

On social networks that support learning¹

Itai Arieli, Fedor Sandomirskiy, and Rann Smorodinsky

Abstract

It is well understood that the structure of a social network is critical to whether or not agents can aggregate information correctly. In this paper, we study social networks that support information aggregation when rational agents act sequentially and irrevocably. Whether or not information is aggregated depends, *inter alia*, on the order in which agents decide. Thus, to decouple the order and the topology, our model studies a random arrival order.

Unlike the case of a fixed arrival order, in our model the decision of an agent is unlikely to be affected by those who are far from him in the network. This observation allows us to identify a *local learning requirement*, a natural condition on the agent's neighborhood that guarantees that this agent makes the correct decision (with high probability) no matter how well other agents perform. Roughly speaking, the agent should belong to a multitude of social circles that, if this agent was absent, would be mutually exclusive.

We illustrate the power of the local learning requirement by constructing a family of social networks that guarantee information aggregation despite that no agent is a social hub (in other words, there are no opinion leaders). Although, the common wisdom of the social learning literature suggests that information aggregation is very fragile, another application of the local learning requirement demonstrates the existence of networks where learning prevails even if a substantial fraction of the agents are not involved in the learning process. On a technical level, the networks we construct rely on the theory of expander graphs, *i.e.*, highly connected sparse graphs with a wide range of applications from pure mathematics to error-correcting codes.

1 Introduction

The way ideas and information propagate in society is critical for engineering political campaigns, evaluating new technologies, marketing products, and introducing new social conventions. It is well known that whether or not the information is properly aggregated is highly dependent on the quality of information accessible to the individual agents, on the way information is transmitted from one agent to another, the order and frequency with which agents take actions, and, finally, on the topology of the underlying social network.

In this paper, we study the role that the topology of the social network plays in information aggregation among rational agents. We do so using a variant of the herding model of Banerjee [1992] and Bikhchandani et al. [1992] adapted to a social network setting. Agents

¹We are grateful (in alphabetic order) to Herve Moulin, Alexander Nesterov, Matthew Jackson, Nicolas Vieille, and Omer Tamuz for inspirational discussions and suggestions. We also thank attendees of the game-theory seminar at the Technion and the Conference on Mechanism and Institution Design 2020 for their feedback. We are grateful to Michael Borns for proofreading the paper.

Arieli's research is supported by the Ministry of Science and Technology grant #2028255. Sandomirskiy's research is partially supported by the Lady Davis Foundation, Grant 19-01-00762 of the Russian Foundation for Basic Research, the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (#740435), and the Basic Research Program of the National Research University Higher School of Economics.

Smorodinsky'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.

decide sequentially on a binary action based on a private (bounded) signal and history of actions taken by their predecessors. In the variant we study, the set of predecessors an agent observes is restricted by his set of neighbors in an exogenously given social network.

Since agents act sequentially a network may properly aggregate information for a particular sequence while failing to do so for most other sequences. To decouple the network’s topology from the order in which agents take their actions, we make a distinction between the *a-priori network*, which is exogenously given, and the *realized network*, which is induced by the aforementioned social network and the ordering of agents. Such a distinction between the social network and realized observability structure is not standard. Our goal in this paper is to find a sufficient condition on the a-priori network that guarantees learning for most sequences. We say that such a network *supports learning*. Our first step in the analysis, which is of independent interest, is to focus on a particular agent in the a-priori network and characterize features of the **local** structure of the network that guarantee that this agent will make the correct decision (with high probability). Indeed, we show that an agent that belongs to a variety of social circles that, among themselves, are mutually exclusive and socially distant has a high probability of taking the correct action. We refer to this condition as the *local learning requirement*. Therefore, social networks with the property that each agent satisfies the local learning requirement will aggregate information. We show, with an example, that this condition is not vacuous. Quite surprisingly, we demonstrate the existence of **symmetric** networks with the aforementioned property.

In the sociology literature, and in particular the literature on mass communication, it is often argued that learning is facilitated vis-à-vis a small number of opinion leaders predetermined by their position in the social network [Katz and Lazarsfeld, 1955]. Thus, the existence of symmetric networks that support learning is quite counterintuitive as it implies that learning obtains without opinion leaders (see the discussion in Appendix A).

When studying learning in social networks it may be quite unrealistic to assume that every decision problem pertains to all agents in the society. In fact, in large societies (such as those present in online social networks) it is far more realistic that only a minority of the agents have a stake in an arbitrary decision problem. For example, a dilemma between two candidates for mayorship may interest one subset of the population while a dilemma between two competing “green” technologies (e.g., hybrid engines vs. electric engines) may be relevant to some other small set of agents. Thus, the proper requirement on a social network is that information be properly aggregated when only a fraction of agents participate. We say that a network *supports robust learning* if any subnetwork containing a constant fraction of the agents supports learning.² To obtain robust learning it is sufficient that in any network induced by a fraction of the agents most of those agents satisfy the local learning requirement. We show, with an example, that this condition is not vacuous. Once again, the example we have satisfies symmetry assumptions on the agents.

To construct social networks that support (robust) learning we tap into the theory of expander graphs and in particular we make use of well-known properties of a family of graphs known as *Ramanujan* expanders.

1.1 Related literature

The literature on information aggregation studies many aspects of social learning. In our literature review, we focus on the connection between the topology of the social network and information aggregation. This connection has received attention in the literature on repeated interaction in social networks, where, in contrast to our model, agents can revise their decision as many times as they want. For example, Mossel et al. [2015] study Bayesian

²An alternative, weaker definition could require that most such subnetworks support learning. We discuss this alternative definition in Appendix F.

learning by rational agents and provide a sufficient condition on the network structure such that asymptotically agents select the optimal action with a probability that approaches one as the network grows. Golub and Jackson [2010] analyze conditions under which a “naive” updating process converges to a fully rational limiting belief, which equals the Bayesian posterior conditional on all agents’ information, in a large society. The conditions require no agent to be too influential.

In the context of sequential social learning, where each agent takes an action only once, it was first suggested by Smith [1991] that learning can be obtained by restricting the social connections among agents. In particular, Smith points out that, for a given sequence of agents, if the deciding agents initially do not observe each other then they collectively form a sample of independent observations (sometimes referred to as the “guinea pigs”). Thus, any subsequent agent observing them is likely to make the correct decision. SgROI [2002] discusses how to choose the set of guinea pigs optimally. Unfortunately, these insights couple the network structure with the order in which agents make their decision and is, therefore, moot to the questions we pose.

The interplay between the network structure and sequential learning is studied by Acemoglu et al. [2010] in a model with rational agents and sequential actions. In their model, prior to making his decision, an agent observes a random sample of predecessors, and the paper focuses on sufficient conditions on this sampling distribution that guarantee learning. The network structure turns out to be coupled with the order in which agents make their decisions. Although the sampling model of Acemoglu et al. [2010] induces a random realized network, this network cannot be derived from and so is not associated with an exogenously given a-priori network.³ Thus, our approach complements that of Acemoglu et al. [2010] in the sense that their model better suits settings where the network structure is generated ad hoc whereas our model fits social networks whose a-priori structure is fixed with no connection to the order in which agents decide in one problem or another and so the same network underlies a variety of decision problems, each with its own order. Arieli and Mueller-Frank [2019] also consider random sampling but assume the edges in the network are limited to the m -dimensional lattice.

Bahar et al. [2020] are the first to study social learning over an exogenously given network with a random arrival order of the agents. They demonstrate the existence of a social network where learning is guaranteed. Their network has a particular structure. It is a bipartite graph with a minority of agents on one side (whom they refer to as “celebrities”) and a majority on the other side (“commoners”); for more details see Example 2.2. Although the family of celebrity graphs supports learning, they are sensitive to the participation of agents. In other words, in decision problems that are irrelevant to a minority of the agents, information need not be properly aggregated and social learning may fail; i.e., the celebrity graphs do not support robust learning.

Recently, various papers have demonstrated how fragile social learning can be (see, e.g., Bohren [2016]; Frick et al. [2020]; Mueller-Frank [2018]). In contrast to these findings we construct a network structure that is robust in the sense that learning prevails even if a majority of agents, chosen adversarially, does not participate.

On a more technical level, the networks we construct rely on the theory of expander graphs, i.e., highly connected sparse graphs with a wide range of applications from solving problems in pure mathematics 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.

³Unlike the independence of the random observation sampling assumed by Acemoglu et al. [2010], in our model the random order generates a high correlation between the observation sets of individuals. Indeed, if an individual observes only a small subset of his friends when he makes his decision he may infer that those friends are early arrivals and, therefore, based their own decisions on very limited information. This is not the case in the random sample model.

[2014] and Feldman et al. [2014] in the context of boundedly rational agents who take actions repeatedly.

1.2 Structure of the paper

Section 2 contains the description of the model. In Section 3, we demonstrate that whether or not an agent learns is determined by the local structure of the a-priori network around him and derive the local learning requirement on this neighborhood, which guarantees that the agent learns the state. Sections 4 and Section 5 are devoted to implications of the local learning requirement. Section 4 shows that learning is possible in egalitarian societies: there are symmetric networks that support learning. Section 5 strengthens this result by constructing symmetric networks where learning is robust to adversarial elimination of groups of agents.

Appendix A discusses alternative interpretations and extensions of the basic model. The results of Sections 4 and 5 heavily rely on the insights from the theory of expander graphs, which are explained in Appendix B. Technical proofs for Sections 3 and 5 are in Appendices C, D, and E. Appendix F discusses an alternative, weaker notion of robustness, where groups of agents are eliminated at random.

2 The model

A social network is an undirected graph $G = (V, E)$, where V is the set of agents and an edge vu is contained in E if v and u are “friends.” Friendship is always mutual, as on Facebook: $vu \in E \Rightarrow uv \in E$. We denote by F_v the set of v ’s friends $\{u \in V : vu \in E\}$.

With each agent $v \in V$, we associate his arrival time t_v ; the arrival times $T = (t_v)_{v \in V}$ are independent random variables uniformly distributed on⁴ $[0, 1]$. The collection of arrival times T induces the orientation on edges of $G = (V, E)$ from late arrivals to early ones and converts it into the direct network $G_T = (V, E_T)$ with $E_T = \{vu \in E : t_u < t_v\}$, which we call *the realized network*. We will refer to G as the *a-priori network* in order to distinguish it from G_T . The directed edge vu in the realized network means that agent v gets to observe agent u ; i.e., v observes his friends who arrived earlier. We denote the set of all such friends of v by $F_{v,T} = \{u \in V : vu \in E_T\} \subset F_v$.

Upon arrival, each agent v takes an action $a_v \in \{0, 1\}$, depending on the information available. The agent gets a payoff of one if $a_v = \theta$, where θ is the 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 p and $1 - \theta$ with probability $1 - p$, where $\frac{1}{2} < p \leq 1$. Signals are independent conditional on θ . In addition to his own signal s_v , each agent v observes the set of his friends who arrived earlier, $F_{v,T}$, and their actions, $(a_u)_{u \in F_{v,T}}$. We denote the information set of an agent v by $I_v = (s_v, F_{v,T}, (a_u)_{u \in F_{v,T}})$. Note that an agent knows neither his arrival time nor the set of agents who arrived before him (except for his friends).

All agents are rational and risk-neutral; the description of the model, probability distributions, and the a-priori network G are common knowledge. By contrast, the realized network is not observed by agents.

⁴Random arrival order allows us to disentangle the topology of the a-priori network 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 vs. Android, public kindergartens vs. private ones). We stress that in our model *arrival* refers to the time when the agent makes a decision, and not to the time when he joins the network as in the random sampling model of Acemoglu et al. [2010].

2.1 Equilibria

A mixed strategy 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 each agent is to maximize his expected payoff, which coincides with 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}_σ and \mathbb{E}_σ for the probability and for the expectation with respect to all the randomness in the problem (the state, signals, arrival times, and actions). An equilibrium exists since it is a finite game; however, it may be non-unique. We omit the subscript 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 $(\theta, (s_v)_{v \in V}, (a_v)_{v \in V}) \rightarrow (1 - \theta, (1 - s_v)_{v \in V}, (1 - a_v)_{v \in V})$, i.e., under simultaneous flipping of 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 $\theta = 1$ 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-breaking; however, it is not the whole truth since $\mathbb{P}(\theta = 1 \mid I_v)$ depends on the equilibrium strategies of other agents, which, in turn, are optimal replies to the strategies of others, including v . In particular, an equilibrium lacks the sequential structure since no pair of agents $v, u \in V$ such that $vu \notin E$ know which of them acted first.

2.2 Learning

The following definition quantifies how well an agent and the a-priori network itself aggregate information.

Definition 2.1 (Learning quality). For an a-priori network $G = (V, E)$ and an equilibrium σ , the *learning quality* of an agent $v \in V$ is the probability that this agent takes the correct action:

$$l_\sigma(v) = \mathbb{P}_\sigma(a_v = \theta).$$

The *learning quality* of the network G is the expected fraction of agents taking the correct action:

$$L_\sigma(G) = \frac{1}{|V|} \cdot \mathbb{E}_\sigma |\{v \in V : a_v = \theta\}| = \frac{1}{|V|} \sum_{v \in V} l_\sigma(v).$$

We say that an agent with $l_\sigma(v)$ close to one *learns the state* and the network with $L_\sigma(G)$ close to one *supports learning*. The following example shows that networks may fail to support learning due to the herding phenomenon.

Example 2.1 (Herding spoils learning). Let the a-priori network G be the n -clique K_n , the complete graph on n vertices. In this case, every agent v observes actions of all those who came earlier. We consider an equilibrium σ , where in the case of indifference, each agent follows his signal (one can show that learning quality in other equilibria can only be worse). If the first two agents take the incorrect action $a = 1 - \theta$, 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 a probability of at least $(1 - p)^2$ all the agents take the incorrect action. Thus $L_\sigma(K_n) \leq 1 - (1 - p)^2$; i.e., the learning quality is bounded away from 1 even for large cliques. By symmetry, $l_\sigma(v) = L_\sigma(K_n)$ for all agents v .

Bahar et al. [2020] consider a similar model of learning with random arrivals and ask whether there exist networks that support learning. They provide an affirmative answer to this question by identifying a family of *celebrity graphs*, the only known family of networks that supports learning with random arrivals. We will construct another such family in Section 4.

Example 2.2 (Celebrity graphs support learning). Consider a two-tier society, where a large set of k “commoners” observe a large but smaller set of m “celebrities,” $1 \ll m \ll k$. The corresponding a-priori network is a complete bipartite graph $B_{k,m}$; see Figure 1.

On average, a set of $\frac{k}{m+1} \gg 1$ commoners arrive before the first celebrity. These commoners take their actions in isolation, i.e., without observing anybody else, and so these actions are dictated solely by their signals. Such isolated commoners play the role of “guinea pigs” [SgROI, 2002]. Since $\frac{k}{m+1} \gg 1$, the law of large numbers suggests that the first celebrity aggregates information from these i.i.d. inputs and thereby takes the correct action with high probability. Subsequent commoners observe this celebrity and, hence, are also likely to make the right choice. This correct action propagates and so we conclude that the whole population except for a negligible fraction of $\frac{1}{m+1}$ commoners takes the right action with high probability. This informal reasoning suggests that for any $\delta > 0$ we can find k and m such that the learning quality of each agent and of the network itself 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.

Although the celebrity graphs support learning they are not robust to elimination of small groups of agents. Indeed, the elimination of all the celebrities would render the commoners completely isolated and so learning would not obtain. We also note that the fact that celebrity graphs support learning hinges on the assumption of a random arrival order. Indeed, if all commoners arrive before celebrities, then all of them make a decision in isolation and learning fails.




Figure 1: The celebrity graph. When the first celebrity arrives, he typically observes $\frac{k}{m+1} \gg 1$ 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, he takes a well-informed action, and then this action propagates.

2.3 Notation

Throughout the paper we use the following notation. For a real number x , the floor and the ceiling are denoted by $\lfloor x \rfloor$ and $\lceil x \rceil$, respectively. The former is the biggest integer n such that $n \leq x$ and the latter is the smallest integer $n \geq x$. The base of the natural logarithm is denoted by e .

Consider a network $G = (V, E)$ that is possibly directed. By $\deg_G(v)$ we denote the total degree of a vertex v , i.e., the number of $u \in V$ such that at least one of the edges, uv or vu , is in E . The network is *D-regular* if $\deg_G(v) = D$ for all $v \in V$. A map $f : V \rightarrow V$ is called an *automorphism* of G if f is a bijection and it preserves edges, i.e., $uv \in E \Leftrightarrow f(u)f(v) \in E$.

A network $G' = (V', E')$ is a *subnetwork* of G (denoted by $G' \subset G$) if $V' \subset V$ and $E' \subset E \cap (V' \times V')$. The subnetwork is *induced* by the set of vertices V' if $E' = E \cap (V' \times V')$; such a subnetwork is denoted by $G^{V'}$. We denote by $G \setminus v = G^{V \setminus \{v\}}$ the induced subnetwork obtained by deletion of a given vertex v .

A *path of length k* in G is a sequence of vertices (u_0, u_1, \dots, u_k) such that all the edges $u_i u_{i+1}$, $i = 0, 1, \dots, k-1$ belong to E . The distance $d_G(v, v')$ between two vertices $v, v' \in V$ is the minimal k such that there is a path of length k with $u_0 = v$ and $u_k = v'$; if there is no such k , the distance is infinite, $d_G(v, v') = +\infty$. The *r-neighborhood* $B_{r,G}(v)$ of a vertex v is the set of all vertices v' such that $d_G(v, v') \leq r$; the number r is the *radius* of the neighborhood. Abusing the notation, we will not distinguish between the *r-neighborhood* $B_{r,G}(v)$ (the set of vertices) and the subnetwork of G induced by this set of vertices. A *cycle of length k* is a path with $u_0 = u_k$ such that all vertices u_0, u_1, \dots, u_{k-1} are distinct (sometimes cycles with no repetitions are called simple). The *girth* g_G is the length of the shortest cycle.

When referring to the a-priori network G , we will omit the dependence of all the objects on G and write simply $\deg(v)$, $d(v, v')$, $B_r(v)$, and g .

3 Learning and the local network structure

In this section we study the connection between the local structure of the a-priori network in the neighborhood of an agent and his learning quality.

In the case of a fixed arrival order, a major impediment for learning is that an agent can affect the decision of those who are far from him in the social network, which results in the possibility of large information cascades involving most of the network. Surprisingly, in our model with random arrival order, the decisions have a *local nature*: in Section 3.1, we show that if a pair of agents are far from each other in the a-priori network G , with high probability the action of one cannot affect the other. This observation suggests that the learning quality of an agent must be determined by the local structure of his neighborhood in the a-priori network G , not by the global topology of G .

The local nature of decisions motivates one to look for a condition on the topology of an agent's neighborhood that ensures that this particular agent learns the state with high probability no matter what the global topology of the network is and the learning qualities of other agents are. We identify such a *local learning requirement* in Section 3.2.

3.1 Local nature of decisions

For a pair of agents v and u , the action of v may be affected by the choice made by u if v observes u , or if v observes somebody who observes u , or if v observes somebody who observes somebody who observes u , and so on.

Given the collection of arrival times T , the agent v can be influenced by u only if there is a path $(v = u_0, u_1, u_2, \dots, u_k = u)$ in the a-priori network G such that $t_{u_i} > t_{u_{i+1}}$ for all i (i.e., if u is reachable from v by a path in the realized network G_T). We define the *realized subnetwork* $G_{v,T}$ of an agent v to be the induced subnetwork of the realized network G_T composed of v and all such u reachable from v in G_T . Whether v learns the state is determined solely by his realized subnetwork: conditional on $G_{v,T}$, the action of v is independent of the signals and arrival times of all the agents outside $G_{v,T}$.

Recall that $B_r(v)$ denotes the r -neighborhood of v in the a-priori network G and $\deg(v)$ denotes the degree of v in G . The following lemma shows that the realized subnetwork of v is contained in the neighborhood of v with high probability provided that the radius of the neighborhood is big enough compared to the maximal degree.

Lemma 3.1 (Local nature of decisions). *The probability that the realized subnetwork of v is contained in the $(r - 1)$ -neighborhood of v enjoys the following lower bound:*

$$\mathbb{P}\left(G_{v,T} \subset B_{r-1}(v)\right) \geq 1 - 2 \cdot \left(\frac{e \cdot \max_{u \in B_r(v)} \deg(u)}{r}\right)^r. \quad (3.1)$$

An extended version of the lemma is proved in Appendix C. Here we present a sketch.

Proof sketch for Lemma 3.1. The realized subnetwork of v belongs to the $(r - 1)$ -neighborhood if and only if no path $(v = u_0, u_1, u_2, \dots, u_k = u)$ connecting v and the sphere $B_r(v) \setminus B_{r-1}(v)$ in the a-priori network G is presented in the realized network G_T . Note that a particular path of this form exists in G_T if and only if u_k arrives first, then u_{k-1} , then u_{k-2} , and so on. Hence, each path of length k in G remains in the realized network with probability $\frac{1}{(k+1)!}$. Each path has length $k \geq r$ and the total number of paths of length k connecting v and the sphere in G is bounded by $(\max_{u \in B_r(v)} \deg(v))^k$. The union bound accompanied by manipulations with factorials similar to the Stirling formula lead to the desired inequality (3.1). \square

3.2 Local learning requirement

The previous subsection suggests that whether or not an agent v learns the state must be determined by the local structure of the a-priori network G around him. Here we formulate a *local learning requirement* on the topology of v 's neighborhood that ensures high learning quality for v no matter how well other agents perform and what the global structure of the network is. In subsequent sections we demonstrate the usefulness of this condition by constructing networks with exceptional learning and robustness properties.

The essence of the requirement is that v must bridge many social circles. To define the requirement formally, we recall the following notation: $G \setminus v$ is the network obtained from the a-priori network G by elimination of an agent v together with all adjacent edges and $B_{r,G \setminus v}(u)$ is the r -neighborhood of an agent u in $G \setminus v$.

Definition 3.1 (Local learning requirement). An agent v satisfies the local learning requirement with parameters (d, r, D) if among his friends we can find d such that each of them has degree at least d , their r -neighborhoods in $G \setminus v$ are disjoint, and the degrees in all these neighborhoods are upper-bounded by D .

Let u_1, \dots, u_d be the friends from the definition. Then their neighborhoods $B_{r,G \setminus v}(u_i)$ can be interpreted as disjoint social circles bridged by v . The intuition why the definition refers to the network $G \setminus v$ relies on the fact that the decision of v depends only on those agents that arrive before him and, hence, do not observe v as if he was absent from the network.

Figure 2 illustrates the definition. The following theorem is our main technical result.

Theorem 3.1. *If an agent v satisfies the local learning requirement with parameters (d, r, D) , then the learning quality of v enjoys the following lower bound:*

$$l_\sigma(v) \geq 1 - \delta(p, d, r, D), \quad (3.2)$$




Figure 2: Agent v satisfies the local learning requirement with parameters $d = 3$, $r = 2$, and $D = 7$. He has 3 friends with degree at least 3 (dark gray nodes), their 2-neighborhoods in the graph $G \setminus v$ are disjoint (shaded areas), and the maximal degree in these neighborhoods is 7 (the agent at the bottom has 5 friends within the neighborhood and 2 outside).

where

$$\delta(p, d, r, D) = \psi + \frac{18}{\sqrt{d-1}(2p-1-\psi)}, \quad \psi = 2d \cdot \left(\frac{e \cdot D}{r}\right)^r. \quad (3.3)$$

The bound holds for any state-symmetric equilibrium σ provided that the probability of the correct signal satisfies $p \geq \frac{1+\psi}{2}$; i.e., the denominator in (3.3) is positive.⁵

Corollary 3.1. The theorem shows that the agent learns the state whenever $\delta(p, d, r, D)$ is close to zero. We note that $\delta(p, d, r, D)$ goes to zero for fixed $p > \frac{1}{2}$ when all the elements of the triplet (d, r, D) go to infinity such that $\liminf \frac{r}{D} > e$.

The intuition behind Theorem 3.1 is simple. Let u_1, \dots, u_d be the friends of v from Definition 3.1. By the local nature of decisions (Lemma 3.1), the realized subnetwork of each u_i , who arrives before v , is likely to be contained within his social circle $B_{r, G \setminus v}(u_i)$ provided that r is large enough compared to $e \cdot D$. Hence, the realized subnetworks are disjoint, which ensures no herding among $(u_i)_{i=1, \dots, d}$. As a result, v learns the state by observing a large sample of d independent sources of information. We see that there is a delicate trade-off: the degrees have to be big enough to guarantee a big sample but not too big in order to avoid a confounding that may be caused by intersection of the realized subnetworks.

The formalization of the above intuition faces some obstacles: the independence is only conditional and the conditioning is on a family of events that do not belong to the information partition⁶ of any of the agents; hence, checking that the independent sources are informative 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 (hence, the condition $\deg(u_i) \geq d$).

Here we present a sketch of the proof; all the details can be found in Appendix D.

⁵The condition of state-symmetry can be dropped at the cost of getting $(4p-3-\psi)$ in the denominator of (3.3) instead of $(2p-1-\psi)$ and of imposing the stricter condition on p to ensure the positivity of the new denominator (see the discussion on the technical assumptions in Appendix A).

⁶The information partition of an agent v is the collection of all events such that the agent knows if these events have occurred. Each such event is given by a certain list of realizations of agent's signal s_v , the set of friends $F_{v, T}$ arriving earlier, and their actions $(a_u)_{u \in F_{v, T}}$.

Sketch of the proof of Theorem 3.1. Denote by $F_{v,T}^d$ the subset of those friends u_1, \dots, u_d from Definition 3.1 who arrive earlier than v . Consider the following deviation a'_v from an equilibrium strategy a_v of v . He repeats the action played by the majority of $F_{v,T}^d$. In equilibrium, the agent cannot benefit from this deviation and, hence, $l_\sigma(v) \geq \mathbb{P}(a'_v = \theta)$.

The probability of a mistake for a'_v can be bounded using the Hoeffding inequality.⁷ This requires two ingredients: independence of the actions a_u for $u \in F_{v,T}^d$ and a lower bound on the probability of a mistake for a_u .

Unfortunately, the requirement of independent actions is not satisfied. However, the actions can be made independent by conditioning on a certain collection of events. This should still be enough to drive the result provided that the probability of this collection is close to one. Let W_v be the event that the realized subnetwork of each u_i who arrived before v is contained in his social circle $B_{r,G \setminus v}(u_i)$. Since these social circles are disjoint, the realized subnetworks do not intersect and thus each u_i aggregates the information from disjoint families of sources conditional on W_v . The local nature of decisions (Lemma 3.1) implies that the probability of W_v is close to one.

Naively, one could argue that conditional on W_v the actions taken by the friends of v are independent. However, this argument is wrong. There are other sources of dependence. For example, the fact that we are interested in agents who arrive *before* v creates dependence between their actions even if their realized subnetworks are disjoint. To see this consider the case where v arrives early. This implies that all the friends he observes are also early arrivals and, hence, are likely to have smaller realized subnetworks compared to the case where v arrives late; in particular, the earlier v arrives, the more mistakes the observed friends make, thus creating dependence between their decisions.

To eliminate all the sources of dependence, we end up conditioning on the arrival time of v , the set $F_{v,T}^d$, the realized state, and the event W_v .

In order to apply the Hoeffding inequality, it remains to show that the conditional probability of a mistake for a_u is bounded away from $1/2$ for $u \in F_{v,T}^d$, i.e., that 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 $\mathbb{P}(a_u \neq \theta \mid A) \leq 1 - p$ 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 on which we condition does not belong to u 's information partition. We overcome this difficulty by approximating the condition to elements of the information partition and showing that the conditional probability of a mistake is at most $\frac{1}{1-\psi} \left(1 - p + \frac{3}{t_v \cdot \sqrt{\deg(u)-1}} \right)$. The bound gets worse for low-degree agents since their information partition is not fine enough for a good approximation. This is why the deviation a'_v of agent v takes into account high-degree friends only.

After these preparations, Theorem 3.1 becomes a corollary of the Hoeffding inequality. \square

4 Implications: Symmetric networks that support learning

Here we demonstrate that there are a-priori networks where *each agent* satisfies the local learning requirement from the previous section and, hence, all agents achieve high learning quality as does the network itself. Moreover, there are such networks with the additional

⁷The Hoeffding inequality [Hoeffding, 1994] states that $\mathbb{P} \left(\frac{1}{N} \sum_{n=1}^N \xi_i \leq \mathbb{E} \left(\frac{1}{N} \sum_{n=1}^N \xi_i \right) - x \right) \leq \exp(-2x^2 N)$ for independent random variables $0 \leq \xi_i \leq 1$ and any $x \geq 0$.

property that any two agents play the same role. Recall that an automorphism of a network $G = (V, E)$ is a bijection $V \rightarrow V$ that preserve edges, i.e., $uv \in E \Leftrightarrow f(u)f(v) \in E$.

Definition 4.1 (Symmetric networks). A network $G = (V, E)$ is symmetric if for any pair $v, v' \in V$ there exists an automorphism f of G such that⁸ $f(v) = v'$.

Symmetric networks represent totally egalitarian societies. The celebrity graphs from Example 2.2 may lead to a conjecture that egalitarian societies cannot aggregate information since one needs a designated minority of agents (like celebrities) in the a-priori network for learning to propagate. The main result of this section states that this intuition is false and symmetric networks can support learning. In Section 5, we will strengthen this result by showing robustness of learning.

Theorem 4.1. *For any $p_0 > \frac{1}{2}$ and any $\delta > 0$ there exists a symmetric network $G = (V, E)$ such that the learning quality $l_\sigma(v) \geq 1 - \delta$ for each agent $v \in V$, any probability of the correct signal $p > p_0$, and any state-symmetric equilibrium⁹ σ .*

In order to prove Theorem 4.1, we use classic results from the theory of expanders to construct networks, where the local learning requirement is satisfied for each agent.

We will need the following lemma, which simplifies checking the local learning requirement. Recall that the *girth* g of a network G is the length of the shortest cycle; the girth of a tree is infinite, $g = +\infty$.

Lemma 4.1. *Consider a network G of girth g and the maximal degree D . In such a network, any agent v satisfies the local learning requirement with parameters $(d, \lfloor \frac{g-3}{2} \rfloor, D)$, where d is a number such that v has at least d friends of degree d or more.*

Proof. Let u_1, u_2, \dots, u_d be v 's friends with degrees $\deg(u_i) \geq d$. Recall that $B_{r, G \setminus v}(u)$ denotes the r -neighborhood of u in the network, where the agent v was eliminated. We aim to select r such that $B_{r, G \setminus v}(u_i)$ and $B_{r, G \setminus v}(u_j)$ are disjoint for $i \neq j$. If $B_{r, G \setminus v}(u_i)$ and $B_{r, G \setminus v}(u_j)$ intersect, this creates a cycle of length $2r + 2$ (or shorter): start from v , then go to u_i , then to an agent in the intersection by the shortest path, then to u_j again by the shortest path, and back to v . Hence, the intersection is possible only if $2r + 2 \geq g$. Thus choosing r such that $2r + 2 < g$, we ensure that the r -neighborhoods are disjoint; $r = \lfloor \frac{g-3}{2} \rfloor$ is the maximal such r . We conclude that v satisfies the local learning requirement with parameters $(d, \lfloor \frac{g-3}{2} \rfloor, D)$. \square

Recall that a network is called *D -regular* if all vertices have degree D . Thanks to Lemma 4.1 and Theorem 3.1, proving Theorem 4.1 reduces to the question of the existence of symmetric D -regular networks with arbitrary high degree and girth. The existence of such networks is demonstrated in the theory of expanders; see Appendix B. The following is the corollary of a theorem by Lubotzky et al. [1988] (Theorem B.1 in Appendix B) that describes a family of so-called Ramanujan expanders.

Corollary 4.1 (of Theorem B.1 by Lubotzky et al. [1988]). There is a sequence $D_k \rightarrow \infty$ such that for any g_0 and k , there exists a symmetric D_k -regular network $G = (V, E)$ such that the girth $g \geq g_0$ and $|\lambda_2| \leq 2\sqrt{D_k - 1}$, where λ_2 is the second-largest eigenvalue of the adjacency matrix.¹⁰

⁸In mathematical literature, such graphs are usually called *transitive* since the group of automorphisms acts transitively on them; i.e., for any pair of vertices there is an automorphism mapping one to the other.

⁹The theorem extends to non-state-symmetric equilibria at the cost of assuming that signals are informative enough, namely, $p_0 > \frac{3}{4}$. This and other extensions are discussed in Section A.2.

¹⁰We do not use the bound on λ_2 in Section 4; in particular, any D -regular network with large D and girth would suffice for the proof of Theorem 4.1. The bound on $|\lambda_2|$ will be critical for the discussion on robust learning in Section 5.

After all these preliminaries, proving Theorem 4.1 becomes easy.

Proof of Theorem 4.1. For a D -regular network of girth g , Theorem 3.1 combined with Lemma 4.1 implies a lower bound on the learning quality $l_\sigma(v) \geq 1 - \delta \left(p, D, \left\lfloor \frac{g-3}{2} \right\rfloor, D \right)$ for any agent v and any state-symmetric equilibrium σ .

Corollary 4.1 allows us to pick a symmetric D -regular network G with arbitrary high degree D and arbitrary high girth/degree ratio $\frac{g}{D}$, while Corollary 3.1 implies that $\delta \left(p, D, \left\lfloor \frac{g-3}{2} \right\rfloor, D \right)$ tends to zero if both the girth g and the degree D tend to infinity and the girth goes to infinity faster: $\frac{g}{D} \rightarrow \infty$. Hence, for any given δ and p_0 , we can choose G to be a symmetric D -regular network with D and g such that $\delta \left(p_0, D, \left\lfloor \frac{g-3}{2} \right\rfloor, D \right) \leq \delta$. Since, δ is decreasing in p , we obtain $\delta \left(p, D, \left\lfloor \frac{g-3}{2} \right\rfloor, D \right) \leq \delta \left(p_0, D, \left\lfloor \frac{g-3}{2} \right\rfloor, D \right) \leq \delta$ for any $p \geq p_0$. Thus the network G satisfies the statement of Theorem 4.1. \square

5 Implications: Robust learning

Here we discuss the learning quality of a network when some of the agents are uninterested or unavailable and, hence, do not participate in the learning process. The celebrity graphs (Example 2.2) demonstrate that high learning quality can be very fragile: if celebrities (a negligible minority of the population) do not show up, this leaves the network totally disconnected, and information aggregation breaks down. We show that there are networks that *support robust learning*, i.e., are free of such a flaw.

In Section 4, we saw that symmetric sparse networks that have high degrees but not short cycles support learning. In such networks, all agents play the same role, which makes it natural to expect that no *small group* of agents is critical in these networks; in particular, adversarial elimination of a small group cannot dramatically spoil the learning outcome. Here we obtain the surprising, much stronger result demonstrating that there are symmetric networks where learning is robust to the adversarial elimination of any group, even a large one.

Theorem 5.1. *For any $p_0 > \frac{1}{2}$ and any $\delta > 0$ there exists a symmetric network $G = (V, E)$ such that for any $\alpha \in (0, 1]$ and any subset $V' \subset V$ with $\lceil \alpha \cdot |V| \rceil$ agents, the learning quality in the induced subnetwork $G^{V'}$ is at least $1 - \frac{\delta}{\alpha^3}$ for any state-symmetric equilibrium and any probability of the wrong signal¹¹ $p > p_0$.*

Robustness of learning turns out to be related to spectral properties of the network. This is captured by Lemma 5.1 below, and Theorem 5.1 easily follows from this lemma combined with known results on expander graphs. Consider a D -regular network and denote by $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_{|V|}|$ the eigenvalues of its adjacency matrix ordered by their absolute values. Expanders are networks with small $|\lambda_2|$ relative to $|\lambda_1|$; see Appendix B.

The next lemma bounds by how much the learning quality can decrease if, instead of the original D -regular network, we consider its arbitrary subnetwork with $\lceil \alpha \cdot |V| \rceil$ agents.

Lemma 5.1. *Let $G = (V, E)$ be a D -regular network of girth g with the second-largest eigenvalue λ_2 . Then for any $\alpha \in (0, 1]$ and any subset $V' \subset V$ of size $|V'| = \lceil \alpha \cdot |V| \rceil$, the learning quality in the induced subnetwork $G^{V'}$ satisfies*

$$L_\sigma \left(G^{V'} \right) \geq 1 - \left(\frac{2}{\sqrt{\alpha}} + \frac{(1-\alpha)|\lambda_2|^2}{\alpha^3 \cdot D^{\frac{3}{2}}} \right) \cdot \delta \left(p, D, \left\lfloor \frac{g-3}{2} \right\rfloor, D \right) \quad (5.1)$$

¹¹Similarly to Sections 3 and 4, the result extends to non-state-symmetric equilibria under the requirement $p_0 > \frac{3}{4}$.

for any state-symmetric equilibrium σ' . Here $\delta(p, d, r, D)$ is given by formula (3.3).

Lemma 5.1 is proved in Appendix E.

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 R. K. 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. *The Review of Economic Studies*, 86(3):913–940, 2019.
- 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. *Quarterly Journal of Economics*, 107(3):797–817, 1992.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom and cultural change as information cascade. *Journal of Political Economy*, 100:992–1026, 1992.
- Aislinn Bohren. Informational herding with model misspecification. *Journal of Economic Theory*, 163:222–247, 2016.
- Xavier Dahan. Regular graphs of large girth and arbitrary degree. *Combinatorica*, 34(4):407–426, 2014.
- Michal Feldman, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. Reaching consensus via non-Bayesian asynchronous learning in social networks. In *Leibniz International Proceedings in Informatics*. Klaus Jansen, José Rolim, Nikhil Devanur, and Cristopher Moore (eds.), pp. 192–208. Dagstuhl Publishers, 2014.
- Mira Frick, Ryota Iijima, and Yuhta Ishii. Misinterpreting others and the fragility of social learning. *Econometrica* (forthcoming), 2020.
- 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*. Free Press, 1955.
- 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, Allan Sly, and Omer Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015.
- Manuel Mueller-Frank. Manipulating opinions in social networks. *Available at SSRN 3080219*, 2018.
- 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.

A Discussion

In Section A.1, we offer two alternative ways to perceive the results on learning in symmetric networks (Section 4) and those on robust learning (Section 5). First, we present the network-designer perspective and then explain the connection to the classic sociological theory of the two-step information flow. Section A.2 discusses our technical assumptions and how to relax them.

A.1 Alternative interpretations of results

Network-designer perspective Consider a designer of a social network who wants to ensure social learning. Sgroi [2002] constructs a network that supports learning for a particular arrival order, and Bahar et al. [2020] propose the celebrity graphs as networks that support learning and are robust to the arrival order of agents. However, the celebrity graphs only support 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 are consequently inactive, then learning may fail. By contrast, our expander-based networks from Theorem 5.1 give a social structure that supports learning and is robust both to the arrival order and to the actual subset of active agents.

In practice, social networks are not designed from scratch and the social connections can be considered as given. However, social networks on the Internet are usually combined with a recommendation system that decides which news to show to a particular user. In so doing, it can intervene in the structure of social connections by concealing some friends’ news and possibly showing some news of non-friends. Our local learning requirement (Theorem 3.1) combined with Lemma 4.1 suggests a possible recipe for improving overall learning quality: eliminating short cycles while keeping high the number of information sources an agent is subject to. This proposal, however, requires additional empirical evaluation.

The paradigm of the two-step information flow The classic sociological theory of the “two-step information flow” by Katz and Lazarsfeld [1955] conveys the idea that learning is always facilitated by a small group of opinion leaders, i.e., influential agents predetermined by the network structure. The original version of the theory, 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. In other words, people’s actions, whether in the form of adopting a new product or voting for a certain political candidate, were not so much a result of what they heard from the TV networks but rather what they heard from influential agents who, for their part, had consumed such content directly.

Let us call a group of agents *influential* if the network supports learning whenever this group is present and fails to aggregate the information when it is not. An example of an influential minority is the group of celebrities in the celebrity graphs of Example 2.2.

The theory of the two-step information flow suggests that networks supporting learning contain an influential minority, which plays the role of intermediary in the information spread over the network. The results of Sections 4 and 5 challenge this thesis (in our stylized theoretical model): high learning quality can be achieved in a-priori symmetric networks and, hence, no group of agents is predetermined by the network structure; moreover, there are symmetric networks where no minority (or even majority) of agents is influential.

A.2 Getting rid of technical assumptions

Some of our modeling assumptions were made to simplify the exposition and can be easily relaxed.

General distributions of arrival times We assumed that agents’ arrival times 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 $[0, 1 - p] \cup [p, 1]$ (the probability of this set of signals will enter into the bound on the learning quality). By “symmetry” we mean that the distribution posteriors $\mathbb{P}(\theta = 1 \mid s)$ is symmetric around $\frac{1}{2}$. Note that signals of unbounded precision are also allowed.

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

Non-symmetric equilibria The symmetry assumption on equilibria and signals can be relaxed. However, for asymmetric equilibria, Theorems 3.1, 4.1, and 5.1 require the probability p of the correct signal to be above $\frac{3}{4}$ (instead of $\frac{1}{2}$). The reason for this is inequality (D.1), where for asymmetric equilibria we get $2(1 - p)$ in the parentheses.

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

More informed agents The results of Sections 3, 4, and 5 hold if in addition to observing the actions of his friends, each agent v gets some information about his realized subnetwork $G_{v,T}$, e.g., a possibly noisy signal about the set of agents, their arrival times, actions, and even their private signals.

Less informed agents All our results can be easily adapted to the case, where observation of friends' actions is noisy. Namely, each agent v , instead of observing the action a_u of his friend $u \in F_{v,T}$, observes either a_u with probability $1 - \varepsilon$ or the flipped action $1 - a_u$ with probability ε . One can consider two variants of this model: the action of u is flipped for all his neighbors at the same time but independently across $u \in V$, or the action is flipped independently for each pair (v, u) .

Both variants require the same straightforward modifications in Sections 3, 4, and 5. They originate from a minor adjustment in the proof of Theorem 3.1: when applying the Hoeffding inequality in formula (D.3), instead of the upper bound on $\mathbb{P}(a_u \neq \theta)$, we will need an upper bound on the probability that the action *observed by v* does not match the state. The latter probability is equal to $(1 - \varepsilon)\mathbb{P}(a_u \neq \theta) + \varepsilon(1 - \mathbb{P}(a_u \neq \theta))$ and does not exceed $(1 - \varepsilon)\mathbb{P}(a_u \neq \theta) + \varepsilon$, where $\mathbb{P}(a_u \neq \theta)$ can be bounded by Lemma D.3 as before.

B 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 $|\lambda_1| \geq |\lambda_2| \geq |\lambda_3| \dots \geq |\lambda_{|V|}|$. The eigenvalue $\lambda_1 = D$ corresponds to the eigenvector representing the uniform distribution over vertices: the smaller the second eigenvalue $|\lambda_2|$ is, 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 B.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 the random walk and, hence, slow down the expansion of the random walk over the graph.

The proof of Theorem B.1 is quite technical and relies on group theory. Lubotzky et al. [1988] and their 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.¹² 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].¹³

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

¹²Given a group \mathcal{G} with a group operation \circ and a subset $S \subset \mathcal{G}$, its Cayley graph is defined in the following way: the set of vertices V coincides with \mathcal{G} and $vu \in E$ if $u = v \circ s$ for some $s \in S$.

¹³Indeed, for any $x \in \mathcal{G}$ the map $f(v) : v \rightarrow x \circ v$ is an automorphism of the Cayley graph. For any pair of vertices $v, v' \in V$, we achieve $f(v) = v'$ by choosing $x = v' \circ v^{-1}$.

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 endpoint in V^1 and another in V^2 . In the Erdos–Renyi model the expected number of such edges is equal to $p \cdot |V^1| |V^2|$. The following result, known as the mixing lemma, shows that $|E(V^1, V^2)|$ for an expander is close to this number.

Lemma B.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:*

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

C Local nature of decisions

Here we prove Lemma 3.1, which ensures that the realized subnetwork $G_{v,T}$ of v is likely to be contained in the neighborhood of v whenever the radius of this neighborhood is large enough compared to the maximal degree. In fact, we prove a slightly more general statement with conditioning on v 's arrival time; this is a technical nuance that we will need in the proof of Theorem 3.1.

Lemma C.1. *The probability that the realized subnetwork of v is contained in the $(r-1)$ -neighborhood of v , conditional on v 's arrival time t_v , enjoys the following lower bound:*

$$\mathbb{P}\left(G_{v,T} \subset B_{r-1}(v) \mid t_v\right) \geq 1 - 2 \cdot \left(\frac{e \cdot D}{r}\right)^r, \quad (\text{C.1})$$

where $D = \max_{u \in B_r(v)} \deg(u)$.

Proof of Lemma C.1. Without loss of generality we can assume that $r > e \cdot D$ since otherwise the bound (C.1) trivializes.

The realized subnetwork of v is contained in the $(r-1)$ -neighborhood if and only if no path connecting v and the sphere $S_r(v) = B_r(v) \setminus B_{r-1}(v)$ in the a-priori network G exists in the realized network G_T . Consider an agent $u \in S_r(v)$ and let $(v = u_0, u_1, u_2, \dots, u_k = u)$ be a path connecting v and u in G . The chance that this path is presented in the realized network G_T is bounded by the probability that u_k, u_{k-1}, \dots, u_1 arrive exactly in this order (we excluded the agent $u_0 = v$ since we condition on his arrival time); this probability is equal to $\frac{1}{k!}$. Each path between v and the sphere $S_r(v)$ has length $k \geq r$. The total number of different paths of length k starting from v is at most D^k since there are at most D options for each of k steps.

Therefore, by the union bound, the chance that there is a path between v and $S_r(v)$ in the realized network G_T is at most

$$\sum_{k=r}^{\infty} \frac{D^k}{k!}.$$

Since $\frac{D^k}{k!} \leq \frac{D^r}{r!} \cdot \left(\frac{D}{r}\right)^{k-r}$, we can bound the sum by the geometric progression

$$\sum_{k=r}^{\infty} \frac{D^k}{k!} \leq \frac{D^r}{r!} \sum_{l=0}^{\infty} \left(\frac{D}{r}\right)^l = \frac{D^r}{r!} \cdot \frac{1}{1 - \frac{D}{r}} \leq \frac{1}{1 - \frac{1}{e}} \cdot \frac{D^r}{r!} \leq 2 \cdot \frac{D^r}{r!},$$

where the last two inequalities follow from the assumption $r \geq e \cdot D$ and the fact that $\frac{1}{1 - \frac{1}{e}} \leq 2$. Using the inequality $n! \geq \left(\frac{n}{e}\right)^n$, we can get rid of the factorial in the denominator:

$$2 \cdot \frac{(D)^r}{r!} \leq 2 \cdot \left(\frac{e \cdot D}{r}\right)^r.$$

Hence, with probability at least $1 - 2 \cdot \left(\frac{e \cdot D}{r}\right)^r$ there are no paths between v and $S_r(v)$ in G_T or, equivalently, $G_{v,T} \subset B_{r-1}(v)$. \square

D Proof of Theorem 3.1

Denote by $F_v^d = (u_i)_{i=1,\dots,d}$ the collection of v 's friends from the local learning requirement and by $F_{v,T}^d = \{u \in F_v^d : t_u \leq t_v\}$ those of them who arrive earlier than v .

Consider the following deviation a'_v of v from his equilibrium action a_v . The agent decides to rely on his friends $F_{v,T}^d$: if the majority of them played action $a \in \{0, 1\}$, the agent v repeats this action; in the 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 a'_v being a mistake can be estimated using the Hoeffding inequality (see Footnote 7).

To apply the Hoeffding inequality, we need independence of actions $(a_u)_{u \in F_{v,T}^d}$ 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.

Recall that $B_{r,G \setminus v}(u)$ denotes the r -neighborhood of u in the network obtained from G by eliminating v . Denote by W_v the event that the realized subnetwork of each $u \in F_{v,T}^d$ is contained in $B_{r-1,G \setminus v}(u)$, i.e., $W_v = \bigcap_{u \in F_{v,T}^d} \{G_{u,T} \subset B_{r-1,G \setminus v}(u)\}$. Note that, given W_v , the realized subnetworks of $u \in F_{v,T}^d$ are disjoint.

Lemma D.1. *Fix $\theta_0 \in \{0, 1\}$, $t \in [0, 1]$, and a subset of friends $F \subset F_v^d$ 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 friends who arrived earlier $F_{v,T}^d = F$, and W_v .*

The second lemma ensures that W_v is a high-probability event. It is a corollary of the local nature of decisions property (Lemma C.1).

Lemma D.2. *The probability of W_v conditional on arrival times of v and all his friends from F_v^d enjoys the following lower bound:*

$$\mathbb{P}(W_v \mid t_v, (t_z)_{z \in F_v^d}) \geq 1 - \psi, \quad \text{where } \psi = 2d \cdot \left(\frac{e \cdot D}{r}\right)^r.$$

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

Lemma D.3. *For any subset F of F_v^d , an agent $u \in F$, and any state-symmetric equilibrium, the conditional probability of the wrong action satisfies the following inequality:*

$$\mathbb{P}(a_u \neq \theta \mid \theta = \theta_0, t_v = t, F_{v,T}^d = F, W_v) \leq \frac{1}{1 - \psi} \left(1 - p + \frac{3}{t \cdot \sqrt{\deg(u) - 1}}\right), \quad (\text{D.1})$$

where ψ is from Lemma D.2.

The lemmas are proved below. With their help, we complete the proof of Theorem 3.1. By the Hoeffding inequality and Lemmas D.1 and D.3, the probability of the wrong action

a'_v for the high state ($\theta = 1$) is bounded as follows:

$$\mathbb{P}(a'_v \neq \theta \mid \theta = 1, t_v, F_{v,T}, W_v) \leq \mathbb{P}\left(\sum_{u \in F_{v,T}^d} a_u \leq \frac{|F_{v,T}^d|}{2} \mid \theta = 1, t_v, F_{v,T}, W_v\right) \leq \quad (\text{D.2})$$

$$\leq \exp\left(-2\left(\frac{1}{2} - \frac{1}{1-\psi}\left(1-p + \frac{3}{t_v \cdot \sqrt{d-1}}\right)\right)^2 \cdot |F_{v,T}^d|\right) = \quad (\text{D.3})$$

$$= \exp\left(-\frac{1}{2}\frac{\left(2p-1-\psi - \frac{6}{t_v \cdot \sqrt{d-1}}\right)^2}{(1-\psi)^2} \cdot |F_{v,T}^d|\right), \quad (\text{D.4})$$

where we used that $\deg(u) \geq d$ for $u \in F_v^d$. The bound (D.4) holds only if t_v is not too small: to apply the Hoeffding inequality, the expression in the parentheses, $\left(2p-1-\psi - \frac{6}{t_v \cdot \sqrt{d-1}}\right)$, must be non-negative (see the requirement $x \geq 0$ in Footnote 7). We will use the bound (D.4) only for t_v such that this expression is at least $\frac{2p-1-\psi}{2}$ or equivalently $t_v \geq \frac{12}{\sqrt{d-1}(2p-1-\psi)}$. For small t_v , we roughly bound the probability by 1 and get the following inequality valid for all t_v :

$$\mathbb{P}(a'_v \neq \theta \mid \theta = 1, t_v, F_{v,T}, W_v) \leq \mathbb{1}_{\left\{t_v < \frac{12}{\sqrt{d-1}(2p-1-\psi)}\right\}} + \exp\left(-\frac{1}{8}\frac{(2p-1-\psi)^2}{(1-\psi)^2} \cdot |F_{v,T}^d|\right), \quad (\text{D.5})$$

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

Let us dispense with conditioning on $\theta = 1$, t_v , $F_{v,T}$, and W_v . For $\theta = 0$, one similarly derives the same bound (D.5) and, hence, it holds unconditionally for θ . Since $\mathbb{P}(\cdot) \leq \mathbb{P}(\cdot \mid W_v) + (1 - \mathbb{P}(W_v))$, we dispense with conditioning on W_v at the cost of increasing the upper bound by ψ . It remains to average the bound over $F_{v,T}^d$ and t_v . The size $|F_{v,T}^d|$ is uniformly distributed over $\{0, 1, \dots, d\}$ because it is the length of the prefix of v in a random permutation of $F_v^d \cup \{v\}$. Thus, averaging the second summand in (D.5) results in the geometric progression of length d . Taking into account that t_v is uniformly distributed on $[0, 1]$ and bounding the geometric progression of length d by the infinite one, we get

$$\mathbb{P}(a'_v \neq \theta) \leq \psi + \frac{12}{\sqrt{d-1}(2p-1-\psi)} + \frac{1}{d+1} \cdot \frac{1}{1 - \exp\left(-\frac{1}{8}\frac{(2p-1-\psi)^2}{(1-\psi)^2}\right)}. \quad (\text{D.6})$$

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 $\frac{2p-1-\psi}{1-\psi} \leq 1$, we obtain

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

where in the second inequality we bounded $(1-\psi)$ by 1 and used that $1 - \exp\left(-\frac{1}{8}\right) \geq \frac{1}{9}$. Denoting by b and c the second and the third summand in (D.6), respectively, we see that $c \leq \frac{1}{2}b^2$. For b below 1, we have $b^2 \leq b$, while for $b \geq 1$ the bound (D.6) becomes trivial anyway. Therefore,

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

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 D.2, we complete the proof of Theorem 3.1. \square

Proof of Lemma D.1: To ensure conditional independence, we make use of the following property of the realized subnetwork $G_{u,T} = (V_{u,T}, E_{u,T})$ of an agent $u \in V$. Conditional on $G_{u,T}$, his action a_u is independent of $(s_z, a_z)_{z \in V \setminus V_{u,T}}$ and $(t_z)_{z \in V}$. In other words, u 's action is solely determined by his realized subnetwork and signals of agents from this subnetwork. This property has an important consequence. Consider a group of agents $U \subset V$ and an event determined by the collection of arrival times $\{T = (t_z)_{z \in V} \in \mathcal{T}\}$ for some $\mathcal{T} \subset [0, 1]^V$. Then actions $(a_u)_{u \in U}$ are independent conditional on $\theta = \theta_0$ and $\{T \in \mathcal{T}\}$ provided that all the realized subnetworks $(G_{u,T})_{u \in U}$ are independent disjoint random networks conditional on $\{T \in \mathcal{T}\}$.

With this general observation, we will demonstrate that actions $(a_u)_{u \in F}$ are conditionally independent given $\theta = \theta_0$, $t_v = t$, $F_{v,T}^d = F$, and W_v . By the definition of W_v , the realized subnetwork of each $u_i \in F$ belongs to $B_{r,G \setminus v}(u_i)$ and thus the realized subnetworks $(G_{u,T})_{u \in F}$ are disjoint. It remains to check their independence.

Recall that the event W_v has the form $\cap_{u \in F_{v,T}^d} \{G_{u,T} \subset B_{r-1,G \setminus v}(u)\}$. We now check that each event $\{G_{u,T} \subset B_{r-1,G \setminus v}(u)\}$ is determined by arrival times t_z of agents $z \in B_{r,G \setminus v}(u) \cup \{v\}$. Indeed, $G_{u,T}$ belongs to $B_{r-1,G \setminus v}(u)$ if and only if u arrives earlier than v and for any path $(u = z_0, z_1, z_2, \dots, z_k)$ not passing through v and connecting u and the sphere $S_{r,G \setminus v}(u) = B_{r,G \setminus v}(u) \setminus B_{r-1,G \setminus v}(u)$ in the a-priori network G , there is an index i such that $t_{z_i} < t_{z_{i+1}}$. Consequently, we can rewrite the event $\{G_{u,T} \subset B_{r-1,G \setminus v}(u)\}$ as $\{(t_z)_{z \in B_{r,G \setminus v}(u)} \subset \mathcal{T}_u(t_v)\}$, where $\mathcal{T}_u(t_v)$ is a certain subset of $[0, 1]^{B_{r,G \setminus v}(u)}$ depending on the arrival time of v .

Conditioning on $F_{v,T}^d = F$, $t_v = t$, and W_v becomes equivalent to conditioning on the following family of events determined by arrival times: $\{(t_z)_{z \in B_{r,G \setminus v}(u)} \subset \mathcal{T}_u(t)\}$ for each $u \in F$, $t_u > t$ for $u \in F_v^d \setminus F$, and $t_v = t$. Arrival times $(t_z)_{z \in V}$ are unconditionally independent, and the condition restricts the values of disjoint subsets (here we use the fact that $B_{r,G \setminus v}(u)$ are disjoint for $u \in F_v^d$). Thus, conditionally on $F_{v,T}^d = F$, $t_v = t$, and W_v , the families of random variables $(t_z)_{z \in B_{r,G \setminus v}(u)}$ are independent across $u \in F$. The realized subnetwork $G_{u,T}$ of u is determined by $(t_z)_{z \in B_{r,G \setminus v}(u)}$ and, hence, the networks $(G_{u,T})_{u \in F}$ are also conditionally independent, which implies the desired conditional independence of actions $(a_u)_{u \in F}$. \square

Proof of Lemma D.2: Recall that $W_v = \cap_{u \in F_{v,T}^d} \{G_{u,T} \subset B_{r-1,G \setminus v}(u)\}$ and consider one of these events. We pick an agent $u \in F_v^d$ with $t_u < t_v$ and demonstrate that

$$\mathbb{P}(G_{u,T} \subset B_{r-1,G \setminus v}(u) \mid t_v, (t_z)_{z \in F_v^d}) \geq 1 - 2 \cdot \left(\frac{e \cdot D}{r}\right)^r. \quad (\text{D.7})$$

Note that the event $\{G_{u,T} \subset B_{r-1,G \setminus v}(u)\}$ is determined by the arrival times t_z of $z \in B_{r,G \setminus v}(u) \cup \{v\}$ only (see the argument in the proof of Lemma D.1). Hence, by the independence of arrival times, we can simplify the condition

$$\mathbb{P}(G_{u,T} \subset B_{r-1,G \setminus v}(u) \mid t_v, (t_z)_{z \in F_v^d}) = \mathbb{P}(G_{u,T} \subset B_{r-1,G \setminus v}(u) \mid t_v, t_u)$$

because $(F_v^d \setminus \{u\}) \cap B_{r,G \setminus v}(u) = \emptyset$. Since we condition on t_u , the only information added by knowing t_v is that *at the time u arrives, v is absent* (recall that we assume that $t_u < t_v$). Thus

$$\mathbb{P}(G_{u,T} \subset B_{r-1,G \setminus v}(u) \mid t_v, t_u) = \mathbb{P}_{G \setminus v}(G_{u,T} \subset B_{r-1,G \setminus v}(u) \mid t_u),$$

where $\mathbb{P}_{G \setminus v}$ refers to the probability with respect to the arrival process in the network $G \setminus v$. By Lemma C.1 applied to u in the network $G \setminus v$, we get

$$\mathbb{P}_{G \setminus v}(G_{u,T} \subset B_{r-1, G \setminus v}(u) \mid t_u) \geq 1 - 2 \cdot \left(\frac{e \cdot D}{r} \right)^r$$

and we deduce (D.7). The desired inequality for the probability of W_v follows from the union bound and (D.7). \square

Proof of Lemma D.3: In the proof of Lemma D.1 we observed that, once the realized subnetwork $G_{u,T}$ of an agent u is given, his action a_u is determined by the signals of agents from this subnetwork. We also saw that, conditional on $F_{v,T}^d = F$, $t_v = t$, and W_v , the realized subnetwork of $u \in F$ is determined by arrival times $(t_z)_{z \in B_r, G \setminus v(u)}$, and the condition is equivalent to $\{(t_z)_{z \in B_r, G \setminus v(u)} \subset \mathcal{T}_u(t)\}$ (we use the notation introduced in that proof). Therefore, the distribution of $G_{u,T}$ (and, hence, a_u) conditional on $\theta = \theta_0$, $F_{v,T}^d = F$, $t_v = t$, and W_v is the same no matter what other agents are in F . This observation allows us to simplify the condition:

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

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

$$\mathbb{P}(a_u \neq \theta \mid t_u < t, v \notin F_{u,T}, W_v) \leq \frac{\mathbb{P}(a_u \neq \theta \mid t_u < t, v \notin F_{u,T})}{1 - \psi} \quad (\text{D.8})$$

by 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 the wrong action $\mathbb{P}(a_u \neq \theta \mid A)$ is at most $1 - p$. Otherwise, the agent can profitably deviate from his equilibrium strategy by following his signal whenever A occurs.

The event $A' = \{t_u < t, v \notin F_{u,T}\}$ is not known to u since u does not observe his arrival time. However, we can approximate the event A' by the event $A = \{\hat{t}_u < t, v \notin F_{u,T}\}$, where $\hat{t}_u = \frac{|F_{u,T}|}{\deg(u) - 1}$ 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, hence, $\mathbb{P}(a_u \neq \theta \mid A) \leq 1 - p$. The conditional probability with respect to A' can be bounded as follows:

$$\begin{aligned} \mathbb{P}(a_u \neq \theta \mid A') &\leq \frac{\mathbb{P}(a_u \neq \theta, A) + \mathbb{P}(A' \setminus A)}{\mathbb{P}(A')} = \\ &= \mathbb{P}(a_u \neq \theta \mid A) \cdot \frac{\mathbb{P}(A)}{\mathbb{P}(A')} + \frac{\mathbb{P}(A' \setminus A)}{\mathbb{P}(A')} \leq (1 - p) \cdot \frac{\mathbb{P}(A)}{\mathbb{P}(A')} + \frac{\mathbb{P}(A' \setminus A)}{\mathbb{P}(A')}. \end{aligned}$$

Let us estimate all the probabilities in this expression. The probability $\mathbb{P}(A')$ can be computed explicitly as

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

To estimate $\mathbb{P}(A)$, we note that conditionally on t_u and $t_v > t_u$, the number of friends observed by u has the 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

$$\begin{aligned}\mathbb{P}\left(\frac{|F_{u,T}|}{\deg(u)-1} \leq t \mid t_u, v \notin F_{u,T}\right) &\leq \exp(-2(t_u - t)^2(\deg(u) - 1)) \quad \text{for } t \leq t_u \\ \mathbb{P}\left(\frac{|F_{u,T}|}{\deg(u)-1} \geq t \mid t_u, v \notin F_{u,T}\right) &\leq \exp(-2(t_u - t)^2(\deg(u) - 1)) \quad \text{for } t \geq t_u.\end{aligned}$$

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

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

Similarly,

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

Putting all the pieces together, we obtain

$$\begin{aligned}\mathbb{P}(a_u \neq \theta \mid A') &\leq (1-p) \left(1 + \frac{1}{t - \frac{t^2}{2}} \sqrt{\frac{\pi}{8(\deg(u) - 1)}}\right) + \frac{1}{t - \frac{t^2}{2}} \sqrt{\frac{\pi}{8(\deg(u) - 1)}} \leq \\ &\leq 1 - p + \frac{3}{t \cdot \sqrt{\deg(u) - 1}}.\end{aligned}$$

In the last inequality we took into account that $t - \frac{t^2}{2} \geq \frac{t}{2}$, $p \geq 0$, and $\sqrt{2\pi} \leq 3$.

Substituting this expression in (D.8) leads to the desired bound (D.1). \square

E Missed proofs for Section 5

Proof of Lemma 5.1. Consider the induced subnetwork $G^{V'}$ of G for some $V' \subset V$ with $|V'| = \lceil \alpha \cdot |V| \rceil$ and denote by $\deg'(v)$ the degree of an agent $v \in V'$ in $G^{V'}$. Fix positive $\gamma \leq \alpha$. 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 (B.1) (see mixing lemma B.1) to these V^1 and V^2 . Since $E(V^1, V^2) > (1 - \gamma)D \cdot |V^1|$ and $(1 - \alpha)|V| - 1 < |V^2| \leq (1 - \alpha)|V|$, we get

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

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

$$(\alpha - \gamma)D \sqrt{|V^1|} \leq |\lambda_2| \sqrt{(1 - \alpha)|V|}$$

and, therefore,

$$|V^1| \leq \frac{(1-\alpha)|\lambda_2|^2}{(\alpha-\gamma)^2 D^2} |V| \leq \frac{(1-\alpha)|\lambda_2|^2}{\alpha(\alpha-\gamma)^2 D^2} |V'|.$$

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

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

Applying Theorem 3.1 combined with Lemma 4.1 to each agent in 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 $G^{V'}$:

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

where $r = \left\lfloor \frac{g-3}{2} \right\rfloor$.

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

Opening the brackets and dispensing with positive terms, we get

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

Picking $\gamma = \frac{2\alpha}{3}$ and assuming that $\frac{\alpha}{3} - \frac{1}{D} \geq \frac{\alpha}{4}$ we obtain the desired bound (5.1). In the complementary case of small $\alpha < \frac{12}{D}$ the bound (5.1) trivializes since $\frac{2}{\sqrt{\alpha}} \delta(p, D, r, D) \geq \frac{\sqrt{D}}{\sqrt{3}} \cdot \frac{18}{\sqrt{D}} > 1$. \square

Lemma E.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 $V_{<\frac{D}{2}, \geq \frac{D}{2}}$ the set of agents with less than $\frac{D}{2}$ friends with degree $\frac{D}{2}$ or higher. We know $|V_{<D}| < \beta|V|$ and want to prove that $|V_{<\frac{D}{2}, \geq \frac{D}{2}}| < 2\beta|V|$. Assume by way of contradiction that $|V_{<\frac{D}{2}, \geq \frac{D}{2}}| \geq 2\beta|V|$. 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 conclude that

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

This contradiction completes the proof. \square

F Randomized robustness

In Section 5, we demonstrated existence of a-priori networks satisfying a very strong notion of adversarial robustness: the network aggregates information even if a subset of agents is eliminated in an adversarial way. Here we consider a weaker notion of *robustness to random elimination* and show that *any network* has this property. Namely, for networks with high learning quality, 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.

We will demonstrate the randomized robustness under an additional assumption about the information available to agents. Recall that $G_{v,T} = (V_{v,T}, E_{v,T})$ denotes the realized subnetwork of an agent v ; see the definition in Section 3.1.

Assumption F.1. Each agent v , in addition to his signal and actions of friends who arrived earlier than him, observes the set of agents $V_{v,T}$ of 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 $I_v = (s_v, (a_u)_{u \in F_{v,T}}, V_{v,T})$.

Remark F.1 (The role of Assumption F.1). We believe that this assumption plays a technical role and the randomized robustness must be ubiquitous without it as well; however, we were unable to get rid of it in our proof. Assumption F.1 ensures the sequential structure of equilibrium. If $V_{v,T} = \emptyset$, agent v follows his signal. If $V_{v,T}$ is non-empty with $|V_{v,T}| = k$, the equilibrium strategy $\sigma_v(s_v, (a_u)_{u \in F_{v,T}}, V_{v,T})$ is the optimal reply to $(\sigma_u)_{u \in V_{v,T}}$ with $|V_{u,T}| \leq k - 1$ (since $V_{u,T}$ is a strict subset of $V_{v,T}$). This sequential structure implies that an agent gains no advantage from learning the set of agents who have not yet arrived, the property critical in the proof of Theorem F.1, the main result of this section.

For a given network G and the probability p of the correct signal, denote by $L(G)$ the learning quality for the best equilibrium: $L(G) = \max_{\sigma} L_{\sigma}(G)$.

Theorem F.1 (Learning is robust to random elimination). *Under the Assumption F.1, consider a network $G = (V, E)$ that has learning quality $L(G) = 1 - \delta$ with some $\delta > 0$.*

Fix $\alpha \in (0, 1)$ and pick a subset $V' \subset V$ with $\lceil \alpha \cdot |V| \rceil$ agents uniformly at random. Then 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 $1 - \sqrt{\frac{\delta}{\alpha}}$ with respect to the choice of V' .

Proof of Theorem F.1. The argument is based on a coupling of the learning process in 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 $\lceil \alpha \cdot |V| \rceil$ earliest arrivals.

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 Assumption¹⁴ F.1.

The constructed coupling allows us 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

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

¹⁴In the game played over V' , all the present agents know that agents from $V \setminus V'$ are absent. In particular, without Assumption F.1 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 in V may no longer be a best reply in the game restricted to V' , even if all other agents maintained their strategy.

Using a rough estimate $\mathbb{P}(a_v = \theta \mid v \notin V') \leq 1$ 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:

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

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

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

Application of the Markov inequality completes the proof. \square

Itai Arieli
Technion-Israel Institute of Technology
Haifa, Israel
Email: iarieli@technion.ac.il

Fedor Sandomirskiy
Technion-Israel Institute of Technology
Haifa, Israel
Higher School of Economics (HSE)
St.Petersburg, Russia
Email: sandomirski@yandex.ru

Rann Smorodinsky
Technion-Israel Institute of Technology
Haifa, Israel
Email: rann@technion.ac.il