Can society learn without opinion leaders?*

Itai Arieli † Fedor Sandomirskiy † ‡ Rann Smorodinsky †

Abstract

A model of learning on a social network with fully-rational Bayesian agents and random arrival order is considered. We show that if for a given large network, most of the agents learn the correct action, then learning persists with high probability even after deleting randomly 99% of the population. In other words, random groups of agents are never critical for learning even if these groups contain a majority of agents.

This robustness does not rule out the existence of a certain small minority of agents critical for learning. The classical sociological theory of “two-step information flow” conveys the idea that learning is always facilitated by such a small group of influencers predetermined by the network structure. We challenge this thesis by showing that learning can also be possible in totally-egalitarian networks, where all agents are ex-ante symmetric and thus equally influential. We construct such networks relying on insights from the theory of expanders and demonstrate that adversarial elimination of large groups of agents does not spoil learning outcomes, i.e., no minority of agents (and even majority) is critical for learning.

For a network-designer, our results offer a recipe for building networks that aggregate information even if a subset of agents do not participate in the learning process. This subset of absent agents can be large and picked in an adversarial way.

1 Introduction

The way ideas and information propagate in society is critical for engineering political campaigns, introducing new products and for formulating new social conventions. The leading theory in mass media communication, dating back to [Katz and Lazarsfeld, 1955] asserts that new content is generated sparsely and is consumed by most agents indirectly.[1] The original version, formulated during the golden era of TV networks, argues that although such networks were responsible for sparking new ideas and introducing new products, this content was consumed by most people in an indirect manner.

*We are grateful to Matthew Jackson, Nicolas Vieille, and Omer Tamuz for inspirational discussions and suggestions. We also thank attendees of the game-theory seminar at Technion and Conference on Mechanism and Institution Design 2020 for their feedback.
Itai Arieli is supported by the Ministry of Science and Technology grant #2028255.
Fedor’s work is partially supported by the Lady Davis Foundation, by Grant 19-01-00762 of the Russian Foundation for Basic Research, by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (#740435), and by the Basic Research Program of the National Research University Higher School of Economics.
Rann’s research is supported by the United States-Israel Binational Science Foundation, National Science Foundation grant #2016734, German-Israel Foundation grant I-1419-118.4/2017, Ministry of Science and Technology grant #19400214, Technion VPR grants, and the Bernard M. Gordon Center for Systems Engineering at the Technion.

†Technion IE&M, Haifa (Israel)
‡Higher School of Economics, St.Petersburg (Russia)

[1] [Katz and Lazarsfeld, 1955] builds on earlier work [Lazarsfeld et al., 1944].

In other words, people's actions, whether in the form of adopting a new product or voting for a certain political candidate was not so much of what they hear from the TV networks but rather what they heard from opinion leaders who, in turn, consumed such content directly (this is often dubbed as the 'influentials hypothesis'). Indeed, Katz and Lazarsfeld [1955] propose to refer to their model as a model of two-step information flow. The original content flows from its originator(s) to intermediaries ('opinion leaders') and thereafter to the masses. The two-step flow model stood in stark contrast with the one-step flow paradigm which argues that most people are effected directly by the media outlets [Bineham, 1988].

The ordinal ideas underlying the two-step information flow were adapted throughout the years as TV networks lost power and influence and peer-to-peer social networks became more central to the generation and flow of information. The importance of opinion leaders as information brokers came to occupy a central place in the literatures of the diffusion of innovations and marketing [Rogers, 1995, Valente and Rogers, 1995, Chan and Misra, 1990, Coulter et al., 2002, Myers and Robertson, 1972, Van den Bulte and Joshi, 2007, Vernette, 2004]. To appreciate the centrality of these ideas we quote an observation by [Weimann, 1994] whereby nearly 4000 studies of influentials, opinion leaders, and personal influence were conducted in the decades after Katz and Lazarsfeld. It is no wonder that, in practice, for many marketing and political campaigns the importance of identifying the opinion leaders became acute.

In their well cited work, Watts and Dodds [2007] study information flow using computer simulations of interpersonal information flows. They ask whether information cascades are necessarily driven by a minority of opinion leaders. Their surprising conclusion is that this hypothesis cannot be tested as the notion of an opinion leader is not well specified. They suggest that the 'influentials hypothesis' requires careful specification: "But what exactly does the two-step flow say about influentials, and how precisely do they exert influence over the (presumably much larger) population of non-influentials?"

To cope with the challenge posed in [Watts and Dodds, 2007] we study a variant of the celebrated model of information cascades [Banerjee, 1992, Bikhchandani et al., 1992]. In our model agents reside on a social network. The social network is fixed and agents arrive in some random order and act sequentially. Each agent observes actions taken by his neighbors who have acted before him. The agent uses this information, alongside a private signal, to choose his action. We keep the model simple and focus on the binary action case with a symmetric binary signal. We say that the society learns if most agents take the correct action with high probability.

In this model there are agents that act in isolation and so purely follow their own signal. These agents will be deemed "content generators" and their identity is random as it depends on the arrival order. All other agents are influenced, through the social network structure, by content generators, either directly (one-step flow) or indirectly (two-step flow). Content generators are always bound to make an error. Thus, if society learns then the proportion of content generators must be small. However, this does not contradict the possibility that all other agents learn by observing a growing number of such agents indirectly.

We interpret the two-step information flow conjecture as follows:

1. Conjecture C1: If society aggregates information then it must be the case that most agents consume information indirectly, i.e., through intermediaries.

On the other hand, we will refer to a subset of agents as opinion leaders (or "influential") if, in their presence, society learns while in their absence it does not.

2. Conjecture C2: A necessary condition for information aggregation is the existence of a (small) subset of influential agents.

We first show that indeed Conjecture C1 is true. In fact, a simple counting argument demonstrates that agents observe on average a single content generator (see Section 3). With just one single observation and one's own signal it is obvious that learning cannot hold. This implies that for learning to hold it must be the case that most content is consumed indirectly.

On the other hand, Conjecture C2 is shown to be false. We first argue that deleting a random subset of agents of a fixed size does not reverse the possibility of learning. In fact, this result holds even if we randomly delete a large proportion (e.g., 99%) of the agents (Theorem 4.1). This, however, does not exclude the possibility of having a certain small influential minority. We go further and demonstrate existence of fully symmetric networks that accommodate learning (Theorem 5.1). Symmetry suggests that in such networks no minority can be influential. We prove a much stronger statement implying existence of networks, where even adversarial elimination of any subset of a fixed proportion (e.g., even 99%) of agents does not spoil information aggregation (Theorem 6.1). Hence, no group of agents is influential. Theorems 5.1 and 6.1 leverage results from the CS literature on expander graphs.

The indirect takeaway from these two contradicting conclusions is that indeed, as hinted in [Watts and Dodds, 2007], the two-step information flow is insufficiently specified. Without more rigor it cannot be put to a test.

Network-designer perspective. An alternative way to perceive our results is as follows. Consider a designer of a social network who wants to ensure social learning. Sgroi [2002] constructs a network that admits learning, Bahar et al. [2020] propose the celebrities graph as a network that admits learning and is robust to the arrival order of agents. However, the celebrities graph only admits learning if *all* agents actually participate and make a decision. In other words, if a minority of agents do not care about the decision at hand and is consequently inactive, then learning may fail. In contrast, our expander-based networks from Theorem 6.1 is a social structure that admits learning and is robust both to the arrival order and to the actual subset of active agents.

1.1 Related literature

Our work is closely tied to the a rich body of informal work (in the mathematical sense) on information flow. We refer the reader to [Watts and Dodds, 2007] for a comprehensive list of related papers. Some formal models have also been used to study information dissemination over networks. Much of this literature, starting with [DeGroot, 1974], focuses on boundedly rational agents, e.g. [Golub and Jackson, 2010, Ellison and Fudenberg, 1993]. More recently a parallel body of literature developed that focuses on rational agents, as we do. The main objective in this body of work is to study necessary or sufficient conditions for information aggregation in terms of the information structure and the topology of the network. For example, Smith [1991] and Sgroi [2002] provide sufficient conditions for information aggregation on the network topology whenever the signals are bounded. In both cases, they assume the order of arrival is known. Mossel et al. [2019] study more comprehensively the interplay between the structure of the network and the structure of information for the agents that is required to achieve learning. Acemoglu et al. [2010] look at information aggregation in a framework of a randomly generated network, where each new arriving agent is randomly connected to the set of existing ones while Arieli and Mueller-Frank [2018] assume the edges in the graph are limited to the m-dimensional lattice. Most closely related is [Bahar et al., 2020] who construct a family of social networks for which information is aggregated when the arrival order is random. Their constriction is essentially a bi-partitive graph with a minority of agents on one side (nicknamed "celebrities") and a majority on the other side. Indeed, for such a construction the majority learns to play optimally indirectly and there exists a clear subset of opinion leaders (the celebrities); our Example 2.2 contains the detailed discussion of this result.

None of these papers, to the best of our knowledge, confronts the two-step information flow conjecture. However, many of these models exhibit conjecture C1 (in some papers this is stated explicitly and some implicitly). This should not surprise us as we show this is true more generally. On the other hand, none of these papers, with the exception of [Bahar et al., 2020], touches on conjecture C2 which relates learning with the existence of a small subset of opinion leaders. In fact, the main construction in [Bahar et al., 2020] could suggest, as implied by the two-step information flow, that such a subset is necessary for learning which we demonstrate to be incorrect.

On a more technical level, our robustness results, which contradict conjecture C2, rely on the theory of expanders, regular sparse graphs having numerous applications from pure math to designing error-correcting codes (see [Lubotzky, 2012] for a survey). Recently, this theory was applied to problems of social learning by Mossel et al. [2014] and Feldman et al. [2014]. The main distinction between these papers and ours is that we consider rational Bayesian agents where these works consider boundedly rational agents that apply majority rule to decide upon their action. Another distinction is the fact that in our case the agents decide only once according to a random order where in both other works agents decides on an action repeatedly. So in our model, agents have only one opportunity to guess the right action and cannot revise it later on.

1.2 Structure of the paper

Section 2 contains the description of the model and examples. In Section 3 we demonstrate necessity of intermediaries for learning. Section 4 shows that even large random subset of agents is not influential with high probability. However, this does not rule out the possibility that a certain minority of agents is crucial. This possibility is examined in Section 5. We show that learning is possible in totally egalitarian societies, where all agents are symmetric. We strengthen these results in Section 6 and demonstrate that there are symmetric networks such that no subset of agents is influential, i.e., even after deleting 99% of agents in an adversarial way remaining group will still aggregate the information very well. Section 7 discusses extensions of the basic model and Section 8 concludes.

Results from Sections 5 and 6 heavily rely on the insights from the theory of expander graphs, which are explained in Appendix A. Technical proofs for these two sections are in Appendices B and C.

2 Preliminaries

2.1 The model

A social network is an undirected graph $G = (V, E)$, where the set of vertices V is the set of agents and an edge $vu = uv$ is contained in E if v and u are "friends", i.e., friendship is always mutual as on Facebook. We denote by F_v the set of v's friends $\{u \in V : vu \in E\}$ and by $\deg_v$ the degree of v in G, i.e., $\deg_v = |F_v|$.

With each agent $v \in V$ we associate his arrival time t_v; all $(t_v)_{v\in V}$ are independent random variables uniformly distributed on $[0, 1]$. Random arrival order allows us to disentangle the network topology and the order in which agents make their decisions. Randomness is also justified by the fact that this order usually differs from one issue to another (e.g, iPhone or Android, public kindergartens or private), see [Bahar et al., 2020].

Upon arrival, each agent v takes an action $a_v \in \{0, 1\}$, depending on the information available. The agent gets a payoff of 1 if $a_v = \theta$, where θ is a random state, a Bernoulli random variable with success probability 1/2; if the action does not match the state, the agent gets zero payoff. Nobody observes θ but every agent receives a binary signal $s_v \in \{0, 1\}$ that equals θ with probability $1-\varepsilon$ and $1-\theta$ with probability ε, where $0 < \varepsilon < \frac{1}{2}$. Signals are independent conditional on θ. In addition to his own signal s_v, an agent v observes the set of his arrived friends $F_{<v} = \{u \in V \,:\, vu \in E,\ t_u < t_v\} \subset F_v$ and their actions $(a_u)_{u\in F_{<v}}$. We denote the information set of an agent v by $I_v = (s_v, (a_u)_{u\in F_{<v}})$.

All agents are rational and risk-neutral; the description of the model, probability distributions, and the graph G are common knowledge.

Arrival times $(t_v)_{v\in V}$ induce an orientation on edges of $G = (V, E)$ from early arrivers to late and convert it into the direct network $G((t_u)_{u\in V})$. This orientation captures possible paths of information transmission: the action a_w of an agent w can affect the choice made by an agent v only if there is a directed path from w to v. The *realized subnetwork* of an agent v is the directed sub-network $G_{<v}((t_u)_{u\in V}) = (V_{<v}, E_{<v})$ of $G((t_u)_{u\in V})$ consisting of all u such that there is a directed path from u to v, i.e., there is a sequence of agents $u_0, u_1, u_2, \ldots, u_k$ such that $u_0 = u$, $u_k = v$, all edges $u_i u_{i+1}$, $i = 0, \ldots k$ belong to E, and $t_{u_i} < t_{u_{i+1}}$ for all i. In other words, $G_{<v}((t_u)_{u\in V})$ is the connected component of v in the directed sub-network composed of v and his predecessors.

In Section 4, we will need the following assumption; all other results are independent of it.

Assumption (★). Each agent v, in addition to his signal and actions of arrived friends, observes the set of agents $V_{<v}$ from his realized subnetwork (without their actions and arrival times). In other words, an agent knows the set of those who can possibly affect his action. The information set of v is, therefore, $I_v = (s_v, (a_u)_{u\in F_{<v}}, V_{<v})$.

This assumption plays a role of simplifying the equilibrium structure, see Subsection 2.1.1. Section 7 discusses how to get rid of this and other technical assumptions.

2.1.1 Equilibria

A mixed strategy of σ_v of an agent v maps his information set I_v to the probability distribution on $\{0, 1\}$, according to which his action a_v is then chosen. The goal of the agent is to maximize the probability of taking the action that matches the state.

We consider an equilibrium $\sigma = (\sigma_v)_{v\in V}$ of the induced Bayesian game and use $\mathbb{P}_\sigma$ and $\mathbb{E}_\sigma$ for the probability and for the expectation with respect to all the randomness in the problem (state, signals, arrival times, and actions). An equilibrium exists since it is a finite game; however, it may be non-unique. We drop the dependence on σ and write $\mathbb{P}$ and $\mathbb{E}$ when this creates no confusion.

We say that an equilibrium is *state-symmetric* if the distribution $\mathbb{P}$ is invariant under the mapping $\big(\theta,\, (s_v)_{v\in V},\, (a_v)_{v\in V}\big) \to \big(1-\theta,\, (1-s_v)_{v\in V},\, (1-a_v)_{v\in V}\big)$, i.e., simultaneous flipping the state, signals, and actions. Since the payoffs enjoy this symmetry, a state-symmetric equilibrium exists.

In any equilibrium, an agent v selects $a_v = 1$ if $\mathbb{P}(\theta = 1 \mid I_v) > \frac{1}{2}$, i.e., if conditionally on the information, the state 1 is more likely. Similarly, $a_v = 0$ if $\mathbb{P}(\theta = 1 \mid I_v) < \frac{1}{2}$. It may look as if equilibria can only differ in tie-braking; however, it is not the whole truth since $\mathbb{P}(\theta = 1 \mid I_v)$ depends on equilibrium strategies of other agents, which, in turn, are optimal replies to the strategies of others, including v. So, without additional assumptions, the equilibrium lacks the sequential structure since no pair of agents $v, u \in V$ such that $vu \notin E$ know who of them acted first.

Assumption ($\bigstar$) ensures the sequential structure of equilibrium. If $V_{<v} = \emptyset$, each agent v follows his signal. If $V_{<v}$ is non-empty with $|V_{<v}| = k$, the equilibrium strategy $\sigma_v\big(s_v, (a_u)_{u \in F_{<v}}, V_{<v}\big)$ is an optimal reply to $(\sigma_u)_{u \in V_{<v}}$ with $|V_{<u}| \leq k-1$ (since $V_{<u}$ is a strict subset of $V_{<v}$). This sequential structure implies that an agent gains no advantage from learning the set of agents that have not yet arrived, the property we need in Section 4.

2.2 Learning

Definition 2.1. For a network $G = (V, E)$ and an equilibrium σ, the *learning quality* $L_\sigma(G)$ is the expected fraction of agents making the correct action: $L_\sigma(G) = \frac{1}{|V|} \cdot \mathbb{E}_\sigma \big|\{v \in V \,:\, a_v = \theta\}\big|$.

We denote the learning quality for the best equilibrium σ by $L(G) = \max_\sigma L_\sigma(G)$. The maximum is attained since the set of equilibria in a finite game is closed.

Example 2.1 (Herding prevents learning)**.** Let G_n be the n-clique, the complete graph on n vertices. In this case, every agent v observes actions of all those who came earlier. If the first two agents take the incorrect action[2] $a = 1 - \theta$, then the third agent will ignore his signal repeating this wrong action (since the chance that both his predecessors are wrong is lower than the chance of him getting the wrong signal), and so on. This phenomenon known as herding ruins information aggregation. Indeed, with probability of at least $(1-\varepsilon)^2$ all the agents take the incorrect action. Thus $L(G_n) \leq 1-(1-\varepsilon)^2$, i.e., the learning quality is bounded away from 1 even for large cliques.

Given an assignment of arrival times, we say that an agent v is *isolated* if $F_{<v} = \emptyset$: he takes his action a_v without observing anybody else, so the action is dictated solely by the signal s_v. Such agents are sometimes called "guinea pigs" [Sgroi, 2002]. We identify the set of isolated agents with the "content generators" from the two-step information flow model since every such agent contributes information to the system. Actions of isolated agents are independent, therefore, the law of large numbers dictates that an agent observing a large enough sample of isolated agents can determine the state with high probability. If such an agent, in his turn, is observed by many others the learning propagates as suggested by the theory of two-step information flow.

This insight underlies the results from [Bahar et al., 2020]. They consider a similar model of learning with random arrivals and ask whether there exist social networks with the learning quality arbitrary close to one. They provide affirmative answer to this question by identifying a family of *celebrity graphs*, which is the only known family of networks with this property. We will construct another family in Section 5.

Example 2.2 (A minority of celebrities facilitates learning)**.** Consider a two-tier society, where a large set of k "commoners" observe a big set of $m \ll k$ "celebrities." The corresponding social network is a complete bipartite graph $B_{k,m}$. When a first celebrity arrives, he observes on average a set of $\frac{k}{m+1}$ isolated commoners. Hence, if $\frac{k}{m+1} \gg 1$, one can expect that each celebrity aggregates the information from these isolated agents an thereby takes the correct action with high probability. All commoners except a negligible fraction of $\frac{1}{m+1}$ observe at least one celebrity on average. Consequently, the overwhelming majority of commoners takes the correct action as well.

[2]The fact that the first two agents are informative does not satisfy in any equilibrium, nevertheless, it is easy to see that $L(G_n) \leq 1-(1-\varepsilon)^2$ holds true in every equilibrium.

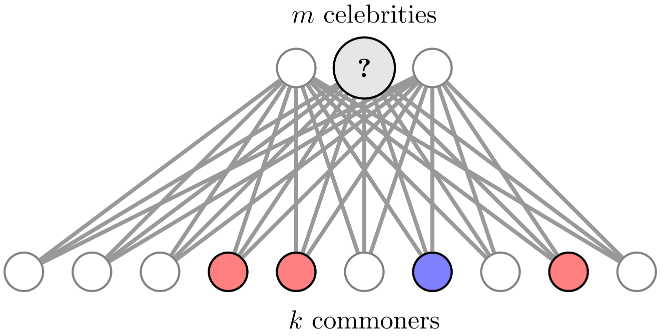

Figure 1: The celebrities graph. When the first celebrity arrives, it typically observes $\frac{k}{m} \gg 1$ commoners who made their decisions in isolation. These decisions match the signals and, hence, the celebrity learns from a number of independent sources. As a result, it takes a well-informed action, and then this action propagates.

This informal reasoning suggests that for any $\delta > 0$ we can find k and m such that the learning quality is at least $1 - \delta$. The formal argument in [Bahar et al., 2020] is tricky since a celebrity must "guess" which observed commoners are isolated and which are probably not.

Observation 2.1 (Random arrival order is crucial)**.** For cliques from Example 2.1, the assumption of random arrival order is inconsequential. However, for arbitrary graphs, such as celebrity graphs of Example 2.2, this assumption is critical. To see this one can consider a celebrity graph, where all commoners arrive before celebrities and thus all of them make a decision in isolation and learning fails.

2.3 Our question

The theory of two-step information flow and Example 2.2 suggest that social learning has two properties:

- Most of agents learn the state "second-hand". They do not observe content-generators (isolated agents) directly but observe intermediaries (celebrities in Example 2.2).
- Learning is facilitated by a minority of agents (e.g., celebrities in Example 2.2 or influentials in the context of the two-step information flow model). Once these agents are deleted from the network, learning does not persist.

Our goal is to understand whether these two properties are indeed immanent features of social learning.

3 Indirect learning

Example 2.2 demonstrates how content-generators are critical for learning over social networks. This is true for most, if not all, of the variants of models of information aggregation over such networks, where observing a large number of such agents is instrumental for learning (see, e.g., [Sgroi, 2002, Acemoglu et al., 2010, Arieli and Mueller-Frank, 2018]). In this section we present a simple argument showing that in our model, learning relies on intermediaries in the sense that the number of isolated agents that is observed directly is at most one on average![3] By "average" we mean both taking expectation with respect to all randomness and averaging over all agents.

Observation 3.1. On average, an agent observes at most one isolated agent directly, i.e.,

$$\frac{1}{|V|} \cdot \mathbb{E}\left(\sum_{v \in V} \left|\{u \in F_{<v} : \ F_{<u} = \emptyset\}\right|\right) \leq 1.$$

Proof. The probability that v is isolated equals $\frac{1}{\deg_v + 1}$, the probability that in a random ordering, involving v and his $\deg_v$ friends, v is the first. Each isolated v is observed by $\deg_v$ friends, therefore, the expected total number of observations of isolated agents is

$$\sum_{v \in V} \frac{\deg_v}{\deg_v + 1} \leq |V|.$$

Thus, on average, there is at most one observation per agent. □

The above reasoning shows that the only possibility for an agent to learn from many content-generators is to observe them "indirectly" through intermediaries. This conclusion provides a theoretical foundation for the first component of the two-step information flow theory.

4 No random opinion leaders

In the celebrities graph from Example 2.2, celebrities are the opinion leaders; deleting this minority from the graph leaves it totally disconnected and learning fails. However, if we eliminate a large set of agents randomly, say 50%, around 50% of celebrities will remain in the network ensuring the information aggregation (with slightly lower learning quality).

Here we show that this robustness of learning to random elimination of groups of agents is a general property. Even if a substantial fraction of agents leaves the network, the remaining agents find a way to learn the state even though most paths of information diffusion (those involving eliminated agents) disappear. Robustness to adversarial elimination is discussed in Section 6.

For a graph $G = (V, E)$ and a subset $V' \subset V$, let $G^{V'}$ be the induced subgraph: $G^{V'} = (V', E')$, where $E' = \{uv \in E : u, v \in V'\}$. The integer part of a real number x is denoted by $\lfloor x \rfloor$.

Theorem 4.1 (Learning is robust to random-subsets elimination)**.** *Consider a network $G = (V, E)$ with the learning quality $L(G) = 1 - \delta$ for some $\delta > 0$.*

Fix $\alpha \in (0, 1)$ and pick a subset $V' \subset V$ with $\lfloor \alpha \cdot |V| \rfloor$ agents uniformly at random. Under the assumption ($\bigstar$), the learning quality for the induced subnetwork $G^{V'}$ enjoys the lower bound $L(G^{V'}) \geq 1 - \sqrt{\frac{\delta}{\alpha}}$ with probability at least $1 - \sqrt{\frac{\delta}{\alpha}}$.

[3] A broader definition of a content-generator would be an agent whose action is determined by his signal: flipping this signal, while keeping actions of observed friends unchanged, flips the action. Intuitively, non-isolated agents of high-degree are almost never content-generators, and an analog of Observation 3.1 must hold for this broader definition. Despite a significant effort, we were unable to prove it.

Proof of Theorem 4.1. The argument is based on a coupling of the arriving process for the original network G and the selection of the random subnetwork $G^{V'}$. Fix an equilibrium $\sigma = (\sigma_v)_{v\in V}$ maximizing $L_\sigma(G)$ and pick a subset V' to be the set of $\lfloor \alpha|V| \rfloor$ earliest arrivers.

For $v \in V'$, the equilibrium strategy σ_v in the original network G can be used as a strategy in $G^{V'}$. The resulting family of strategies $(\sigma_v)_{v\in V'}$ constitutes an equilibrium in $G^{V'}$, which we denote by $\sigma^{V'}$. Note that here we use the assumption ($\bigstar$).[4]

The constructed coupling allows to link the learning quality for G and for $G^{V'}$ under equilibria σ and $\sigma^{V'}$, respectively. By the formula of total probability, the learning quality for G can be represented as

$$L(G) = \frac{1}{|V|}\sum_{v\in V}\Big(\mathbb{P}(a_v = \theta \mid v \in V')\cdot \mathbb{P}(v\in V') + \mathbb{P}(a_v = \theta \mid v \notin V')\cdot\mathbb{P}(v \notin V')\Big).$$

Using rough estimate $\mathbb{P}(a_v = \theta \mid v \notin V') \le 1$ on the probability of the correct action outside V' and taking into account that $\mathbb{P}(v \in V')$ is bounded from below by α, we get the following inequality

$$L_\sigma(G) \le \alpha \cdot \mathbb{E}_{V'} L_{\sigma^{V'}}\left(G^{V'}\right) + (1-\alpha),$$

where $\mathbb{E}_{V'}$ denotes expectation with respect to the choice of V'. Since the left-hand side is equal to $1-\delta$, we obtain

$$1 - \mathbb{E}_{V'} L_{\sigma^{V'}}\left(G^{V'}\right) \le \frac{\delta}{\alpha}.$$

Application of the Markov inequality completes the proof. □

5 Learning in egalitarian societies

While Theorem 4.1 from the previous section shows that most of the agents are unimportant for social learning and can be eliminated without significantly spoiling information aggregation, it leaves the possibility that there is always a certain negligible minority of agents, predetermined by the structure of social network (like celebrities in Example 2.2) that is responsible for effective information aggregation. In this section, we show that such minority need not exist and learning is possible even in totally egalitarian societies, where any two agents in the network are indistinguishable.[5]

Definition 5.1. A network $G = (V, E)$ is symmetric if for any pair of $v, v' \in V$ there exists a bijection $f : V \to V$ such that $f(v) = v'$ and for any $u, w \in V$ the edge uw belongs to E if and only if $f(u)f(w)$ is in E.

Theorem 5.1. *For any $\varepsilon_0 < \frac{1}{2}$ and any $\delta > 0$ there exists a symmetric network $G = (V, E)$ such that the learning quality $L_\sigma(G) \ge 1-\delta$ for any probability of a wrong signal $\varepsilon < \varepsilon_0$ and any state-symmetric equilibrium σ.*

The theorem extends to non-state-symmetric equilibria at the cost of assuming that signals are informative enough, namely $\varepsilon_0 < \frac{1}{4}$. This and other extensions are discussed in Section 7.

[4] In the game played over V' agents know that those from $V \setminus V'$ are absent. In particular, without assumption ($\bigstar$) an agent v in V' would know more about his predecessors than the same agent in the game played over V. Thus, a best reply within V may no longer be a best reply in the game restricted to V', even if all others maintained their strategy.

[5] Results of this section do not require the assumption ($\bigstar$). However, both results and proofs remain unchanged under ($\bigstar$) or even more general observation structures, see the discussion in Section 7.

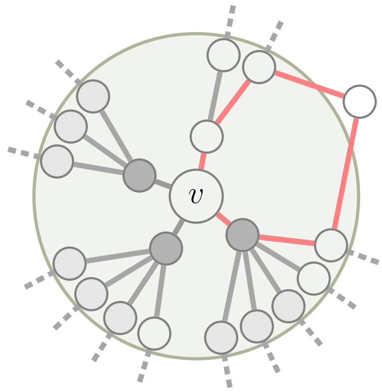

Figure 2: Agent v has $d = 3$ friends with degree 3 and higher. The 2-neighborhood of v (the shaded area) is a tree despite the red cycle. The maximal degree in the 2-neighborhood is 5. Hence the graph around v is characterized by the triplet $(d, r, D) = (3, 2, 5)$.

Theorem 5.1 relies on Proposition 5.1 stated in the next subsection. This proposition provides a useful sufficient condition for an agent to guess the state with high probability: the agent must have many friends with high degree and there must be no short cycles in the network. We prove Theorem 5.1 in Subsection 5.2 by combining Proposition 5.1 with classic results from the theory of expanders.

5.1 Sufficient condition for high-quality decisions: many high-degree friends and no short cycles

Recall some graph-theoretic notation. A *path of length* k in a graph $G = (V, E)$ is a sequence of vertices $(v_0, v_1, \dots, v_k)$ such that all edges $v_i v_{i+1}$, $i = 0, 1, \dots, k$, belong to E. The distance $d(v, v')$ between two vertices $v, v' \in V$ is defined as the length of the shortest path connecting them ($d(v, v') = +\infty$ if no such path exists). The r*-neighborhood* $B_r(v)$ of a vertex v is the set of all vertices u with $d(v, u) \leq r$; abusing the notation, we also refer to the induced subnetwork of G spanned by $B_r(v)$ as the r-neighborhood of v. A *cycle of length* k is a path with $v_0 = v_k$ such that all vertices $v_0, v_1, \dots, v_{k-1}$ are distinct (sometimes cycles with no repetitions are called simple). A graph is a tree if it has no cycles.

We capture a local structure of the network around an agent v by a triplet (d, r, D) such that

- agent v has at least d friends with the degree d or higher;
- the r-neighborhood of v is a tree;
- each agent from the r-neighborhood has at most D friends in G.

This definition is illustrated by Figure 2.

The local quality of decisions turns out to be related to this local structure of the network.

Proposition 5.1. *For any probability of a wrong signal* ε *and any state-symmetric equilibrium, the chance that an agent* v *takes the correct action satisfies*

$$\mathbb{P}(a_v = \theta) \geq 1 - \delta(\varepsilon, d, r, D), \tag{5.1}$$

where

$$\delta(\varepsilon, d, r, D) = \psi + \frac{18}{\sqrt{d-1}\,(1-\psi-2\varepsilon)}, \qquad \textit{and} \quad \psi = r \cdot \left(\frac{e \cdot D}{r}\right)^r.$$

Corollary 5.1. The parameter $\delta(\varepsilon, d, r, D)$ tends to zero whenever all elements of the triplet (d, r, D) go to infinity in a way that $\liminf \frac{r}{D} < e$. This means that an agent can guess the state if he has many high-degree friends and a tree-like neighborhood that is big enough compared to the maximal degree.

The intuition behind Proposition 5.1 is simple: absence of short cycles ensures that there is no herding among v's friends and thus v learns the state by observing a large sample of independent sources of information. However, formalization of this intuition faces some obstacles: the independence is only conditional and the conditioning is on a family of events, which does not belong to the information set of any of the agents; hence, checking that the independent sources are reliable requires the approximation of the condition by elements of information partitions; we are able to carry out this approximation only for high-degree friends of v since their partitions are finer.

We provide a sketch of Proposition 5.1; all the details are contained in Appendix B.

Sketch of the proof of Proposition 5.1 (see Appendix B for the details). Consider the following deviation a'_v from an equilibrium strategy a_v of agent v. He repeats the action played by the majority of his arrived friends with high degree $F^d_{<v} = \{u \in F_{<v} : \deg_u \geq d\}$. In equilibrium, the agent cannot benefit from this deviation and therefore $\mathbb{P}(a_v = \theta) \geq \mathbb{P}(a'_v = \theta)$.

The probability of mistake for a'_v can be bounded using the Hoeffding inequality.[6] This requires two ingredients: independence of actions a_u, $u \in F^d_{<v}$ and the bound on the individual probability of mistake for a_u.

The actions are not independent; however, they can be made independent by conditioning on a certain collection of events.

The most important one is the event $W_{<v}$ that the realized subnetwork of v is contained in his $(r-1)$-neighborhood. Since this neighborhood is a tree, all the realized graphs of $u \in F^d_{<v}$ do not intersect and thus each u aggregates the information from disjoint families of independent sources conditional on $W_{<v}$. The event $W_{<v}$ has high probability: $\mathbb{P}(W_{<v}) \geq 1 - \psi$, where ψ is given in the statement of the proposition. The intuition behind this bound is simple. The chance that a particular path $(z_0, z_1, \ldots z_{r-1}, v)$ of length r is contained in the realized subnetwork is extremely low for big r: it is equal to $\frac{1}{r!}$, the probability that all these agents arrived in exactly this order. Since r-neighborhood is a tree with the maximal degree D, the total number of such paths is at most $(D)^r$ and the result follows from the union bound.

Rather counterintuitively, there are other sources of dependence despite intersecting realized graphs. For example, the fact that we are interested in agents that arrive *before* v creates dependence between their actions even if their realized subnetworks are disjoint: when v arrives early, all the friends he observes are also early-comers and, hence, are likely to have smaller realized graphs compared to the case when v arrives late; in particular, the earlier v arrives, the more mistakes the observed friends make thus creating the dependence.

To cut all the sources of dependence, we end up conditioning to the arrival time of v, his arrived friends, the realized state, and the event $W_{<v}$.

[6] The Hoeffding inequality [Hoeffding, 1994] states that $\mathbb{P}\left(\frac{1}{N}\sum_{n=1}^{N}\xi_i \leq \mathbb{E}\left(\frac{1}{N}\sum_{n=1}^{N}\xi_i\right) - x\right) \leq \exp\left(-2x^2N\right)$ for independent random variables $0 \leq \xi_i \leq 1$ and any $x \geq 0$.

In order to apply the Hoeffding inequality, it remains to show that the individual conditional probability of mistake for a_u is bounded away from $1/2$ for $u \in F^d_{<v}$, i.e., the independent sources observed by v are informative. We use the following idea. For any event A from the information partition of u, we have $\mathbb{P}(a_u \neq \theta \mid A) \leq \varepsilon$ because otherwise u is better off following his signal whenever A occurs; additional conditioning on the realized state does not change the bound because we assume a state-symmetric equilibrium. Unfortunately, the family of events that we are conditioning at does not belong to u's information partition. We overcome this difficulty approximating the condition by elements of the information partition and show that the conditional probability of mistake is at most $\frac{1}{1-\psi}\left(\varepsilon + \frac{3}{t_v \cdot \sqrt{\deg_u - 1}}\right)$. The bound gets worse for low-degree agents since their information partition is not fine enough for good approximation. This is why the deviation a'_v of agent v takes into account high-degree friends only.

After these preparations, Proposition 5.1 becomes a corollary of the Hoeffding inequality. □

5.2 Proof of Theorem 5.1: D-regular graphs with no short cycles

The *girth* $g = g(G)$ of a graph G is the length of a shortest cycle; $g(G) = +\infty$ for a tree. The girth of a graph induces a local property on each vertex, namely, in a high-girth graph each vertex spans a local tree.

Observation 5.1. Let G be a graph with girth g. Then for radius $r \leq \frac{g}{2} - 1$, the r-neighborhood of any agent v is a tree.[7]

Recall that a graph is called *D-regular* if all vertices have degree D. For a D-regular graph of girth g, Proposition 5.1 with Observation 5.1 imply a lower bound on the learning quality

$$L_\sigma(G) = \frac{1}{|V|} \cdot \sum_{v \in V} \mathbb{P}(a_v = \theta) \geq 1 - \delta\left(\varepsilon, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right). \tag{5.2}$$

for any state-symmetric σ. Theorem 5.1 follows from the existence of symmetric D-regular graphs with arbitrary high degree and girth. Existence of such graphs is itself a non-trivial question, which is resolved in the theory of expanders, see Appendix A. The following is the corollary of a theorem by Lubotzky et al. [1988] (Theorem A.1 from Appendix A) that describes a family of the so-called Ramanujan expanders.

Corollary 5.2 (of Theorem A.1 by Lubotzky et al. [1988])**.** There is a sequence $D_k \to \infty$ such that for any g_0 and k, there exists a symmetric D_k-regular graph $G = (V, E)$ such that $g(G) \geq g_0$ and $|\lambda_2| \leq 2\sqrt{D_k - 1}$, where λ_2 is the second-largest eigenvalue of the adjacency matrix.[8]

Corollary 5.2 allows to pick a symmetric D-regular graph G with arbitrary high degree D and arbitrary high girth/degree ratio $\frac{g}{D}$, while Corollary 5.1 implies that $\delta\left(\varepsilon, D, \lfloor \frac{g}{2} - 1 \rfloor, D\right)$ tends to zero if both the girth g and the degree D tend to infinity and the girth goes to infinity faster: $\frac{g}{D} \to \infty$. Hence, for any given δ and ε_0, there is a symmetric D-regular graph with D and g such that $\delta(\varepsilon_0, D, g, D) \leq \delta$.

[7] If r-neighborhood of v is not a tree, it has some cycle C. Let u be a vertex from C farthest from v. Consider the following path P: it starts from v, goes optimally to u, makes one step along the cycle to some not yet visited u' and comes back to v optimally. By the choice of u, the vertex u' cannot be a part of the shortest path from v to u. Hence, P contains a cycle C'. The length of P is at most $r + 1 + r = 2r + 1$; it gives an upper bound on the length of C'. Therefore, the girth satisfies $g \leq 2r + 1$ or, equivalently, $\frac{g}{2} - 1 < r$.

[8] We use this bound on $|\lambda_2|$ in Section 6.

The lower bound (5.2) implies that for any probability of a wrong signal $\varepsilon < \varepsilon_0$, the learning quality for G is at least

$$1 - \delta\left(\varepsilon, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right), \quad \text{where} \quad \delta\left(\varepsilon, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right) \leq \delta\left(\varepsilon_0, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right) \leq \delta,$$

which completes the proof of Theorem 5.1. □

6 Learning without opinion leaders

In Section 5, we saw that learning is possible in symmetric sparse networks that have high degrees but no short cycles. In such networks, all agents play the same role, which makes it natural to expect that no *small group* of agents is critical, in particular, adversarial elimination of a small group cannot spoil the learning outcome.

In this section we obtain a surprising much stronger result demonstrating that there are symmetric networks, where leaning is robust to adversarial elimination of any group, possibly, large.[9]

Theorem 6.1. *For any $\varepsilon_0 < \frac{1}{2}$ and any $\delta > 0$ there exists a symmetric network $G = (V, E)$ such that for any $0 < \alpha \leq 1$ and any subset $V' \subset V$ with $\lfloor \alpha \cdot |V| \rfloor$ agents, for the induced subnetwork $G^{V'}$ spanned by V' the learning quality is at least $1 - \frac{\delta}{\alpha^3}$ for any state-symmetric equilibrium and any probability of a wrong signal $\varepsilon < \varepsilon_0$.*

Similarly to Section 5, all the results of this section can be extended to non-symmetric equilibria. However, the analog of Theorem 6.1 requires $\varepsilon_0 < \frac{1}{4}$ in this case.

It would seem that Theorem 6.1 follows as a corollary from Theorems 4.1 and 5.1. Theorem 5.1 claims the existence of a symmetric networks that sustains social learning. Due to symmetry arguments one could say that if there exists a subset whose deletions spoils learning then the deletion of any subset does so as well, thus contradicting Theorem 4.1. Unfortunately, this argument is incorrect as the symmetry of the network does not lead to symmetry of its subnetworks and, in particular, does not induce equality on the learning qualities of the networks after the deletion. Two different subsets may have completely different implications when deleted even though the initial network is symmetric.

Robustness of learning quality to adversarial elimination turns out to be related to spectral properties of the network. This is captured by Proposition 6.1 below, and Theorem 6.1 easily follows from this proposition and known results on expander graphs.

Consider a D-regular graph and denote by $|\lambda_1| \geq |\lambda_2| \geq \ldots \geq |\lambda_{|V|}|$ eigenvalues of its adjacency matrix ordered by their absolute values. Expanders are graphs with small $|\lambda_2|$ relative to $|\lambda_1|$, see Appendix A.

The next proposition bounds by how much the average number of mistakes can increase if instead of the original D-regular network, we consider its arbitrary subnetwork with $\lfloor \alpha \cdot |V| \rfloor$ agents. It can be seen as a generalization of Proposition 5.1 in D-regular case.

Proposition 6.1. *Let $G = (V, E)$ be a D-regular graph with girth g and the second largest eigenvalue λ_2. Then for any $\alpha \in (0, 1]$ and any subset $V' \subset V$ of size $|V'| = \lfloor \alpha \cdot |V| \rfloor$, the learning quality in the induced subnetwork $G^{V'}$ spanned by V' satisfies*

$$L_{\sigma'}\left(G^{V'}\right) \geq 1 - \left(\frac{2}{\sqrt{\alpha}} + \frac{(1-\alpha)|\lambda_2|^2}{\alpha^3 \cdot D^{\frac{3}{2}}}\right) \cdot \delta\left(\varepsilon, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right) \qquad (6.1)$$

[9] As in Section 5, the results of this section do not require the assumption (★) and hold both with or without it. They also don't change if an agent learns any additional information about members of his realized subnetwork.

for any state-symmetric equilibrium σ'. Here $\delta(\varepsilon, d, r, D)$ is given by formula (5.1)*.*

Proposition 6.1 is proved in Appendix C; here we present the main idea and then prove Theorem 6.1.

Idea of the proof of Proposition 6.1. The key tool is the mixing lemma from the theory of expanders (Lemma **??**). This lemma is applicable to any D-regular graph and provides a bound in terms of $|\lambda_2|$ on how much the number of edges $|E(V^1, V^2)|$ between any two disjoint subsets of vertices V_1, V_2 deviates from the expected number of edges in the Erdos-Renyi model:

$$\left| |E(V^1, V^2)| - \frac{D}{|V|} \cdot |V^1| \cdot |V^2| \right| \leq |\lambda_2| \sqrt{|V^1||V^2|}. \tag{6.2}$$

This lemma allows to bound the fraction of agents with low degree in $G^{V'}$. Indeed, for $\gamma < \alpha$, let V^1 be the set of agents with degree less than $\gamma \cdot D$ in $G^{V'}$ and V^2 be the set of eliminated agents $V \setminus V'$. Since each agent has degree D in the original graph, there are at least $(1-\gamma) D|V^1|$ edges between V^1 and V^2. The Erdos-Renyi model prescribes $(1-\alpha)D|V^1|$ edges in expectation. This discrepancy is compatible with (6.2) only for relatively small subsets V^1. We obtain that the fraction of agents in $G^{V'}$ with degree $\gamma \cdot D$ or higher is at least $1-\theta$, where $\theta = \theta(\alpha, D, \gamma, \lambda_2)$ is small.

By Lemma C.1 from the appendix, the lower bound on the fraction of high-degree agents implies a bound on the number of agents having high number of high-degree friends: the fraction of agents having at least $\frac{\gamma}{2}D$ friends with degrees $\frac{\gamma}{2}D$ and higher is above $1-2\theta$. Application of formula (5.1) from Proposition 5.1 on this set completes the proof. □

Proof of Theorem 6.1. As in the proof of Theorem 5.1, we can pick the Ramanujan expander such that $\delta\left(\varepsilon, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right)$ is at most $\frac{\delta}{6}$. The second eigenvalue of the Ramanujan expander satisfies $|\lambda_2| \leq 2\sqrt{D-1} \leq 2\sqrt{D}$ (Corollary 5.2) and thus, by Proposition 6.1, any sub-network $G^{V'}$ with $|V'| = \lfloor \alpha \cdot |V| \rfloor$ has learning quality at least

$$1 - \left(\frac{2}{\alpha} + \frac{4(1-\alpha)}{\alpha^3 \sqrt{D}} \right) \cdot \delta\left(\varepsilon, D, \left\lfloor \frac{g}{2} - 1 \right\rfloor, D\right).$$

Since $\frac{2}{\alpha} + \frac{4(1-\alpha)}{\alpha^3\sqrt{D}} \leq \frac{2}{\alpha^3} + \frac{4}{\alpha^3} \leq \frac{6}{\alpha^3}$, the learning quality is bounded from below by $1 - \frac{\delta}{\alpha^3}$. □

7 Getting rid of technical assumptions

Some of our modeling assumptions were made to simplify exposition. Relaxing the following assumptions will not alter our results and the conclusions:

General distributions of arrival times. We assumed that agents' arrival time are i.i.d. and, in particular, uniformly distributed on the unit interval. However, any non-atomic distribution leads to an equivalent model (equivalence is obtained by a monotone reparameterization), as long as the i.i.d. assumption is maintained. An important robustness result would be to extend our conclusions to some approximate notion of i.i.d. as some local dependence among agents is a realistic assumption.

Non-binary signals Our results hold for any non-binary signaling device as long as signals are informative and symmetry is maintained. By informativeness we mean that for a positive probability set of signals, s, the posteriors $\mathbb{P}(\theta = 1 \mid s)$ belong to the union of intervals $[0, \varepsilon] \cup [1 - \varepsilon, 1]$ (the probability of this set of signals will enter into the bound on the learning quality). By symmetry we mean the distribution posteriors $\mathbb{P}(\theta = 1 \mid s)$ is symmetric with respect to $\frac{1}{2}$ (note that signals of unbounded precision are also allowed).

Heterogeneous agents Agents can also be heterogeneous, i.e., the signaling device can be agent-specific and so each agent v has his own signal precision ε_v. In this case, all the results hold with $\varepsilon = \max_{v \in V} \varepsilon_v$.

Non-symmetric equilibria The symmetry assumption on equilibria and signals can be relaxed. However, for non-symmetric equilibria, Theorems 5.1 and 6.1 require the probability of a wrong signal to be below $\frac{1}{4}$ (instead of $\frac{1}{2}$). The reason for that is the inequality (B.1), where for non-symmetric equilibria we get 2ε in the numerator.

General states Extension to non-equiprobable states is straightforward. Note, however, that this breaks the state-symmetry of equilibria and therefore we get 2ε in (B.1) (see the comment above). Extension to a non-binary state does not lead to any technical difficulties.

More informed agents Results of Sections 5 and 6 hold unchanged if in addition to observing actions of friends each agent v gets some information about his realized subnetwork $G_{<v}$, e.g., a possibly noisy signal about the set of members of the realized subnetwork, their arrival times, actions, and even about their private signals. In particular, these results are independent of the assumption ($\bigstar$) and hold both with or without it. This assumption is used only in Theorem 4.1. There it helps to show that an equilibrium for the original network induces a certain equilibrium in a random sub-network. We don't know how to relax this assumption, but believe that the theorem remains valid without it.

8 Summary

It is quite straightforward that when signals are bounded social learning requires many isolated agents that establish a large enough sample to point at the correct action. Agents must observe such samples and so a necessary condition for learning is that networks are sufficiently connected. What we show is that most of these observations must take place indirectly. This, in turn, means that the network may not be too connected.

We then go on and ask whether minorities play a crucial role for social learning in the sense that learning does not obtain in their absence. Although the existence of such critical minorities cannot be ruled out for some networks, it turns out that they need not necessarily exist. In order to design a network with high learning quality, it is enough to ensure that agents have many friends (so they learn from many sources) and there are no short cycles in the network (this prevents herding). Both properties are robust to eliminating groups of agents: this creates no new cycles and most of remaining agents have high degree. This suggests that learning on networks with high degrees and large girth must be robust to elimination of any minority.

A prototypical example of high-girth networks is given by the Ramanujan expanders and our paper contributes one more item to the list of surprising properties of these networks. We showed that expanders have the exceptional combination of symmetry with learning and robustness properties: ex-ante (i.e., before the realization of arrival times) all agents are symmetric (i.e., there is no selected minority) and, even if 99% agents are deleted from the network in an adversarial way, remaining agents learn the state.

References

Daron Acemoglu, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *Review of Economic Studies*, 78:1–34, 2010.

Noga Alon and Fan RK Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988.

Itai Arieli and Manuel Mueller-Frank. Multidimensional social learning. *Review of Economic Studies*, 0:1–28, 2018.

Gal Bahar, Itai Arieli, Rann Smorodinsky, and Moshe Tennenholtz. Multi-issue social learning. *Mathematical Social Sciences*, 104:29–39, 2020.

Abhijit V Banerjee. A simple model of herd behavior. *The quarterly journal of economics*, 107(3): 797–817, 1992.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom and cultural change as information cascade. *The Journal of Political Economy*, 100:992–1026, 1992.

Jeffery L Bineham. A historical account of the hypodermic model in mass communication. *Communications Monographs*, 55(3):230–246, 1988.

Kenny K Chan and Shekhar Misra. Characteristics of the opinion leader: A new dimension. *Journal of advertising*, 19(3):53–60, 1990.

Robin A Coulter, Lawrence F Feick, and Linda L Price. Changing faces: Cosmetics opinion leadership among women in the new hungary. *European Journal of Marketing*, 36(11/12):1287–1308, 2002.

Xavier Dahan. Regular graphs of large girth and arbitrary degree. *Combinatorica*, 34(4):407–426, 2014.

Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121, 1974.

Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of political Economy*, 101(4):612–643, 1993.

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)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.

Benjamin Golub and Matthew O Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–49, 2010.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.

Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

Elihu Katz and Paul F Lazarsfeld. Personal influence: the part played by people in the flow of mass communications. 1955.

Paul Felix Lazarsfeld, Bernard Berelson, and Hazel Gaudet. The people's choice. 1944.

Alexander Lubotzky. Expander graphs in pure and applied mathematics. *Bulletin of the American Mathematical Society*, 49(1):113–162, 2012.

Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3): 261–277, 1988.

Moshe Morgenstern. Existence and explicit constructions of q+ 1 regular ramanujan graphs for every prime power q. *Journal of Combinatorial Theory, Series B*, 62(1):44–62, 1994.

Elchanan Mossel, Joe Neeman, and Omer Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3):408–429, 2014.

Elchanan Mossel, Manuel Mueller-Frank, Allan Sly, and Omer Tamuz. Social learning equilibria. *Available at SSRN 3124385*, 2019.

James H Myers and Thomas S Robertson. Dimensions of opinion leadership. *Journal of marketing research*, 9(1):41–46, 1972.

Everett M Rogers. Diffusion of innovations: modifications of a model for telecommunications. In *Die diffusion von innovationen in der telekommunikation*, pages 25–38. Springer, 1995.

Daniel Sgroi. Optimizing information in the herd: Guinea pigs, profits, and welfare. *Games and Economic Behavior*, 39:137–166, 2002.

Lones A Smith. *Essays on dynamic models of equilibrium and learning.* PhD thesis, University of Chicago, Department of Economics, 1991.

Thomas W Valente and Everett M Rogers. The origins and development of the diffusion of innovations paradigm as an example of scientific growth. *Science communication*, 16(3):242–273, 1995.

Christophe Van den Bulte and Yogesh V Joshi. New product diffusion with influentials and imitators. *Marketing science*, 26(3):400–421, 2007.

Eric Vernette. Targeting women's clothing fashion opinion leaders in media planning: an application for magazines. *Journal of Advertising Research*, 44(1):90–107, 2004.

Duncan J Watts and Peter Sheridan Dodds. Influentials, networks, and public opinion formation. *Journal of consumer research*, 34(4):441–458, 2007.

Gabriel Weimann. *The influentials: People who influence people.* SUNY Press, 1994.

A Expanders

There are several equivalent definitions of expanders, see [Hoory et al., 2006]. The "spectral" definition is the most convenient for our needs. Consider a D-regular ($\deg_v = D$ for all vertices) graph $G = (V, E)$ and denote by $(\lambda_k)_{k=1,\dots,|V|}$ eigenvalues of its adjacency matrix ordered by absolute values $|\lambda_1| \geq |\lambda_2| \geq |\lambda_3|.... \geq |\lambda_{|V|}|$. The eigenvalue $\lambda_1 = D$ corresponds to the eigenvector representing the uniform distribution over vertices; the smaller is the second eigenvalue $|\lambda_2|$, the faster the distribution of a random walk started at some vertex converges to the uniform distribution, see Proposition 1.6. in [Lubotzky, 2012]. A graph G is an *expander* if $|\lambda_2|$ is small relative to $|\lambda_1|$, i.e., a random walk on G forgets the starting point fast.

Theorem A.1 (Lubotzky et al. [1988])**.** *For any N and D such that $D-1$ is a prime number and $D-2$ is divisible by 4, there exists a D-regular symmetric graph $G = (V, E)$ with at least N vertices, $|\lambda_2| \leq 2\sqrt{D-1}$, and girth $g \geq \frac{2}{3}\log_{D-1}|V|$.*

Graphs with $|\lambda_2| \leq 2\sqrt{D-1}$ are called Ramanujan graphs; by the Alon-Boppana theorem (see Theorem 5.3 in [Hoory et al., 2006]), this value of λ_2 is essentially the best possible. It is quite intuitive that the best expanders cannot have short cycles, which lead to recurrences in a random walk and hence slow down its expansion over the graph.

The proof of Theorem A.1 is quite technical and relies on group-theory. Lubotzky et al. [1988] and successors (e.g., Morgenstern [1994], who relaxed the condition of divisibility by 4, and Dahan [2014], who extended the result to arbitrary $D \geq 11$) construct G as a Cayley graph of a certain group.[10] While the symmetry of the constructed graph is not mentioned explicitly in these papers, it comes for free because any Cayley graph is symmetric, see Claim 11.4 in [Hoory et al., 2006].[11]

In addition to large girth, we need another property of expanders demonstrating their similarity to random graphs in the Erdos-Renyi model with the probability p of an edge between two given vertices equal to $\frac{D}{|V|}$. For a pair of disjoint subsets $V^1, V^2 \subset V$ denote by $E(V^1, V^2)$ the set of edges with one end-point in V^1 and another in V^2. In Erdos-Renyi model the expected number of such edges is equal to $p \cdot |V^1||V^2|$. The following result, known as mixing lemma, shows that $\left|E(V^1, V^2)\right|$ for an expander is close to this number.

Lemma A.1 (Mixing lemma, [Alon and Chung, 1988])**.** *For any D-regular graph $G = (V, E)$ and any two disjoint subsets $V^1, V^2 \subset V$ the following inequality holds*

$$\left|\left|E(V^1, V^2)\right| - \frac{D}{|V|} \cdot |V^1| \cdot |V^2|\right| \leq |\lambda_2|\sqrt{|V^1||V^2|}.$$

B Proof of Proposition 5.1

Denote by F_v^d the subset of v's friends having degree at least d; those of them who arrived earlier than v are denoted by $F_{<v}^d = F_v^d \cap F_{<v}$.

[10] Given a group $\mathcal{G}$ with group-operation $\circ$ and a subset $S \subset \mathcal{G}$, its Cayley graph is defined in the following way: the set of vertices $V = \mathcal{G}$ and $gg' \in E$ if $g' = g \circ s$ for some $s \in S$.

[11] Indeed, for any $x \in \mathcal{G}$ the map $g \to x \circ g$ defines an edge-preserving bijection: $gg' \in E$ if and only if $(x \circ g)(x \circ g') \in E$.

Consider the following deviation a'_v of agent v from his equilibrium action a_v. The agent decides to rely on the majority of his high-degree friends who arrived earlier: if the majority of $F^d_{<v}$ played action $a \in \{0,1\}$, agent v repeats this action; in case of a tie or v being isolated, $a_v = s_v$, i.e., v follows his signal. In equilibrium, no deviation is profitable and hence $\mathbb{P}(a_v = \theta) \geq \mathbb{P}(a'_v = \theta)$. The probability of mistake for a'_v can be estimated using the Hoeffding inequality (see Footnote 6).

In order to apply the Hoeffding inequality, we need independence of actions $(a_u)_{u \in F^d_{<v}}$ and an upper bound on the probability that $a_u \neq \theta$. To make the actions independent, we need conditioning on the appropriate family of events. The following lemmas describe the desired family of events and establish the bound.

Denote by $W_{<v}$ the event that the realized subnetwork of agent v is contained in his $(r-1)$-neighborhood. Note that, given $W_{<v}$, the subnetwork of v is a tree.

Lemma B.1. *Fix $\theta_0 \in \{0,1\}$, $t \in [0,1]$, and a subset of friends $F \subset F_v$ of agent v. The actions $(a_u)_{u \in F}$ are independent conditional on $\theta = \theta_0$, the arrival time $t_v = t$ of v, the set of observed friends $F_{<v} = F$, and $W_{<v}$.*

The second lemma ensures that $W_{<v}$ is a high-probability event if the radius r is large enough compared to the maximal degree.

Lemma B.2. *The probability that the realized subnetwork of v is contained in the $(r-1)$-neighborhood of v, conditional on arrival times of v and all his friends, enjoys the following lower bound*

$$\mathbb{P}\big(W_{<v} \mid t_v,\ (t_u)_{u \in F_v}\big) \geq 1 - \psi, \quad \textit{where} \quad \psi = r \cdot \left(\frac{e \cdot D}{r}\right)^r.$$

The third lemma provides an upper bound on the probability of a wrong action.

Lemma B.3. *For any subset F of v's friends, an agent $u \in F$, and any state-symmetric equilibrium, the conditional probability of a wrong action satisfies the following inequality*

$$\mathbb{P}(a_u \neq \theta \mid \theta = \theta_0,\ t_v = t,\ F_{<v} = F,\ W_{<v}) \leq \frac{1}{1-\psi}\left(\varepsilon + \frac{3}{t \cdot \sqrt{\deg_u - 1}}\right), \tag{B.1}$$

where ψ is from Lemma B.2.

The lemmas are proved below. With their help, we complete the proof of Proposition 5.1. By the Hoeffding inequality and Lemmas B.1 and B.3, the probability of a wrong action a'_v for the high state $(\theta = 1)$ is bounded as follows

$$\mathbb{P}\big(a'_v \neq \theta \mid \theta = 1,\ t_v,\ F_{<v},\ W_{<v}\big) \leq \mathbb{P}\left(\sum_{u \in F^d_{<v}} a_u \leq \frac{\left|F^d_{<v}\right|}{2} \;\middle|\; \theta = 1,\ t_v,\ F_{<v},\ W_{<v}\right) \leq \tag{B.2}$$

$$\leq \exp\left(-2\left(\frac{1}{2} - \frac{1}{1-\psi}\left(\varepsilon + \frac{3}{t_v \cdot \sqrt{d-1}}\right)\right)^2 \cdot \left|F^d_{<v}\right|\right) = \exp\left(-\frac{1}{2}\frac{\left(1 - \psi - 2\varepsilon - \frac{6}{t_v \cdot \sqrt{d-1}}\right)^2}{(1-\psi)^2} \cdot \left|F^d_{<v}\right|\right), \tag{B.3}$$

where we used that $\deg_u \geq d$ for any $u \in F^d_{<v}$. The bound (B.3) holds only if t_v is not too small: to apply the Hoeffding inequality, the expression in the parentheses, $\left(1-\psi-2\varepsilon-\frac{6}{t\cdot\sqrt{d-1}}\right)$, must be non-negative (see the requirement $x \geq 0$ in Footnote 6). We will use the bound (B.3) only for t_v such that this expression is at least $\frac{1-\psi-2\varepsilon}{2}$ or equivalently $t_v \geq \frac{12}{\sqrt{d-1}(1-\psi-2\varepsilon)}$. For small t_v, we roughly bound the probability by 1 and get the following inequality valid for all t_v

$$\mathbb{P}(a'_v \neq \theta \mid \theta = 1,\ t_v,\ F_{<v},\ W_{<v}) \leq \mathbb{1}_{\left\{t_v < \frac{12}{\sqrt{d-1}(1-\psi-2\varepsilon)}\right\}} + \exp\left(-\frac{1}{8}\frac{(1-\psi-2\varepsilon)^2}{(1-\psi)^2}\cdot\left|F^d_{<v}\right|\right), \tag{B.4}$$

where $\mathbb{1}_A$ denotes the indicator of an event A.

Let's get rid of conditioning to $\theta = 1,\ t_v,\ F_{<v}$, and $W_{<v}$. For $\theta = 0$, one similarly derives the same bound (B.4) and hence it holds unconditionally on θ. Since $\mathbb{P}(\cdot) \leq \mathbb{P}(\cdot \mid W_{<v}) + (1-\mathbb{P}(W_{<v}))$, we eliminate conditioning to $W_{<v}$ at the cost of increasing the upper bound by ψ. It remains to average the bound over $F^d_{<v}$ and t_v. The size $|F^d_{<v}|$ is uniformly distributed over $\{0, 1, \ldots, |F^d_v|\}$ because it is the length of the prefix of v in a random permutation of $F^d_v \cup \{v\}$. Thus, averaging the second summand in (B.4) results in the geometric progression of length $|F^d_v| \geq d$. Taking into account that t_v is uniformly distributed on $[0, 1]$, we get

$$\mathbb{P}(a'_v \neq \theta) \leq \psi + \frac{12}{\sqrt{d-1}(1-\psi-2\varepsilon)} + \frac{1}{d+1}\cdot\frac{1}{1-\exp\left(-\frac{1}{8}\frac{(1-\psi-2\varepsilon)^2}{(1-\psi)^2}\right)}. \tag{B.5}$$

This upper bound can be simplified if we note that $1-\exp(-x) \geq x\cdot\frac{1-\exp(-a)}{a}$ for $x \in [0, a]$. Since $\frac{1-\psi-2\varepsilon}{1-\psi} \leq 1$, we obtain

$$\frac{1}{1-\exp\left(-\frac{1}{8}\frac{(1-\psi-2\varepsilon)^2}{(1-\psi)^2}\right)} \leq \frac{8(1-\psi)^2}{\left(1-\exp\left(-\frac{1}{8}\right)\right)(1-\psi-2\varepsilon)^2} \leq \frac{72}{(1-\psi-2\varepsilon)^2},$$

where in the second inequality we estimated $(1-\psi)$ by one and used that $1-\exp\left(-\frac{1}{8}\right) \geq \frac{1}{9}$. Denoting by b and c the second and the third summands in (B.5), respectively, we see that $c \leq \frac{1}{2}b^2$. For b below one, we have $b^2 \leq b$, while for $b \geq 1$ the bound (B.5) becomes trivial anyway. Therefore,

$$\mathbb{P}(a'_v \neq \theta) \leq \psi + \frac{18}{\sqrt{d-1}(1-\psi-2\varepsilon)}.$$

Taking into account that $\mathbb{P}(a_v = \theta) \geq \mathbb{P}(a'_v = \theta) = 1 - \mathbb{P}(a'_v \neq \theta)$ and substituting the expression for ψ from Lemma B.2, we complete the proof of Proposition 5.1. $\square$

Proof of Lemma B.1: To ensure the conditional independence, we make use of the following property of a realized subnetwork. For any agent $w \in V$, his equilibrium action a_w is determined by the collection of signals s_z of agents z from w's realized subnetwork $G_{<w}$, while signals of other agents play no role. This property has an important consequence. Consider a group of agents $W \subset V$ and an event determined by the collection of arrival times $B = \left\{(t_y)_{y\in V} \subset T\right\}$ for some $T \subset [0, 1]^V$. Then actions $(a_w)_{w\in W}$ are independent conditional on $\theta = \theta_0$ and B provided that all the realized subnetworks $(G_{<w})_{w\in W}$ are independent disjoint random graphs conditional on B.

With this observation, it is easy to check that actions $(a_u)_{u\in F}$ are conditionally independent given $\theta = \theta_0,\ t_v = t,\ F_{<v} = F$, and $W_{<v}$. By the definition of $W_{<v}$, the realized subnetwork of v is a tree and

thus the subnetworks of his arrived friends $(G_{<u})_{u\in F}$ are disjoint. It remains to check independence of $(G_{<u})_{u\in F}$.

Since r-neighborhood of v is a tree, after eliminating agent v it splits into $\deg_v$ trees. Denote the subtree containing u by Π_u, and define its boundary and interior by $\partial\Pi_u = \Pi_u \cap \big(B_r(v) \setminus B_{r-1}(v)\big)$ and $\Pi_u^{\text{int}} = \Pi_u \setminus \partial\Pi_u$, respectively. The realized subnetwork of v belongs to $(r-1)$-neighborhood if and only if the realized subnetwork of each $u \in F_{<v}$ is in Π_u^{int}. Therefore, the event $W_{<v}$ can be represented as $\cap_{u\in F_{<v}}\{G_{<u} \subset \Pi_u^{\text{int}}\}$. Each event $\{G_{<u} \subset \Pi_u^{\text{int}}\}$ is determined by arrival times t_z of agents $z \in \Pi_u \cup \{v\}$. Indeed, $G_{<u}$ belongs to Π_u^{int} if and only if u arrives earlier than v and $G_{<u}$ has no paths connecting the boundary of Π_u and u, i.e., for any $z_0 \in \partial\Pi_u$ and any path $(z_0, z_1, z_2, ..., z_{r-1} = u)$ there is an index i such that $t_{z_i} > t_{z_{i+1}}$. Consequently, we can rewrite the event $\{G_{<u} \subset \Pi_u^{\text{int}}\}$ as $\{(t_z)_{z\in\Pi_u} \subset T_u(t_v)\}$, where $T_u(t_v)$ is a certain subset of $[0,1]^{\Pi_u}$ depending on the arrival time of v.

Conditioning on $F_{<v} = F$, $t_v = t$, and $W_{<v}$ becomes equivalent to conditioning on the following family of events determined by arrival times: $\{(t_z)_{z\in\Pi_u} \subset T_u(t)\}$ for each $u \in F$, $t_u > t$ for u outside F, and $t_v = t$. Arrival times $(t_y)_{y\in V}$ are unconditionally independent, and the condition restricts the values of disjoint subsets. Thus, conditionally on $F_{<v} = F$, $t_v = t$, and $W_{<v}$, the families of random variables $(t_z)_{z\in\Pi_u}$ are independent across $u \in F$. The realized subnetwork $G_{<u}$ of u is determined by $(t_z)_{z\in\Pi_u}$ and, therefore, the graphs $(G_{<u})_{u\in F}$ are also conditionally independent, which implies the desired conditional independence of actions $(a_u)_{u\in F}$. □

Proof of Lemma B.2. Consider an agent z such that the distance from v to z is equal to r, i.e., z is on the boundary of the r-neighborhood of v. Since this neighborhood is a tree, there is only one path $(z = z_0, z_1, ...z_r = v)$ connecting z and v. The chance that this path is presented in the realized subnetwork of v is bounded by the probability that $z_0, z_1, ..., z_{r-2}$ arrive exactly in this order (we excluded agents $z_r = v$ and z_{r-1} since we are conditioning on their arrival times); this probability is equal to $\frac{1}{(r-1)!}$. The total number of agents z on the boundary of r-neighborhood is at most $\big(D\big)^r$. Therefore, by the union bound, no boundary agent belongs to the realized subnetwork of v with probability at least $1 - \frac{\left(D\right)^r}{(r-1)!}$. Since no agent at the distance r is in $G_{<v}$, the realized subnetwork is contained in the $(r-1)$-neighborhood. Using the inequality $n! \geq \left(\frac{n}{e}\right)^n$, we can get rid of the factorial: $\frac{\left(D\right)^r}{(r-1)!} = r \cdot \frac{\left(D\right)^r}{r!} \leq r \cdot \left(\frac{e\cdot D}{r}\right)^r$. □

Proof of Lemma B.3: The action of an agent u is determined by signals of agents from his realized subnetwork $G_{<u}$. In the proof of Lemma B.1 we saw that, conditional on $F_{<v} = F$, $t_v = t$, and $W_{<v}$, the realized subnetwork of $u \in F$ is determined by arrival times $(t_z)_{z\in\Pi_u}$, and the condition is equivalent to $\{(t_z)_{z\in\Pi_u} \subset T_u(t)\}$ (we use the notation introduced in that proof). Therefore, the distribution of $G_{<u}$ (and hence a_u) conditional on $\theta = \theta_0$, $F_{<v} = F$, $t_v = t$, and $W_{<v}$ is the same no matter what are other agents in F. This observation allows to simplify the condition:

$$\mathbb{P}(a_u \neq \theta \mid \theta = \theta_0,\ t_v = t,\ F_{<v} = F,\ W_{<v}) = \mathbb{P}(a_u \neq \theta \mid \theta = \theta_0,\ t_u < t,\ v \notin F_{<u},\ W_{<v}).$$

By the state-symmetry of equilibrium, the latter probability does not change if we eliminate conditioning to $\theta = \theta_0$. This probability can be bounded as follows

$$\mathbb{P}(a_u \neq \theta \mid t_u < t,\ v \notin F_{<u},\ W_{<v}) \leq \frac{\mathbb{P}(a_u \neq \theta \mid\ t_u < t,\ v \notin F_{<u})}{1-\psi}, \tag{B.6}$$

which follows from the formula of total probability and the lower bound $\mathbb{P}(W_{<v} \mid t_u, t_v) \geq 1 - \psi$.

It remains to estimate the numerator. We use the following observation: for any event A that belongs to the information partition of agent u, the conditional probability of a wrong action $\mathbb{P}(a_u \neq \theta \mid A)$ is at most ε. Otherwise, the agent can profitably deviate from his equilibrium strategy by following his signal whenever A occurs.

The event $A' = \{t_u < t,\ v \notin F_{<u}\}$ is not known to u since u does not observe his arrival time. However, we can approximate the event A' by the event $A = \{\hat{t}_u < t,\ v \notin F_{<u}\}$, where $\hat{t}_u = \frac{|F_{<u}|}{\deg_u - 1}$ is a proxy for u's arrival time (for large-degree agents $\hat{t}_u \approx t_u$ by the law of large numbers). Agent u knows when A occurs and, therefore, $\mathbb{P}(a_u \neq \theta \mid A) \leq \varepsilon$. The conditional probability with respect to A' can be bounded as follows

$$\mathbb{P}(a_u \neq \theta \mid A') \leq \frac{\mathbb{P}(a_u \neq \theta,\ A) + \mathbb{P}(A' \setminus A)}{\mathbb{P}(A')} = \mathbb{P}(a_u \neq \theta \mid A) \cdot \frac{\mathbb{P}(A)}{\mathbb{P}(A')} + \frac{\mathbb{P}(A' \setminus A)}{\mathbb{P}(A')} \leq \varepsilon \cdot \frac{\mathbb{P}(A)}{\mathbb{P}(A')} + \frac{\mathbb{P}(A' \setminus A)}{\mathbb{P}(A')}.$$

Let's estimate all the probabilities in this expression. For $\mathbb{P}(A')$ we get

$$\mathbb{P}(A') = \mathbb{P}(t_u < t,\ t_v > t_u) = \int_0^t dt_u \int_{t_u}^1 dt_v = \int_0^t (1 - t_u)\, dt_u = t - \frac{t^2}{2}.$$

To compute $\mathbb{P}(A)$, we note that conditionally on t_u and $t_v > t_u$, the number of friends observed by u has binomial distribution with parameters $\deg_u - 1$ and t_u (there are $\deg_u - 1$ friends and each of them arrives before time t_u independently with probability t_u). By the Hoeffding inequality, we get

$$\mathbb{P}\left(\frac{|F_{<u}|}{\deg_u - 1} \leq t \mid t_u,\ v \notin F_{<u}\right) \leq \exp\left(-2(t_u - t)^2(\deg_u - 1)\right) \quad \text{for } t \leq t_u$$
$$\mathbb{P}\left(\frac{|F_{<u}|}{\deg_u - 1} \geq t \mid t_u,\ v \notin F_{<u}\right) \leq \exp\left(-2(t_u - t)^2(\deg_u - 1)\right) \quad \text{for } t \geq t_u.$$

Now we are ready to estimate $\mathbb{P}(A)$:

$$\mathbb{P}(A) = \mathbb{P}\left(\frac{|F_{<u}|}{\deg_u - 1} \leq t,\ t_v > t_u\right) = \int_0^1 dt_v \int_0^{t_v} \mathbb{P}\left(\frac{|F_{<u}|}{\deg_u - 1} \leq t \mid t_u,\ t_v > t_u\right) dt_u \leq$$
$$\leq \int_0^t dt_v \int_0^{t_v} 1\, dt_u + \int_t^1 dt_v \int_0^t 1\, dt_u + \int_t^1 dt_v \int_t^{t_v} \exp\left(-2(t_u - t)^2(\deg_u - 1)\right) dt_u \leq$$
$$\leq \frac{t^2}{2} + t(1-t) + \int_{-\infty}^0 \exp\left(-2s^2(\deg_u - 1)\right) ds = t - \frac{t^2}{2} + \sqrt{\frac{\pi}{8(\deg_u - 1)}}.$$

Similarly,

$$\mathbb{P}(A' \setminus A) = \mathbb{P}\left(t_u < t,\ t_v > t_u,\ \frac{|F_{<u}|}{\deg_u - 1} > t\right) = \int_0^t dt_u \left(\int_{t_u}^1 dt_v\right) \cdot \mathbb{P}\left(\frac{|F_{<u}|}{\deg_u - 1} > t \mid t_u,\ v \notin F_{<u}\right) \leq$$
$$\leq \int_{-\infty}^0 \exp\left(-2s^2(\deg_u - 1)\right) ds = \sqrt{\frac{\pi}{8(\deg_u - 1)}}.$$

Putting all the pieces together, we obtain

$$\mathbb{P}(a_u \neq \theta \mid A') \leq \varepsilon\left(1 + \frac{1}{t - \frac{t^2}{2}}\sqrt{\frac{\pi}{8(\deg_u - 1)}}\right) + \frac{1}{t - \frac{t^2}{2}}\sqrt{\frac{\pi}{8(\deg_u - 1)}} \leq \varepsilon + \frac{3}{t \cdot \sqrt{\deg_u - 1}}.$$

In the last inequality we took into account that $t - \frac{t^2}{2} \geq \frac{t}{2}$, $\varepsilon \leq 1$, and $\sqrt{2\pi} \leq 3$.

Substituting this expression in (B.6) leads to the desired bound (B.1). $\square$

C Missed proofs for Section 6

Proof of Proposition 6.1. Denote by $\deg'_v$ the degree of an agent $v \in V'$ in the sub-network $G^{V'}$ of G with $|V'| = \lfloor \alpha \cdot |V| \rfloor$. Fix positive $\gamma \le \alpha$. Our goal is to bound the fraction of agents that have $\deg'_v < \gamma \cdot D = \gamma \cdot \deg_v$. Denote by V^1 the set of all such agents $v \in V'$ and by V^2, the set of eliminated agents $V \setminus V'$. Apply inequality (6.2) (see Mixing Lemma A.1) to these V^1 and V^2. Since $E(V^1, V^2) > (1-\gamma)D \cdot |V^1|$ and $(1-\alpha)|V| - 1 < |V^2| \le (1-\alpha)|V|$, we get

$$(1-\gamma)D \cdot |V^1| - \frac{D}{|V|} \cdot |V^1| \cdot (1-\alpha)|V| \le |\lambda_2| \sqrt{|V^1| \cdot (1-\alpha)|V|}.$$

Dividing both sides by $\sqrt{|V^1|}$ and rearranging terms we get

$$(\alpha-\gamma)D\sqrt{|V_1|} \le |\lambda_2|\sqrt{(1-\alpha)|V|} \Longrightarrow |V^1| \le \frac{(1-\alpha)|\lambda_2|^2}{(\alpha-\gamma)^2 D^2}|V| \le \frac{(1-\alpha)|\lambda_2|^2}{\alpha(\alpha-\gamma)^2 D^2}|V'|.$$

Thus at least $\left(1 - \frac{(1-\alpha)|\lambda_2|^2}{\alpha(\alpha-\gamma)^2 D^2}\right)$ fraction of agents $v \in V'$ has $\deg_v \ge \gamma \cdot D$.

By Lemma C.1 contained below, at least $\left(1 - \frac{2(1-\alpha)|\lambda_2|^2}{\alpha(\alpha-\gamma)^2 D^2}\right)|V'|$ agents has at least $\frac{\gamma}{2} \cdot D$ friends with degree $\frac{\gamma}{2} \cdot D$ or higher.

Applying Proposition 5.1 to this set and estimating the chance of the correct action outside this set by zero, we obtain the following bound on the learning quality for any state-symmetric equilibrium σ' in $G^{V'}$:

$$L_{\sigma'}\left(G^{V'}\right) \ge \left(1 - \frac{2(1-\alpha)|\lambda_2|^2}{\alpha(\alpha-\gamma)^2 D^2}\right)\left(1 - \delta\left(\varepsilon, \frac{\gamma}{2} \cdot D, r, D\right)\right),$$

where $r = \left\lfloor \frac{g}{2} - 1 \right\rfloor$ by Observation 5.1.

Taking into account that $\delta(\varepsilon, D, r, D) > \frac{18}{\sqrt{D}}$, we see that the expression in the first parentheses is greater than $1 - \frac{(1-\alpha)|\lambda_2|^2}{9\alpha(\alpha-\gamma)^2 D^{\frac{3}{2}}} \cdot \delta(\varepsilon, D, r, D)$. It is easy to check that $\delta(\varepsilon, \beta D, r, D) \le \frac{1}{\sqrt{\beta - \frac{1}{D}}}\delta(\varepsilon, D, r, D)$ and therefore the expression in the second parentheses is at least $1 - \frac{1}{\sqrt{\frac{\gamma}{2} - \frac{1}{D}}} \cdot \delta(\varepsilon, D, r, D)$.

Opening the brackets and throwing away positive terms, we get

$$L_{\sigma'}\left(G^{V'}\right) \ge 1 - \left(\frac{(1-\alpha)|\lambda_2|^2}{9\alpha(\alpha-\gamma)^2 D^{\frac{3}{2}}} + \frac{1}{\sqrt{\frac{\gamma}{2} - \frac{1}{D}}}\right) \cdot \delta(\varepsilon, D, r, D).$$

Picking $\gamma = \frac{2\alpha}{3}$ and assuming that $\frac{\alpha}{3} - \frac{1}{D} \ge \frac{\alpha}{4}$ we get the desired bound (6.1). In the complement case of small $\alpha < \frac{12}{D}$ the bound (6.1) since $\frac{2}{\sqrt{\alpha}}\delta(\varepsilon, D, r, D) \ge \frac{\sqrt{D}}{\sqrt{3}} \cdot \frac{18}{\sqrt{D}} > 1$. $\square$

Lemma C.1. *If in a network $G = (V, E)$ at least $(1-\beta)|V|$ agents have degree D or higher, then at least $(1-2\beta)|V|$ agents have at least $\frac{D}{2}$ friends with degree at least $\frac{D}{2}$.*

Proof. Denote by $V_{<D}$ the set of agents with less than D friends and by $V_{<\frac{D}{2}, \ge \frac{D}{2}}$ the set of agents that have less than $\frac{D}{2}$ friends with degree $\frac{D}{2}$ or higher. We know $|V_{<D}| < \beta|V|$ and want to prove that $\left|V_{<\frac{D}{2}, \ge \frac{D}{2}}\right| < 2\beta|V|$. Assume towards a contradiction that $\left|V_{<\frac{D}{2}, \ge \frac{D}{2}}\right| \ge 2\beta|V|$. Therefore $V' =$

$V_{<\frac{D}{2},\geq\frac{D}{2}} \cap (V \setminus V_{<D})$ contains at least $\beta|V|$ agents. Each $v \in V'$ has at least D friends and at least $D - \frac{D}{2} = \frac{D}{2}$ of them have degree less than $\frac{D}{2}$; denote the set of such low-degree friends by $F_{v,<\frac{D}{2}}$. The set $V_{<D}$ contains the union of $F_{v,<\frac{D}{2}}$ over $v \in V'$. We obtain

$$\frac{2}{D}\sum_{v\in V'}|F_{v,<\frac{D}{2}}| \leq |V_{<D}|,$$

where the factor $\frac{2}{D}$ originates because no $u \in F_{v,<D}$ is counted more than $\deg_u < \frac{D}{2}$ times. Using the bounds on cardinalities of all the sets in this expression, we get

$$\frac{2}{D}\cdot\beta|V|\cdot\frac{D}{2} < \beta|V| \Longleftrightarrow 1 < 1.$$

This contradiction completes the proof. $\square$