_2$ and the proof of the theorem is finished. □

3.3 Not Very Sparse Random Graphs

In this subsection, we provide a relatively easy argument that works for random graphs with $pn \gg \log n$. In particular, we show that a.a.s. after round T_2 but before round $T_3 = n/\sqrt{\omega}$ all nodes that are selected for the first time announce 1. Moreover, after round T_3 every node selected announces 1 a.a.s.

Proof of Theorem 1.1. Let $\omega = \omega(n) \ll \min\{(pn/\log n)^{1/2}, pn, (1/p)^{1/2}\}$ be any function that tends to infinity as $n \to \infty$. In particular, $pn \ge \omega^2 \log n$. Fix $T_2 = T_2(n) = n/\omega$. It follows from Theorem 3.5 that a.a.s. at the end of the second phase, there are $Y_1^{T_2} \ge (1/2 + \delta/2)T_2$ nodes that announced opinion 1, and so $Y_0^{T_2} \le (1/2 - \delta/2)T_2$ nodes announced opinion 0; moreover, $Y_1^{T_2} + Y_0^{T_2} = T_2(1 + \mathcal{O}(1/\omega))$.

Let $V_i = \{v : C^t(v) = i\}$ be the set of nodes with opinion $i \in \{0, 1\}$ at time T_2. Note that, by Property 3.2, we may assume that only edges within $V_0 \cup V_1$ are exposed at that stage of the process. We will first show that a.a.s. all nodes $v \notin V_0 \cup V_1$ have substantially more neighbours in V_1 than in V_0. Indeed, this is a simple consequence of the Chernoff bounds (1) and (2): for any $v \notin V_0 \cup V_1$:

$$\begin{aligned}
\mathbb{P}\Big(|N(v) \cap V_1| \le |N(v) \cap V_0| + \delta T_2 p/2\Big) \\
&\le \mathbb{P}\Big(|N(v) \cap V_1| \le (1/2 + \delta/4) T_2 p\Big) + \mathbb{P}\Big(|N(v) \cap V_0| \ge (1/2 - \delta/4) T_2 p\Big) \\
&= \mathbb{P}\Big(\mathrm{Bin}(|V_1|, p) \le (1/2 + \delta/4) T_2 p\Big) + \mathbb{P}\Big(\mathrm{Bin}(|V_0|, p) \ge (1/2 - \delta/4) T_2 p\Big) \\
&\le 2 \exp\Big(-\Theta(T_2 p)\Big) \\
&= 2 \exp\left(-\Omega\left(\frac{n}{\omega} \cdot \frac{\omega^2 \log n}{n}\right)\right) \\
&= \mathcal{O}(1/n^2),
\end{aligned}$$

since $\mathbb{E}[\mathrm{Bin}(|V_1|, p)] \ge (1/2 + \delta/2) T_2 p$ and $\mathbb{E}[\mathrm{Bin}(|V_0|, p)] \le (1/2 - \delta/2) T_2 p$. The desired property holds by the union bound over all nodes $v \notin V_0 \cup V_1$.

Fix $T_3 = T_3(n) = n/\sqrt{\omega}$. The third phase will last till time T_3. Let $V_1' \subseteq V_1$ be the set of nodes from V_1 that were selected during the third phase. Note that each node from V_1 is selected during the third phase with probability at most $(T_3 - T_2)/n \le 1/\sqrt{\omega}$. Hence, $\mathbb{E}[|V_1'|] \le |V_1|/\sqrt{\omega}$ and so a.a.s. $|V_1'| \le |V_1|/\omega^{1/3}$ by Markov's inequality. A simple but important observation is that V_1' is determined exclusively by the selection process (coupon collector process); in particular, it does not depend on the random graph nor the opinion dynamics. Hence, we can use Chefnoff's bound again to show that a.a.s. all nodes $v \notin V_0 \cup V_1$ have very few neighbours in V_1'. Indeed, note that for any $v \notin V_0 \cup V_1$, the number of neighbours of n in V_1' can be stochastically upper bounded by $\mathrm{Bin}(|V_1|/\omega^{1/3}, p)$ with expectation $|V_1|p/\omega^{1/3} = \Theta(np/\omega^{4/3}) = \Omega(n^{2/3} \log n) \gg \log n$. Hence, $|N(v) \cap V_1'| = \mathcal{O}(|V_1|p/\omega^{1/3}) = \mathcal{O}(T_2 p/\omega^{1/3}) = o(T_2 p)$ with probability $1 - \mathcal{O}(1/n^2)$, and so a.a.s. all nodes $v \notin V_0 \cup V_1$ satisfy this property.

Combining the two properties together, we get that a.a.s. for all nodes $v \notin V_0 \cup V_1$ we have

$$|N(v) \cap (V_1 \setminus V_1')| > |N(v) \cap (V_0 \cup V_1')|. \tag{15}$$

Let W_1 be the set of nodes outside of $V_0 \cup V_1$ that were selected during the third phase (possibly multiple times). If property (15) is satisfied, then (deterministically) all nodes in W_1 announce 1 in this phase. Indeed, even if all nodes from V_1' changed their opinion to 0 in the meantime, nodes in V_1 still have majority of their neighbours with opinion 1.

Let us summarize the situation at the beginning of the fourth (and the last) phase. Recall that W_1 consists of nodes that were selected for the first time during the third phase. Let $W_0 = V_0 \cup V_1$ be the set of nodes that were selected before the third phase (that is, during the first or the second phase). A.a.s. nodes in W_1 have opinion 1 and $|W_1| = (T_3 - T_2) + \mathcal{O}(T_3^2/n) \sim T_3$. We may assume that nodes in W_0 have opinion 0 and a.a.s. $|W_0| = T_2(1 + \mathcal{O}(1/\omega)) \sim T_2 = o(T_3)$. Again, it is important to notice that W_1 and W_0 are determined exclusively by the selection process. (V_1 and V_0 do not posses this property and that was the main reason we needed to consider the third phase.) We may then use Chernoff's bound again to show that a.a.s. all nodes (not only outside of $W_1 \cup W_0$!) have more neighbours in W_1 than in W_0. It means that every node that is selected during this last phase announces opinion 1.

Since a.a.s. every node is selected at least once during the next $n(\log n + \omega'/2)$ rounds, the process is over after at most that many rounds with everyone converging to opinion 1. Hence, a.a.s. the entire process takes at most $T_3 + n(\log n + \omega'/2) \le n(\log n + \omega')$ rounds. In fact, the expected number of nodes that were selected before the last phase but were not selected in the first $T_4' = n(\log n - \log \omega/4)$ rounds of the last phase is equal to

$$T_3 \left(1 - \frac{1}{n}\right)^{T_4'} \le \frac{n}{\sqrt{\omega}} \exp(-\log n + \log \omega/4) = \omega^{-1/2} = o(1),$$

and so a.a.s. all nodes selected before the last phase are selected again during the first T_4' rounds of the last phase. On the other hand, a.a.s. there are still some nodes not selected at all after $T_3 + T_4'$ rounds. Indeed, this follows immediately from the well studied coupon collector concentration bound for $\hat{T}$: $\mathbb{P}(\hat{T} < n \log n - cn) < e^{-c}$. The conclusion is that a.a.s. all nodes are selected at least once between round T_3 and $\hat{T}$, and the proof is finished. □

3.4 Very Sparse Random Graphs

In this subsection, we investigate random graphs that are close to the threshold for connectivity but are still connected, that is, we assume that $pn \le \omega \log n$ and $pn \ge \log n + \omega$ for some $\omega = \omega(n) \to \infty$ as $n \to \infty$.

First, we will show that at time $T_3 = T_3(n) = 2n \log n$, every node announced its opinion at least once, and only at most $n\omega / \log n = o(n)$ nodes have opinion 0.

Theorem 3.8. *Let $\omega = \omega(n) = o(\log n)$ be any function that tends to infinity (sufficiently slowly) as $n \to \infty$. Suppose that $p = p(n) \le \omega \log n / n$ and $p \ge (\log n + \omega)/n$. Set $T_3 = T_3(n) = 2n \log n$ and $s = s(n) = n\omega / \log n$. Then, a.a.s. all nodes announced their opinion at time T_3, and at most s of them have opinion 0.*

Proof. Fix $T_2 = T_2(n) = n/\omega$. It follows from Theorem 3.5 that a.a.s. $Y_1^{T_2} \ge (1/2 + \delta/2)T_2$, $Y_0^{T_2} \le (1/2 - \delta/2)T_2$, and trivially $Y_1^{T_2} + Y_0^{T_2} \le T_2$. Let V_i ($i \in \{0, 1\}$) be the set of nodes with opinion i at time T_2. Since we aim for a statement that holds a.a.s., we may assume that the above inequalities are satisfied at time T_2 and continue the process from there. In fact, as explained in Subsection 2.3, we may assume that $|V_1| = Y_1^{T_2} = (1/2 + \delta/2)T_2$ and $|V_0| = Y_0^{T_2} = (1/2 - \delta/2)T_2$.

Note that during the next $T_3 - T_2 \sim 2n \log n$ rounds, a.a.s. all nodes announce their opinion at least once. Indeed, the coupon collector problem is well understood and it is known that a.a.s. it happens after $(1 + o(1))n \log n$ rounds. We consider the nodes that announce opinion 0 at some point during this phase. In particular, let v_1 be the first node that announced opinion 0 during this phase, let $v_2 \ne v_1$ be the second such node, etc. For a contradiction, suppose that at time T_3 there are more than s nodes with opinion 0. It means that the sequence we just constructed consists of more than s nodes; let $S = \{v_1, \ldots, v_s\}$ be the set of the first s nodes in this sequence. Note that all neighbours of v_i in $V_1 \setminus \{v_1, \ldots, v_{i-1}\} \supseteq V_1 \setminus S$ had opinion 1 when v_i announced opinion 0. On the other hand, no neighbour of v_i outside of $V_0 \cup \{v_1, \ldots, v_{i-1}\} \subseteq V_0 \cup S$ had opinion 0 at that point. It follows that for any $v_i \in S$,

$$|N(v_i) \cap (V_1 \setminus S)| \le |N(v_i) \cap (V_0 \cup S)|. \tag{16}$$

In fact, we will relax this property and conclude that for any $v_i \in S$, at least one of the following three properties holds:

$$\begin{aligned} \text{Property (a):} \quad & |N(v_i) \cap (V_1 \setminus S)| \le (1/2 + \delta/4) T_2 p & (17) \\ \text{Property (b):} \quad & |N(v_i) \cap (V_0 \setminus S)| \ge (1/2 - \delta/4) T_2 p & (18) \\ \text{Property (c):} \quad & |N(v_i) \cap S| \ge (\delta/2) T_2 p. & (19) \end{aligned}$$

(Indeed, if none of properties (17)–(19) holds, then (16) does not hold.) We partition the set S into $S = S_a \cup S_b \cup S_c$: nodes in S_x satisfy Property (x). We will show that a.a.s. in $\mathcal{G}(n,p)$ there are no sets V_0, V_1, S, and partition $S = S_a \cup S_b \cup S_c$ such that Properties (a)–(c) hold. This will finish the proof of the theorem.

Let us fix $V_1 \subseteq V$ with $|V_1| = (1/2 + \delta/2)T_2$, $V_0 \subseteq V \setminus V_1$ with $|V_0| = (1/2 - \delta/2)T_2$, $S \subseteq V$ with $|S| = s = n\omega/\log n$, and partition $S = S_a \cup S_b \cup S_c$. For any node $v_i \in S_a$, $|N(v_i) \cap (V_1 \setminus S)|$ is a binomial random variable with expectation $|V_1 \setminus S| p \sim (1/2 + \delta/2)T_2 p$. (Note that $s = o(T_2)$.) It follows from Chernoff's bound (2) that $v_i \in S_a$ satisfies Property (a) with probability at most $\exp(-\Theta(T_2 p))$. Similarly, Chernoff's bound (1) implies that $v_i \in S_b$ satisfies Property (b) with probability at most $\exp(-\Theta(T_2 p))$. More importantly, the events associated with different nodes $v_i \in S_a \cup S_b$ are independent. Unfortunately, this is not the case for events associated with nodes $v_i \in S_c$. To deal with them, we need to consider all of them together. There are $\binom{|S_c|}{2} + |S_c|(s - |S_c|) \le |S_c| s$ pairs of nodes from S such that at least one of them is in S_c. In order for nodes in S_c to satisfy Property (c), at least $(|S_c|/2)(\delta/2)T_2 p$ of such pairs must generate an edge in $\mathcal{G}(n,p)$. Since the expected number of edges is at most $|S_c| s p = o(|S_c| T_2 p)$, by Chernoff's bound (1) we get that it happens with probability at most $\exp(-\Theta(|S_c| T_2 p))$.

Note that by the union bound the probability that there exist sets V_0, V_1, S, and partition $S = S_a \cup S_b \cup S_c$ such that Properties (a)–(c) hold can be upper bounded by

$$\begin{aligned} & \binom{n}{|V_1|} \binom{n - |V_1|}{|V_0|} \binom{n}{s} (2^s)^2 \exp\Big(-\Theta\Big((|S_a| + |S_b| + |S_c|) T_2 p\Big)\Big) \\ & \le \binom{n}{T_2}^2 \binom{n}{s} 2^{2s} \exp\left(-\Theta(s T_2 p)\right) \\ & \le \left(\frac{en}{n/\omega}\right)^{2n/\omega} \left(\frac{en}{s}\right)^s 2^{2s} \exp\left(-\Theta(s T_2 p)\right) \\ & \le \exp\left(\mathcal{O}\left(\frac{2n \log \omega}{\omega} + s \log\log n + s\right) - \Omega\left(\frac{n\omega}{\log n} \cdot \frac{n}{\omega} \cdot \frac{\log n}{n}\right)\right) \\ & \le \exp\left(\mathcal{O}\left(\frac{2n \log \omega}{\omega}\right) - \Omega\left(n\right)\right) = o(1), \end{aligned}$$

which finishes the proof of the theorem. □

We will call a node v to be of *small degree*, if its degree is at most $k = 5 \log n / (\log\log n)^{1/2}$. Nodes of degree larger than k will be called of *large degree*. Before we continue investigating the process, we need to show a well-known fact that small degree nodes are not too close to each other.

Lemma 3.9. *Let $\omega = \omega(n) = o(\log n)$ be any function that tends to infinity (sufficiently slowly) as $n \to \infty$. Suppose that $p = p(n) \le \omega \log n / n$ and $p \ge (\log n + \omega)/n$. Then, the following property holds a.a.s. in $\mathcal{G}(n,p)$: any two small degree nodes are at distance at least 2 from each other.*

Proof. Since $np > \log n$ and $k = o(\log n)$, $\binom{n}{i}p^i$ is an increasing sequence for $0 \le i \le k$ and so

$$\mathbb{P}(\deg(v) \le k) \le \sum_{i=0}^{k} \binom{n}{i} p^i (1-p)^{n-i} \le (k+1)\binom{n}{k} p^k (1-p)^{n-k}.$$

We obtain the following upper bound on the probability that a node v has small degree:

$$\begin{aligned}
\mathbb{P}(\deg(v) \le k) &\le (k+1)\binom{n}{k} p^k (1-p)^{n-k} \\
&\le (k+1)\left(\frac{en}{k} \cdot \frac{\omega \log n}{n}\right)^k \exp\Big(-pn + pk)\Big) \\
&\le (k+1)\Big(\omega(\log\log n)^{1/2}\Big)^k \exp\Big(-\log n - \omega + o(1)\Big) \\
&\le (k+1)\exp\left(\frac{5\log n}{(\log\log n)^{1/2}} \cdot \log\log\log n - \log n\right) \\
&= n^{-1+o(1)}.
\end{aligned}$$

Hence, we expect $n^{o(1)}$ small degree nodes and so a.a.s. we have only $n^{o(1)}$ of them. More importantly, using similar computations one can show that the expected number of small degree nodes that are adjacent to each other is equal to

$$\binom{n}{2} \cdot p \cdot \left(n^{-1+o(1)}\right)^2 = n^{-1+o(1)} = o(1).$$

Similarly, the expected number of small degree nodes that are at distance two from each other is equal to

$$\binom{n}{2} \cdot n \cdot p^2 \cdot \left(n^{-1+o(1)}\right)^2 = n^{-1+o(1)} = o(1).$$

Hence, a.a.s. any two nodes of small degree are at distance at least two from each other, and the proof of the lemma is finished. □

Our next observation is that the number of large degree nodes that have opinion 0 is decreasing.

Theorem 3.10. *Let $\omega = \omega(n) = o(\log n)$ be any function that tends to infinity (sufficiently slowly) as $n \to \infty$. Suppose that $p = p(n) \le \omega \log n/n$ and $p \ge (\log n + \omega)/n$. Then, the following property holds a.a.s. for all phases.*

Suppose that at the beginning of a phase, $s = (n\omega/\log n) \cdot (\log\log n)^{-(i-1)/4}$ large degree nodes have opinion 0 for some $i \in \mathbb{N}$. Then, after $2n\log n$ rounds all nodes announced their opinion at least once more, and at most $u = s/(\log\log n)^{1/4} = (n\omega/\log n) \cdot (\log\log n)^{-i/4}$ large degree nodes have opinion 0.

Proof. First, note that the expected number of nodes that were not selected in $2n\log n$ rounds is

$$n\left(1 - \frac{1}{n}\right)^{2n\log n} \le n\exp(-2\log n) = 1/n,$$

so with probability $1 - \mathcal{O}(1/\log n)$ all of them are selected at least once in any phase consisting of $2n\log n$ rounds. Since we will iteratively apply the argument for $\mathcal{O}(\log n/\log\log\log n) = o(\log n)$ phases, all of them have the desired property a.a.s.

Since we aim for a statement that holds a.a.s., we may assume that the graph satisfies property stated in Lemma 3.9. For a contradiction, suppose that some phase fails, that is, at the beginning of this phase s large degree nodes have opinion 0, and at the end of this phase more than $u = s/(\log\log n)^{1/4}$ large degree nodes have opinion 0. As in the proof of Theorem 3.8, we consider a sequence of distinct nodes, $v_1, v_2, \ldots$, in which large degree nodes announce opinion 0: v_1 announced opinion 0 first, then $v_2 \neq v_1$, etc. Let $U = \{v_1, \ldots, v_u\}$ be the set of the first u nodes in this sequence and let S be the set of large degree nodes that have opinion 0 at the beginning of this phase. Recall that each large degree node has degree at least $k = 5(\log n)(\log\log n)^{-1/2}$ and at most one neighbour of small degree (Lemma 3.9). Small degree nodes may (or may not) have opinion 0 but no large degree node outside of $S \cup U$ has opinion 0 at the time node v_i announced opinion 0. We conclude that for all $i \in [u]$, v_i has at least $k/2 - 1 \geq 2(\log n)(\log\log n)^{-1/2}$ neighbours in $S \cup U$.

We say that set U satisfies Property (a) if the following holds:

Property (a): at least $u/2$ nodes in U have at least $(\log n)(\log\log n)^{-1/2}$ neighbours in U.

If U does not satisfy Property (a), then less than $u/2$ of nodes in U have at least $(\log n)(\log\log n)^{-1/2}$ neighbours in U, which implies that set U (together with S) satisfies the following property:

Property (b): at least $u/2$ nodes in U have at least $(\log n)(\log\log n)^{-1/2}$ neighbours in $S \setminus U$.

We will deal with each property independently and show that it is not present in $\mathcal{G}(n,p)$ with the desired probability.

If Property (a) is satisfied for some set U of size u, then U induces at least $u(\log n)(\log\log n)^{-1/2}/4$ edges. Hence, the probability that some set of size u has this property is at most

$$
\begin{aligned}
&\binom{n}{u}\binom{\binom{u}{2}}{u(\log n)(\log\log n)^{-1/2}/4} p^{u(\log n)(\log\log n)^{-1/2}/4} \\
&\leq \binom{n}{u}\left(\frac{eu^2/2}{u(\log n)(\log\log n)^{-1/2}/4} \cdot \frac{\omega \log n}{n}\right)^{u(\log n)(\log\log n)^{-1/2}/4} \\
&\leq \binom{n}{u}\left(\frac{2e\omega u(\log\log n)^{1/2}}{n}\right)^{u(\log n)(\log\log n)^{-1/2}/4} \\
&\leq \binom{n}{u}\left(\frac{2e\omega s(\log\log n)^{1/4}}{n}\right)^{u(\log n)(\log\log n)^{-1/2}/4} \\
&\leq n^u\left(\frac{2e\omega^2(\log\log n)^{1/4}}{\log n}\right)^{u(\log n)(\log\log n)^{-1/2}/4} \\
&\leq \exp\left(u\log n - \frac{u(\log n)}{4(\log\log n)^{1/2}} \cdot (1+o(1))\log\log n\right) \\
&= \mathcal{O}(1/\log n).
\end{aligned}
$$

If Property (b) is satisfied for some set U of size u and some set S of size s, then there exists a subset $U' \subseteq U$ of size $u/2$ such that each $v_i \in U'$ has at least $(\log n)(\log\log n)^{-1/2}$ neighbours in

$S \setminus U$. The probability that a given $v_i \in U'$ has this property is at most

$$\begin{aligned} &\binom{s}{(\log n)(\log\log n)^{-1/2}} p^{(\log n)(\log\log n)^{-1/2}} \\ &\qquad \le \left(\frac{es}{(\log n)(\log\log n)^{-1/2}} \cdot \frac{\omega \log n}{n} \right)^{(\log n)(\log\log n)^{-1/2}} \\ &\qquad \le \left(\frac{e\omega^2 (\log\log n)^{1/2}}{(\log n)} \right)^{(\log n)(\log\log n)^{-1/2}} \\ &\qquad \le \exp\Big(-(\log n)(\log\log n)^{-1/2} \cdot (1+o(1)) \log\log n \Big) \\ &\qquad = \exp\Big(-(1+o(1))(\log n)(\log\log n)^{1/2} \Big). \end{aligned}$$

Moreover, the events associated with different $v_i \in U'$ are independent. Hence, by the union bound, the probability that there exist a pair of sets U, S, and a partition $U = U' \cup (U \setminus U')$ can be upped bounded by

$$\begin{aligned} &\binom{n}{s}\binom{n}{u} 2^u \exp\Big(-(1+o(1))(\log n)(\log\log n)^{1/2} \Big) \\ &\qquad \le \exp\Big(s\log n + u \log n + u - (1+o(1))(\log n)(\log\log n)^{1/2} \cdot (u/2) \Big) \\ &\qquad = \exp\Big((1+o(1)) s \log n - (1+o(1))(\log n)(\log\log n)^{1/4} \cdot (s/2) \Big) \\ &\qquad = \mathcal{O}(1/\log n). \end{aligned}$$

This finishes the proof of the theorem as the argument has to be (iteratively) applied only for $\mathcal{O}(\log n / \log\log\log n) = o(\log n)$ phases. □

Finally, we are ready to show that all nodes eventually converge to opinion 1.

Proof of Theorem 1.2. The proof is an easy consequence of Theorems 3.8, 3.10, and Lemma 3.9. Indeed, a.a.s. at time $T_3 = T_3(n) = 2n \log n$, all but at most $s = s(n) = n\omega/\log n$ nodes have opinion 1 (Theorem 3.8). Most of them are of large degree but some of them may be of small degree. By Theorem 3.10, the number of large degree nodes that have opinion 0 decreases: a.a.s. at time $2n \log n \cdot \mathcal{O}(\log n / \log\log n) = \mathcal{O}(n (\log n)^2 / (\log\log n))$ no large degree node has opinion 0. There could possibly be still some nodes of small degree that have opinion 0 but everyone converges to opinion 1 after additional $\mathcal{O}(n \log n)$ rounds. Indeed, every node is selected at least once during that time period a.a.s. Large degree nodes have many neighbours but at most one neighbours of small degree (Lemma 3.9). So they will not change their opinion and stay with opinion 1. On the other hand, by the same lemma, no small degree node has a neighbour of small degree. Hence, such nodes will switch to opinion 1 once they are selected again. This finishes the proof of the theorem. □

4 Dense Random Graphs

In this section, we prove that for dense graphs (that is, when $p \in (0, 1]$ is a constant) it is not true that all nodes converge to the correct opinion a.a.s. On the contrary, there maybe an information cascade where all the nodes converge to the wrong opinion with constant probability.

Proof of Theorem 1.3. Fix any $p \in (0, 1)$. We will consider the case $p = 1$ (easy case) at the end of the proof.

Trivially, the first node announces its private belief, that is, it announces opinion 1 with probability $1/2 + \delta$; otherwise, it announces 0. Since nodes are selected by the process ("coupon collector") independently of the graph, we may postpone exposing edges of the random graph till the first time a node is selected. Each time this happens, we expose edges from v^t to all nodes that already announced their opinion. If every single time at least one edge is present, then all nodes are going to announce the opinion of the very first node. It follows that

$$p_1 \geq (1/2 + \delta) \prod_{i=1}^{n} \Big(1 - (1-p)^i\Big).$$

It is easy to see that for any $x \in [0, 1-p]$,

$$f(x) = 1 - x \geq \exp\left(-\frac{\log(1/p)}{1-p} x\right) = g(x).$$

(Note that $f(0) = g(0)$, $f(1-p) = g(1-p)$, and $g(x)$ is convex.) Hence,

$$\begin{aligned} p_1 &\geq (1/2 + \delta) \exp\Big(- \frac{\log(1/p)}{1-p} \sum_{i=1}^{n} (1-p)^i \Big) \\ &\geq (1/2 + \delta) \exp\Big(- \log(1/p) \sum_{i=0}^{\infty} (1-p)^i \Big) \\ &= (1/2 + \delta) \exp\Big(- \log(1/p)(1/p) \Big). \end{aligned}$$

The same argument works for p_0 with the only difference that the probability of the first node announcing 1 ($1/2 + \delta$) needs to be replaced with the probability of announcing 0 ($1/2 - \delta$).

Finally, note that if $p = 1$, then the graph is (deterministically) the complete graph and (again, deterministically) all nodes are going to adopt the opinion of the very first node. Thus, we immediately get $p_1 = 1/2 + \delta$ and $p_0 = 1/2 - \delta$ (which matches the general formula that works for $p \in (0, 1]$). This finishes the proof of the theorem. □

References

[1] Daron Acemoglu, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *Review of Economic Studies*, 78(4):1201–1236, 2011.

[2] Aris Anagnostopoulos, Luca Becchetti, Emilio Cruciani, Francesco Pasquale, and Sara Rizzo. Biased opinion dynamics: when the devil is in the details. *Information Sciences*, 593:49–63, 2022.

[3] Maryam Bahrani, Nicole Immorlica, Divyarthi Mohan, and S Matthew Weinberg. Asynchronous majority dynamics in preferential attachment trees. *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, 2020.

[4] Abhijit Banerjee, Emily Breza, Arun G Chandrasekhar, and Markus Mobius. Naive learning with uninformed agents. *American Economic Review*, 111(11):3540–3574, 2021.

[5] Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. *Games and Economic Behavior*, 46(1):1–22, January 2004. URL: http://ideas.repec.org/a/eee/gamebe/v46y2004i1p1-22.html.

[6] Abhijit V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.

[7] Luca Becchetti, Andrea Clementi, and Emanuele Natale. Consensus dynamics: An overview. *SIGACT News*, 51(1):58–104, March 2020. doi:10.1145/3388392.3388403.

[8] Luca Bechetti, Andrea E.F. Clementi, Emanuele Natale, Francesco Pasquale, and Luca Trevisan. Stabilizing consensus with many opinions. In *ACM Symposium on Discrete Algorithms (SODA)*, 2016.

[9] Itai Benjamini, Siu-On Chan, Ryan O'Donnell, Omer Tamuz, and Li-Yang Tan. Convergence, unanimity and disagreement in majority dynamics on unimodular graphs and random graphs. *Stochastic Processes and their Applications*, 126(9):2719–2733, 2016.

[10] Sushil Bikhchandani, David Hirshleifer, Omer Tamuz, and Ivo Welch. Information cascades and social learning. Technical report, National Bureau of Economic Research, 2021.

[11] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy*, 100(5):992–1026, October 1992.

[12] Béla Bollobás. *Random graphs*. Cambridge University Press, 2001.

[13] Boğaçhan Çelen and Shachar Kariv. Observational learning under imperfect information. *Games and Economic behavior*, 47(1):72–86, 2004.

[14] Debsoumya Chakraborti, Jeong Han Kim, Joonkyung Lee, and Tuan Tran. Majority dynamics on sparse random graphs. *Random Structures & Algorithms*, 63(1):171–191, 2023.

[15] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973. URL: http://www.jstor.org/stable/2335008.

[16] Emilio Cruciani, Hlafo Alfie Mimun, Matteo Quattropani, and Sara Rizzo. Phase transitions of the k-majority dynamics in a biased communication model. In *Proceedings of the 22nd International Conference on Distributed Computing and Networking*, pages 146–155, 2021.

[17] Morris H. DeGroot. Reaching a consensus. *Review of Economic Studies*, 69(345):118–121, 1974.

[18] Michal Feldman, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. Reaching consensus via non-bayesian asynchronous learning in social networks. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2014*, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.192.

[19] Nikolaos Fountoulakis, Mihyun Kang, and Tamás Makai. Resolution of a conjecture on majority dynamics: Rapid stabilization in dense random graphs. *Random Structures & Algorithms*, 57(4):1134–1156, 2020.

[20] Douglas Gale and Shachar Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329 – 346, 2003. Special Issue in Honor of Robert W. Rosenthal. URL: http://www.sciencedirect.com/science/article/pii/S0899825603001441, doi:https://doi.org/10.1016/S0899-8256(03)00144-1.

[21] Mohsen Ghaffari and Johannes Lengler. Nearly-tight analysis for 2-choice and 3-majority consensus dynamics. In *ACM Symposium on Principles of Distributed Computing (PODC)*, 2018.

[22] Mohsen Ghaffari and Merav Parter. A polylogarithmic gossip algorithm for plurality consensus. In *ACM Symposium on Principles of Distributed Computing (PODC)*, 2016.

[23] Benjamin Golub and Matthew O. Jackson. Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.

[24] Richard A. Holley and Thomas M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, 3(4):643–663, 1975. URL: http://www.jstor.org/stable/2959329.

[25] C Douglas Howard. Zero-temperature ising spin dynamics on the homogeneous tree of degree three. *Journal of applied probability*, 37(3):736–747, 2000.

[26] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random graphs*, volume 45. John Wiley & Sons, 2011.

[27] Y. Kanoria and O. Tamuz. Tractable bayesian social learning on trees. *IEEE Journal on Selected Areas in Communications*, 31(4):756–765, 2013.

[28] Yashodhan Kanoria, Andrea Montanari, et al. Majority dynamics on trees and the dynamic cavity method. *The Annals of Applied Probability*, 21(5):1694–1748, 2011.

[29] Michal Karoński and Alan Frieze. *Introduction to Random Graphs*. Cambridge University Press, 2016.

[30] Marcos Kiwi, Lyuben Lichev, Dieter Mitsche, and Paweł Prałat. Label propagation on binomial random graphs. *arXiv preprint arXiv:2302.03569*, 2023.

[31] S. Matwin, A. Milios, P. Prałat, A. Soares, and F. Théberge. *Generative Methods for Social Media Analysis*. SpringerBriefs in Computer Science. Springer Nature Switzerland, 2023. URL: https://books.google.ca/books?id=wPbJEAAAQBAJ.

[32] Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pages 195–248. Springer, 1998.

[33] Markus Mobius and Tanya Rosenblat. Social learning in economics. *Annual Review of Economics*, 6(1):827–847, 2014. arXiv:https://doi.org/10.1146/annurev-economics-120213-012609, doi:10.1146/annurev-economics-120213-012609.

[34] E. Mossel, N. Olsman, and O. Tamuz. Efficient bayesian learning in social networks with gaussian estimators. In *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2016.

[35] Elchanan Mossel, Joe Neeman, and Omer Tamuz. Majority dynamics and aggregation of information in social networks. In *Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2013.

[36] Elchanan Mossel, Allan Sly, and Omer Tamuz. Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, 2014.

[37] Manuel Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1):1–40, 2013.

[38] Dinah Rosenberg, Eilon Solan, and Nicolas Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979 – 994, 2009.

[39] Lones Smith and Peter Sorensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, March 2000. URL: http://ideas.repec.org/a/ecm/emetrp/v68y2000i2p371-398.html.

[40] Omer Tamuz and Ran Tessler. Majority dynamics and the retention of information. In *Working paper*, 2013.

[41] Linh Tran and Van Vu. Reaching a consensus on random networks: The power of few. *arXiv preprint arXiv:1911.10279*, 2019.

5 Appendix

Proof of Lemma 3.7. Let us first consider any odd value of $k \geq 3$. We get that

$$\begin{aligned} q_k &= \sum_{i\geq(k+1)/2} \binom{k}{i} (1/2+\delta/2)^i (1/2-\delta/2)^{k-i} \\ &\geq \frac{1/2+\delta/2}{1/2-\delta/2} \cdot \binom{k}{(k+1)/2} (1/2-\delta/2)^{(k+1)/2} (1/2+\delta/2)^{(k-1)/2} \\ &\quad + \left(\frac{1/2+\delta/2}{1/2-\delta/2}\right)^3 \sum_{i\geq(k+3)/2} \binom{k}{i} (1/2-\delta/2)^i (1/2+\delta/2)^{k-i} \\ &= \frac{1/2+\delta/2}{1/2-\delta/2} \cdot A + \left(\frac{1/2+\delta/2}{1/2-\delta/2}\right)^3 \cdot B, \end{aligned}$$

where

$$\begin{aligned} A &= \binom{k}{(k+1)/2} (1/2-\delta/2)^{(k+1)/2} (1/2+\delta/2)^{(k-1)/2} \\ B &= \sum_{i\geq(k+3)/2} \binom{k}{i} (1/2-\delta/2)^i (1/2+\delta/2)^{k-i}. \end{aligned}$$

Note that $A+B = 1-q_k \leq 1/2$. More importantly, if $q_k \geq 1/2+51\delta/100$, then the desired property holds and there is nothing to prove. Hence, we may assume that $1-q_k \geq 1/2-51\delta/100 \geq 449/1000$.

It follows that

$$\begin{aligned}
A &= \binom{k}{(k+1)/2}\Big((1/2-\delta/2)(1/2+\delta/2)\Big)^{(k-1)/2}(1/2-\delta/2)\\
&= \binom{k}{(k+1)/2}\Big((1/4-\delta^2/4)\Big)^{(k-1)/2}(1/2-\delta/2)\\
&\leq \binom{k}{(k+1)/2}(1/2)^k \leq \frac{3}{8} \leq \frac{375}{449}(A+B) \leq \frac{9}{10}(A+B).
\end{aligned}$$

As a consequence,

$$\begin{aligned}
q_k &\geq \frac{1/2+\delta/2}{1/2-\delta/2}\cdot A + \left(\frac{1/2+\delta/2}{1/2-\delta/2}\right)^3\cdot B\\
&\geq \left(\frac{1/2+\delta/2}{1/2-\delta/2}\cdot\frac{9}{10} + \left(\frac{1/2+\delta/2}{1/2-\delta/2}\right)^3\cdot\frac{1}{10}\right)(A+B),
\end{aligned}$$

where the last inequality follows from the fact that the coefficient in front of A is smaller than the one in front of B, so a linear combination of the coefficient's are minimized when A is the largest it can possibly be (which is $9/10(A+B)$). Moreover, since $A+B = 1-q_k$ we have,

$$\begin{aligned}
\frac{q_k}{1-q_k} &\geq \frac{1+\delta}{1-\delta}\cdot\frac{9}{10} + \left(\frac{1+\delta}{1-\delta}\right)^3\cdot\frac{1}{10}\\
&= (1+\delta)(1+\delta+\mathcal{O}(\delta^2))\cdot\frac{9}{10} + (1+3\delta+\mathcal{O}(\delta^2))(1+3\delta+\mathcal{O}(\delta^2))\cdot\frac{1}{10}\\
&= (1+2\delta+\mathcal{O}(\delta^2))\cdot\frac{9}{10} + (1+6\delta+\mathcal{O}(\delta^2))\cdot\frac{1}{10}\\
&= (1+2\delta+\mathcal{O}(\delta^2))\cdot\frac{9}{10} + (1+6\delta+\mathcal{O}(\delta^2))\cdot\frac{1}{10}\\
&= 1+\frac{12}{5}\delta+\mathcal{O}(\delta^2) = \frac{1+6\delta/5}{1-6\delta/5}+\mathcal{O}(\delta^2) = \frac{1/2+3\delta/5}{1/2-3\delta/5}+\mathcal{O}(\delta^2).
\end{aligned}$$

Clearly,

$$\frac{q_k}{1-q_k} \geq \frac{1/2+51\delta/100}{1/2-51\delta/100}$$

for sufficiently small δ but one can show it holds for $\delta \in (0, 1/10]$. This implies $q_k \geq 1/2 + 51\delta/100$ for odd values of $k \geq 3$.

Let us now consider any even value of $k \geq 2$. This time we get that

$$\begin{aligned}
q_k &= \sum_{i\geq k/2+1}\binom{k}{i}(1/2+\delta/2)^i(1/2-\delta/2)^{k-i} + \binom{k}{k/2}(1/2+\delta/2)^{k/2}(1/2-\delta/2)^{k/2}(1/2+\delta)\\
&\geq \left(\frac{1/2+\delta/2}{1/2-\delta/2}\right)^2\sum_{i\geq k/2+1}\binom{k}{i}(1/2-\delta/2)^i(1/2+\delta/2)^{k-i}\\
&\quad+\frac{1/2+\delta}{1/2-\delta}\cdot\binom{k}{k/2}(1/2+\delta/2)^{k/2}(1/2-\delta/2)^{k/2}(1/2-\delta)\\
&= \left(\frac{1+\delta}{1-\delta}\right)^2\cdot A + \frac{1/2+\delta}{1/2-\delta}\cdot B,
\end{aligned}$$

where

$$A = \sum_{i \geq k/2+1} \binom{k}{i} (1/2 - \delta/2)^i (1/2 + \delta/2)^{k-i}$$

$$B = \binom{k}{k/2} (1/2 + \delta/2)^{k/2} (1/2 - \delta/2)^{k/2} (1/2 - \delta),$$

and $A + B = 1 - q_k$. Since

$$\left(\frac{1+\delta}{1-\delta}\right)^2 = \frac{1 + 2\delta + \delta^2}{1 - 2\delta + \delta^2} \geq \frac{1 + 19\delta/10}{1 - 19\delta/10} = \frac{1/2 + 19\delta/20}{1/2 - 19\delta/20} \geq \frac{1/2 + 51\delta/100}{1/2 - 51\delta/100}$$

and, trivially,

$$\frac{1/2 + \delta}{1/2 - \delta} \geq \frac{1/2 + 51\delta/100}{1/2 - 51\delta/100},$$

we get that

$$q_k \geq \frac{1/2 + 51\delta/100}{1/2 - 51\delta/100}(1 - q_k),$$

which implies $q_k \geq 1/2 + 51\delta/100$ for even values of k too. □