

Strategic Learning and the Topology of Social Networks

Elchanan Mossel*, Allan Sly† and Omer Tamuz‡

May 19, 2015

Abstract

We consider a group of strategic agents who must each repeatedly take one of two possible actions. They learn which of the two actions is preferable from initial private signals, and by observing the actions of their neighbors in a social network.

We show that the question of whether or not the agents learn efficiently depends on the topology of the social network. In particular, we identify a geometric “egalitarianism” condition on the social network that guarantees learning in infinite networks, or learning with high probability in large finite networks, in any equilibrium. We also give examples of non-egalitarian networks with equilibria in which learning fails.

Keywords: Social learning, informational externalities, social networks, aggregation of information.

1 Introduction

Consider a group in which each agent faces a repeated choice between two actions. Initially, the information available to each agent is a private signal, which gives a noisy indication of which is the correct action. As time progresses, the agents learn more by observing the actions of their neighbors in a social network. They do not, however, obtain any direct indication of the payoffs from their actions. For example, their choice could be one of lifestyle, where one can learn by observing the actions of others, but where payoffs (e.g., longevity) are only revealed after a large amount of time¹.

We are interested in the question of *learning*, or *aggregation of information*: When is it the case that, through observing each other, the agents exchange enough information to converge to the correct action? In particular, we are interested in the role that the geometry of the social network plays in this process, and in its effect on learning. Which social

*University of Pennsylvania and University of California, Berkeley. E-mail: mossel@stat.berkeley.edu. Supported by NSF award DMS 1106999, by ONR award N000141110140 and by ISF grant 1300/08.

†University of California, Berkeley. Supported by a Sloan Research Fellowship in mathematics and by NSF award DMS 1208339

‡Massachusetts Institute of Technology and Microsoft Research New England. This research was supported in part by a Google Europe Fellowship.

¹Consider parents who, each night, decide whether to lay their baby to sleep on its back or on its stomach. They can learn by observing the actions of their peers, but presumably do not receive any direct feedback regarding the effect of their actions on the baby’s health.

networks enable the flow of information, and which impede it? This problem has been studied extensively in the literature, using mostly boundedly-rational or heuristic approaches [10, 12, 5, 11, 16, 17]. However, the basic question of how *strategic* agents behave in this setting has been largely ignored², perhaps because the model is mathematically difficult to approach, or because strategic behavior seems unfeasible³. This article aims to fill this gap. We define a notion of egalitarianism for social networks, and show that when agents are strategic, learning always occurs on egalitarian social networks, and may not occur on those that are not egalitarian. Interestingly, these results broadly resemble those of some of the heuristic models (see, e.g., Golub and Jackson [16, Theorem 1]).

We call a social network graph (d, L) -*egalitarian* if it satisfies the following two conditions: (1) At most d edges leave each node (that is, each agent observes at most d others), and (2) whenever there is an edge from node i to j , there is a path *back* from j to i , of length at most L (that is, no agent is too far removed from those who observe her). In this article we show that on connected (d, L) -egalitarian graphs the agents learn the correct action, and give examples of non-egalitarian graphs in which learning fails.

Our model is a discounted, repeated game with incomplete information. We consider a state of nature S which is equal to either 0 or 1, with equal probability. Each agent receives a private signal that is independent and identically distributed conditioned on S , and is correlated with S . In each discrete time period t , each agent i chooses an action A_t^i taking values in $\{0, 1\}$. The information available to her is her own private signal, as well as the actions of her social network neighbors in the previous time periods. Agent i 's stage utility at time period t is equal to 1 if $A_t^i = S$ and to 0 otherwise, and is discounted exponentially, by a common rate. We consider general Nash equilibria, and show that they indeed exist (Theorem D.5); this does not follow from standard results.

We say that agent i *learns* S when A_t^i is equal to S from some time on, and that *learning* takes place when all agents learn S . Our main result (Theorem 1) is that on connected (d, L) -egalitarian graphs, in any equilibrium, learning occurs with high probability on large graphs, and with probability one on infinite graphs. We do not impose unbounded likelihood ratios: learning occurs in egalitarian networks even for weak - but informative - signals (contrast this with the sequential learning case of Smith and Sørensen [28]). Note that this applies to all Nash equilibria, and therefore in particular to any perfect Bayesian equilibria. We also provide examples of equilibria on large non-egalitarian networks in which, with non-vanishing probability, the agents do not learn.

Our results require a smoothness condition on private signals: each private belief (the probability that $S = 1$, conditioned on the private signal) must have a non-atomic distribution. This ensures that agents are never (i.e., with zero probability) indifferent. Our results do not, in general, hold without this condition; indifference can impede the flow of information (see, e.g., [23, Example A.1]). While real life signals are arguably always discrete or even finite, we propose that even with this requirement it is still possible to model or approximate a large range of signals.

The model we study makes heavy demands on the agents in terms of rationality, common knowledge, and human computation: agents are assumed to maximize a complicated

²Notable exceptions are [23] and [3]; we discuss these below.

³See, e.g., Bala and Goyal [5]: “to keep the model mathematically tractable... this possibility [strategic agents] is precluded in our model... simplifying the belief revision process considerably.”

expected utility function, to know the structure of the entire social network, and to precisely make complicated inferences regarding the state of nature. While our approach is standard in this literature (see, e.g., [13, 26, 1]), these features of our model prompt us to present our results as benchmarks, rather than as predictive statements about the world.

The rest of this article proceeds as follows. In section 2 we discuss an example of a (d, L) -egalitarian graph, using it to provide intuition into the ideas behind our main result (Theorem 1). In Section 3 we provide two examples of non-egalitarian graphs on which the agents fail to learn. In Section 4 we introduce our model formally. In Section 5 we explore the question of agreement and show that indeed the agents all converge to the same action. Section 6 includes our main technical contribution: a topology on equilibria of this game, as seen from the point of view of a particular agent. In Section 7 we prove Theorem 1, and Section 8 provides a conclusion.

1.1 Related literature

Learning on social networks is a widely studied field; a complete overview is beyond the scope of this paper, and so we shall note only a few related studies.

Bala and Goyal [5] study a similar model, and show results of learning or non-learning in different cases. Their model is boundedly-rational, with agents not taking into account the choices of their neighbors when forming their beliefs. Other notable bounded rationality models of learning through repeated social interaction are those of DeGroot [10], Ellison and Fudenberg [12], DeMarzo, Vayanos and Zwiebel [11], Golub and Jackson [16] and recently Jadbabaie, Molavi and Tahbaz-Salehi [17]. Interestingly, a recurring theme is that learning is facilitated by graphs which are egalitarian, although notions of egalitarianism differ across models (see, e.g., Golub and Jackson [16, Property 2]).

In a previous paper [23], we consider the same question, but for myopic agents. The analysis in that case is far simpler and does not require the technical machinery that we construct in this article. More importantly, the conditions for learning are qualitatively different for myopic agents, as compared to those for strategic agents: in the myopic setting, the upper bound on the number of observed neighbors is not needed. In fact, myopic agents learn with high probability on networks with no uniform upper bound. Thus there are examples of graphs on which myopic agents learn but strategic agents do not. We elaborate on this in our second example of non-learning, in Section 3.

In concurrent work by Arieli and Mueller-Frank [3], learning results are derived in a strategic setting with richer actions spaces; they study models in which actions are rich enough to reveal beliefs, and show that in that case learning occurs under general conditions, and in particular for any graph topology. To the best of our knowledge, no previous work considers learning, in repeated interaction, on social networks, in a fully rational, strategic setting.

The study of *agreement* (rather than learning) on social networks is also related to our work, and in fact we make crucial use of the work of Rosenberg, Solan and Vieille [26], who prove an agreement result for a large class of games with informational externalities played on social networks. This is a field of study founded by Aumann’s “Agreeing to disagree” paper [4], and elaborated on by Sebenius and Geanakoplos [27], McKelvey and Page [20],

Parikh and Krasucki [25], Gale and Kariv [13], Ménager [21] and recently Mueller-Frank [24], to name a few. The moral of this research is that, by-and-large, rational agents eventually reach consensus, even in strategic settings. We elaborate on the work of Rosenberg, Solan and Vieille [26] and show that when private signals are non-atomic then, asymptotically, agents agree on best responses (Theorem 5.1). This agreement result is an important ingredient of our main learning result (Theorem 1).

Another strain of related literature is that of *herd behavior*, started by Banerjee [6] and Bikhchandani, Hirshleifer and Welch [8], with significant generalizations and further analysis by Smith and Sørensen [28], Acemoglu, Dahleh, Lobel and Ozdaglar [1] and recently Lobel and Sadler [19]. Here, the state of nature and private signals are as in our model, and agents are rational. However, in these models agents act sequentially rather than repeatedly. The same informational framework is also shared by models of committee behavior and committee mechanism design (cf. Laslier and Weibull [18], Glazer and Rubinstein [14]).

1.2 Acknowledgments

We would like to thank Shahar Kariv for introducing us to this field. For commenting on drafts of this paper we would like to thank Nageeb Ali, Ben Golub, Eva Lyubich, Markus Mobius, Ariel Rubinstein, Ran Shorrer, Glen Weyl, and especially Scott Kominers.

2 An illustrative example




Figure 1: Learning in symmetric equilibria on the two dimensional grid.

To provide some intuition for why agents learn on egalitarian graphs (Theorem 1) we consider the simple, particular case that the graph is the undirected, infinite two dimensional grid, in which each agent has four neighbors. This is a $(4, 1)$ -egalitarian graph, and so Theorem 1 says that the agents learn S , or that, equivalently, in any equilibrium the actions of every agent converge to S . We now explain why this is indeed the case, under the further restriction to symmetric equilibria.

The first step in proving that all agents converge to S is to show that all agents converge to the same action, which we do in Theorem 5.1. This result uses - and perhaps elucidates - an important theorem of Rosenberg, Solan and Vieille [26], who consider the question of when agents eventually agree, regardless of whether or not they learn. For a large class of

games which includes the one we consider, they show that agents can disagree only if they are indifferent. Our additional requirement of non-atomic private signals allows us to rule out the possibility of indifference, and show that all agents converge to the same action⁴.

Having established that all agents converge to the same action, we use the fact that the graph is symmetric, as is the equilibrium. Hence all agents converge at the same ex-ante rate, and therefore, at some large enough time T , any particular agent will have converged, except with some very small probability ϵ . Of course, since the graph is infinite, there will be at time T many agents who have yet to converge. However, if we consider any one agent (or two, as we do immediately below), the probability of non-convergence is negligible.

Now, consider two agents which are more than $2T$ edges apart on the graph (see Figure 1), and condition on the state of nature S equaling one. The two agents' actions at time T are independent random variables (conditioned on the state of nature), as they are too far apart for any information to have been exchanged between them. On the other hand, since all agents converge to the same action, these independent random variables are equal (except with probability $\sim 2\epsilon$); the two agents somehow, with high probability, reach the same conclusion independently.

Now, two independent random variables that are equal must be constant. The agents' actions at time T are equal with high probability and independent conditioned on S , and so are with high probability equal to some fixed action. Since the agents' signals are informative, this action is more likely to equal the state of nature than not (Claim G.4). Since this holds for every $\epsilon > 0$, every agent's limit action must equal the state of nature.

2.1 General egalitarian graphs

The formalization and extension of this intuition to general egalitarian graphs and general (i.e., non-symmetric) equilibria requires a significant technical effort, and in fact the construction of novel tools for the analysis of games on networks; to this we devote most of the rest of this article. We now provide an overview of the main ideas.

The main notion we use is one of *compactness*. The two dimensional grid graph “looks the same” from the point of view of every node: there is only one “point of view” in this graph. Such graphs as known as *transitive graphs* in the mathematics literature. Note that this is the only property of the grid that we used in the proof sketch above, and therefore the same idea can be applied to all symmetric equilibria on infinite, connected transitive graphs.

We formalize a notion of an “approximate points of view”. We show that in particular, in an egalitarian graph, the nodes of the graph can be grouped into a finite number of sets, where from each set the graph “looks *approximately* the same”. Formally, we construct a topology in which the set of points of view in a graph is precompact if and only if the graph is (d, L) -egalitarian for some d and L (Theorem A.3). In this sense, egalitarianism, in which the set of points of view is precompact, is a relaxation of transitivity, in which the set of points of view is a singleton. Indeed, transitivity is an extreme notion of egalitarianism, by any reasonable definition of an egalitarian graph.

This property of egalitarian graphs allows us to apply the intuition of the above example (or, more precisely, a similar intuition) to any infinite, (d, L) -egalitarian graph (Theorem I.1).

⁴In fact, Theorem 5.1 does not exclude the case that no agent converges at all; we will, for now, ignore this possibility.

The fact that general equilibria are not symmetric is similarly treated by establishing that the space of equilibria is compact (Claim D.4). The theorem on finite graphs is proved by reduction to the case of infinite graphs.

Our main technical innovation is the construction of a topology on equilibria of this game, as seen from the point of view of a particular agent (Section 6.2). In this topology, an equilibrium has a finite number of “approximate points of view” if and only if the graph is egalitarian (Claim D.3). This topology is also useful for showing that equilibria exist in the case of an infinite number of agents, which requires a non-standard argument (Theorem D.5). This technique should be applicable to the analysis of a large range of repeated, discounted games on networks.

3 Non-learning




Figure 2: The Royal Family. Each member of the public (on the left), observes each royal (on the right), as well as her next door neighbors. The royals observe each other, and one royal observes one member of the public.

We provide two example of non-egalitarian graphs in which the agents do not learn. In the first example (Figure 2), inspired by Bala and Goyal’s royal family graph [5], the social network has two groups of agents: a “royal family” clique of R agents who all observe each other, and n agents - the “public” - who are connected in an undirected chain, and additionally can observe all the agents in the royal family. Finally, a single member of the royal family observes one of the public, so that the graph is connected⁵. We think of R as fixed and consider the case of arbitrarily large n , or even infinite n .

While this graph satisfies condition (1) of egalitarianism, it violates condition (2). Therefore, Theorem 1 does not apply. Indeed, in the online appendix we construct an equilibrium for the game on this network, in which the agents of the public ignore their own private signals after observing the first action of the royal family, which provides a much stronger indication of the correct action. However, the probability that the royal family is wrong is independent of n : since the size of the royal family is fixed, with some fixed probability every one of its members is mislead by her private signal to choose the wrong action in the first period. Hence, regardless of how large society is, there is a fixed probability that learning does not occur.

⁵The graph is, in fact, *strongly connected*, meaning that there is a directed path connecting every ordered pair of agents.




Figure 3: The mad king. In this social network all edges are bi-directional. Each member of the people is connected only to the king, as is each member of the court. The members of the bureaucracy are connected only to the regent.

We construct a second example of non-learning which we call “the mad king”. Here, the graph is undirected, so that whenever i observes j then j observes i ; the graph therefore satisfies condition (2) of egalitarianism with $L = 1$, but not condition (1). The graph (see Figure 3) consists of five types: the people, the king, the regent, the court and the bureaucracy. There is one king, one regent, a fixed number of members of the court and a fixed number of members of the bureaucracy, which is much larger than the court. The number of people is arbitrarily large. They are connected as follows:

- The king is connected to the regent, the court and the people.
- The regent is connected to the king and to the bureaucracy.
- The members of the court are each connected only to the king.
- The members of the bureaucracy are each connected only to the regent.
- Importantly, the people are each connected only to the king, and not to each other.

For an appropriate choice of private signal distributions and discount factor, we construct an equilibrium in which all agents act myopically in the first two rounds, except for the people, who choose the constant action 0. This is enforced by a threat from the king, who, if any of the people deviate, will always play 1, denying them any information he has learned; the prize for complying is the exposure to a well informed action which first aggregates the information available to the court, and later aggregates the information available to the (larger) bureaucracy. The result is that the information in the people’s private signals is lost, and so we have non-learning with probability bounded away from 0, for graphs of arbitrarily large size.

The equilibrium path can be succinctly described as follows; we provide a complete description in the online appendix:

- The members of the bureaucracy act myopically in round 0, as do the members of the court.
- The regent, who learns by observing the bureaucracy, acts myopically at time 0 and at time 1. Therefore, and since the bureaucracy is large, his action at time 1 will be correct with some fixed high probability.
- The king acts myopically in round 0. At round 1, after having learned from the court's actions, the king again acts myopically, unless any of the people chose action 1 at round 0, in which case he chooses action 1 at this time and henceforth.
- The people choose action 0 in round 0. They have no incentive to deviate, since they stand to learn much from the king's actions, which, at the next round, will aggregate the information in the court's actions.
- By round 2, the king has learned from the regent's well informed action of round 1. He therefore, at round 2, emulates the regent's action of round 1, unless any of the people chose action 1 at rounds 0 or 1, in which case he again chooses action 1 at this time and henceforth.
- In round 1 the people again choose action 0. They again have no reason to deviate, this time because they wish to learn the regent's action, through the king; this information - which originates from the bureaucracy - is much more precise than that which the king collected from the court in the previous round and reveals to them in this round.
- At round 2 (and henceforth) the people emulate the king's previous action, and therefore the king will not learn from them.

It follows that the private signals of the people are lost, and so, regardless of the number of people, there is a fixed probability of non-learning.

We were not able to prove - or to disprove - that this equilibrium is a perfect Bayesian equilibrium. However, it intuitively seems likely that if the people were to deviate from the equilibrium, then the king would not have an incentive to carry out his threat. If this intuition holds then this is not a perfect Bayesian equilibrium.

An interesting phenomenon is that on this graph, the agents do learn S with high probability when they discount the future sufficiently, or in the limiting case that they are myopic (i.e., fully discount the future). This is thus an example - and perhaps a counter-intuitive one - of how strategic agents may learn less effectively than myopic ones.

4 Model

4.1 Informational structure

The structure of the private information available to the agents is the standard one used in the herding literature (see, e.g., Smith and Sørensen [28]).

We denote by V the set of agents, which we take to equal $\{1, 2, \dots, n\}$ in the finite case and $\mathbb{N} = \{1, 2, \dots\}$ in the (countably) infinite case. Let $\{0, 1\}$ be the set of possible values of the *state of nature* S , and let $\mathbb{P}[S = 1] = \mathbb{P}[S = 0] = 1/2$. Let Ω be a measurable space, called the space of *private signals*. Let $W_i \in \Omega$ be agent i 's private signal, and denote $\bar{W} = (W_1, W_2, \dots)$. Fix μ_0 and μ_1 , two mutually absolutely continuous probability measures on Ω . Conditioned on $S = 0$, let W_i be i.i.d. μ_0 , and conditioned on $S = 1$ let W_i be i.i.d. μ_1 .

The assumption that $\mathbb{P}[S = 1] = 1/2$ can be relaxed; in particular, for every choice of private signals there exist $p_1 < 1/2 < p_2$ such that our results apply when $\mathbb{P}[S = 1]$ is taken to be in (p_1, p_2) . However, when agents are myopic (or more generally discount the future enough), when priors are skewed, and when signals are weak, then regardless of the graph, and in any equilibrium, agents will disregard their private signals and only play the more a priori probable action. Indeed, if for example the prior is $\mathbb{P}[S = 1] = 0.8$, then for weak enough private signals it will be the case that $\mathbb{P}[S = 1|W_i] > 0.7$ with probability one. Myopic agents always choose the action that they deem more likely to equal the state of the world, and therefore will choose 1, as will agents who are not myopic but sufficiently discount the future. It follows that in this case observing others' actions will reveal no information, and no learning will occur. Therefore, to ensure learning, one must impose some conditions on the prior and the strength of the private signals; for example, a sufficient condition would be that $\mathbb{P}[S = 1|W_i]$ has positive probability of both being above half and of being below half. To avoid encumbering this paper with an additional layer of technical complexity, we focus on the case in which $\mathbb{P}[S = 1] = \mathbb{P}[S = 0] = 1/2$.

Agent i 's *private belief* I_i is the probability that the state of the world is 1, given i 's private signal:

$$I_i = \mathbb{P}[S = 1|W_i].$$

Since I_i is a sufficient statistic for S given W_i , we assume below without loss of generality that an agent's actions depend on W_i only through I_i (see, e.g., Smith and Sørensen [28]).

We consider only μ_0 and μ_1 such that the distribution of I_i is non-atomic. This is the condition that we refer to above as *non-atomic private beliefs*. This is an additional restriction that we impose, beyond what is standard in the herding literature.

4.2 The social network

The agents' social network defines which of them observe the actions of which others. We do not assume that this is a symmetric relation: it may be that i observes j while j doesn't observe i . Formally, the social network $G = (V, E)$ is a directed graph: V is the set of agents, and E is a relation on V , or a subset of the set of ordered pairs $V \times V$. The set of neighbors of $i \in V$ is

$$N(i) = \{j : (i, j) \in E\},$$

and we consider only graphs in which $i \in N(i)$; that is, we require that an agent observes her own actions. The *out-degree* of i is given by $|N(i)|$, and will always be finite. This means that an agent observes the actions of a finite number of other agents. We do allow infinite

in-degrees; this corresponds to agents whose actions are observed by infinitely many other agents.

Let $G = (V, E)$ be a directed graph. A (directed) *path* of length k from $i \in V$ to $j \in V$ in G is sequence of $k + 1$ nodes i_1, \dots, i_{k+1} such that $(i_n, i_{n+1}) \in E$ for $n = 1, \dots, k$, and where $i_1 = i$ and $i_{k+1} = j$.

A directed graph is *strongly connected* if there exists a directed path between every ordered pair of nodes; we restrict our attention to such graphs. Strong connectedness is natural in the contexts of agreement and learning; as an extreme example, consider a graph in which some agent observes no-one. In this graph we cannot hope for that agent to learn the state of nature.

A directed graph G is *L-locally-connected* if, for each $(i, j) \in E$, there exists a path of length at most L in G from j to i . Equivalently, G is *L-locally-connected* if whenever there exists a path of length k from i to j , there exists a path of length at most $L \cdot k$ from j back to i . Note that 1-locally-connected graphs are commonly known as undirected graphs.

As defined above, a graph is said to be (d, L) -egalitarian if all out-degrees are bounded by d , and if it is *L-locally-connected*.

4.3 The game

To model the agents' strategic behavior we consider the following game of incomplete information. This framework, with some variations, has been previously used, for example, by Gale and Kariv [13] and Rosenberg, Solan and Vieille [26].

We consider the discrete time periods $t = 0, 1, 2, \dots$, where in each period each agent $i \in V$ has to choose one of the actions in $\{0, 1\}$. The information available to i at time t is her own private signal (of which the relevant information is her private belief, taking values in $[0, 1]$), and the actions of her neighbors in previous time periods, taking values in $\{0, 1\}^{|N(i)| \cdot t}$. This action is hence calculated by some function from $[0, 1] \times \{0, 1\}^{|N(i)| \cdot t}$ to $\{0, 1\}$.

A *pure strategy at time t* of an agent $i \in V$ is therefore a Borel-measurable function $q_t^i : [0, 1] \times \{0, 1\}^{|N(i)| \cdot t} \rightarrow \{0, 1\}$. A *pure strategy* of an agent i is the sequence of functions $q^i = (q_0^i, q_1^i, \dots)$, where q_t^i is i 's pure strategy at time t . We endow the space of pure strategies with the topology derived from the weak topology on functions from $[0, 1]$ to $\{0, 1\}$.

A *mixed strategy* Q^i of agent i is a pure-strategy-valued random variable; this is the standard notion of a mixed strategy, and we shall henceforth refer to mixed strategies simply as strategies. A (mixed) *strategy profile* is a set of strategies $\bar{Q} = \{Q^i : i \in V\}$, where the random variables Q^i are independent of each other and of the private signals.

The *action* of agent i at time t is denoted by $A_t^i \in \{0, 1\}$. Denote the *history* of actions of the neighbors of i before time t by $A_{[0,t)}^{N(i)} = \{A_s^j : s < t, j \in N(i)\}$; this depends on the social network G . The action that agent i plays at time t under strategy profile \bar{Q} is

$$A_t^i = A_t^i(G, \bar{Q}) = Q_t^i \left(I_i, A_{[0,t)}^{N(i)} \right).$$

Note again that we (without loss of generality) limit the action to be a function of the private belief I_i , as opposed to the private signal W_i .

Let $0 < \lambda < 1$ denote the agents' common *discount factor*. Given a social network G and strategy profile \bar{Q} , agent i 's *stage utility at time t* , $U_{i,t}$, is 1 if her action matches S , and 0 otherwise:

$$U_{i,t} = U_{i,t}(G, \bar{Q}) = \mathbf{1}_{A_t^i(G, \bar{Q})=S}.$$

Her *expected stage utility at time t* , $u_{i,t}$, is therefore given by

$$u_{i,t} = u_{i,t}(G, \bar{Q}) = \mathbb{E} [U_{i,t}(G, \bar{Q})] = \mathbb{P} [A_t^i(G, \bar{Q}) = S].$$

Agent i 's *expected utility* u_i is given by

$$u_i = u_i(G, \bar{Q}) = (1 - \lambda) \sum_{t=0}^{\infty} \lambda^t u_{i,t}(G, \bar{Q}).$$

Note that $u_i \in [0, 1]$, due to the normalization factor $(1 - \lambda)$. A *game* \mathcal{G} is a 4-tuple $(\mu_0, \mu_1, \lambda, G)$ consisting of two measures, a discount factor and a social network graph, satisfying the conditions of the definitions above.

4.4 Equilibria

Our equilibrium concept is the standard Nash equilibrium in games of incomplete information: \bar{Q} is an equilibrium if no agent can improve her expected utility $u_i(\bar{Q})$ by deviating from \bar{Q} .

Formally, in a game $\mathcal{G} = (\mu_0, \mu_1, \lambda, G)$, strategy profile \bar{Q} is an equilibrium if, for every agent $i \in V$ it holds that

$$u_i(G, \bar{Q}) \geq u_i(G, \bar{R}),$$

for any \bar{R} such that $R^j = Q^j$ for all $j \neq i$ in V .

5 Agreement

Let the *infinite action set* C_i of agent i be defined by

$$C_i = C_i(G, \bar{Q}) = \{s \in \{0, 1\} : A_t^i(G, \bar{Q}) = s \text{ for infinitely many values of } t\}.$$

There could be more than one action that i takes infinitely often. In that case we write $C_i = \{0, 1\}$. Otherwise, with a slight abuse of notation, we write $C_i = 0$ or $C_i = 1$, as appropriate.

In this section we show that the agents reach consensus in any graph, in the following sense:

Theorem 5.1. *Let \mathcal{G} be a game with either finitely many players or countably infinitely many players, and let \bar{Q} be an equilibrium strategy profile of \mathcal{G} . Then, with probability one, $C_i = C_j$ for all agents $i, j \in V$.*

This theorem is a crucial ingredient in the proof of the main result of this article. Indeed, learning occurs if $C_i = S$ for all i , and so a prerequisite is that $C_i = C_j$ for all i, j .

Recall that a strategy of agent i at time t is a function of her private belief I_i and the actions of her neighbors in previous time periods, $A_{[0,t]}^{N(i)}$. Hence we can think of the sigma-algebra generated by these random variables as the “information available to agent i at time t ”. Denote the information available to agent i at time t by

$$\mathcal{F}_t^i = \mathcal{F}_t^i(G, \bar{Q}) = \sigma \left(I_i, Q^i, A_{[0,t]}^{N(i)} \right),$$

and denote by

$$\mathcal{F}_\infty^i = \mathcal{F}_\infty^i(G, \bar{Q}) = \sigma \left(\cup_{t=0}^\infty \mathcal{F}_t^i \right)$$

the information available to agent i at the limit $t \rightarrow \infty$. Note that \mathcal{F}_t^i includes the sigma-algebra generated by i 's private belief, the actions of i 's neighbors before time t , and i 's pure strategy; i knows which pure strategy she has chosen.

Since the expected stage utility of action s at time t is $\mathbb{P}[s = S]$, a myopic agent would take an action s in $\{0, 1\}$ that maximizes $\mathbb{P}[s = S | \mathcal{F}_t^i]$. This motivates the following definition. Denote the best response of agent i at time t by

$$B_t^i = B_t^i(G, \bar{Q}) = \operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[s = S | \mathcal{F}_t^i(G, \bar{Q})].$$

Likewise denote the *set* of best responses of agent i at the limit $t \rightarrow \infty$ by

$$B_\infty^i = B_\infty^i(G, \bar{Q}) = \operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[s = S | \mathcal{F}_\infty^i].$$

At any time t there is indeed almost surely only one action that maximizes $\mathbb{P}[s = S | \mathcal{F}_t^i(G, \bar{Q})]$, since we require that the distribution of private beliefs be non atomic. This does not necessarily hold at the limit $t \rightarrow \infty$, and so we let B_∞^i take the values 0, 1 or $\{0, 1\}$. Note that a reasonable conjecture is that the probability that $B_\infty^i = \{0, 1\}$ is zero, but we are not able to prove this. This does not, however, prevent us from proving our results, but it does complicate the proofs.

The following theorem is a restatement, in our notation, of Proposition 2.1 in Rosenberg, Solan and Vieille [26].

Theorem 5.2 (Rosenberg, Solan and Vieille). *For any agent i it holds that $C_i \subseteq B_\infty^i$ almost surely, in any equilibrium.*

That is, any action that i takes infinitely often is optimal, given all the information agent i eventually learns. Note that this theorem is stated in [26] for a finite number of agents. However, a careful reading of the proof reveals that it holds equally for a countably infinite set of agents. The same holds for their Theorem 2.3, in which they further prove the following agreement result.

Theorem 5.3 (Rosenberg, Solan and Vieille). *Let j be a neighbor of i . Then $C_j \subseteq B_\infty^i$ almost surely, in any equilibrium.*

Equivalently, if i observes j , and j takes an action a infinitely often, then a is an optimal action for i . If we could show that $B_\infty^i = C_i$ for all i , it would follow from these two theorems, and from the fact that the graph is strongly connected, that $C_i = C_j$ for all agents i and j ; the agents would agree on their optimal action sets. This is precisely what we show in Theorem E.3. Our agreement theorem (Theorem 5.1) is a direct consequence.

6 Topologies on graphs and strategy profiles

6.1 Rooted graphs and their topology

A *rooted graph* is a pair (G, i) , where $G = (V, E)$ is a directed graph, and $i \in V$ is a vertex in G .

Rooted graphs are a basic mathematical concept, and are important to the understanding of this game. This section starts with some basic definitions, continues with the definition of a metric topology on rooted graphs, and culminates in a novel theorem on compactness in this topology, which may be of independent interest. In this we follow our previous work [23], which builds on the work of others such as Benjamini and Schramm [7] and Aldous and Steele [2].

Intuitively, a rooted graph is a graph, as seen from the “point of view” of a particular vertex - the root. Two rooted graphs will be close in our topology if the two graphs are similar, as seen from the roots.

Before defining our topology we will need a number of standard definitions. Let $G = (V, E)$ and $G' = (V', E')$ be graphs, and let (G, i) and (G', i') be rooted graphs. A *rooted graph isomorphism* between (G, i) and (G', i') is a bijection $h : V \rightarrow V'$ such that

1. $h(i) = i'$.
2. $(j, k) \in E \Leftrightarrow (h(j), h(k)) \in E'$.

If there exists a rooted graph isomorphism between (G, i) and (G', i') then we say that they are isomorphic, and write $(G, i) \cong (G', i')$. Informally, isomorphic graphs cannot be told apart when vertex labels are removed; equivalently, one can be turned into the other by an appropriate renaming of the vertices. The isomorphism class of (G, i) is the set of rooted graphs that are isomorphic to it, and will be denoted by $[G, i]$.

Let j, k be vertices in a graph G . Denote by $\Delta(j, k)$ the length of the shortest (directed) path from j to k . In general, $\Delta(j, k) \neq \Delta(k, j)$, since the graph is directed. The (directed) *ball* $B_r(G, i)$ of radius r of the rooted graph (G, i) is the rooted graph, with root i , induced in G by the set of vertices $\{j \in V : \Delta(i, j) \leq r\}$.

We now proceed to define our topology on the space of isomorphism classes of strongly connected rooted graphs, which is an extension of the Benjamini-Schramm [7] topology on undirected graphs. We define this topology by a metric⁶.

⁶ This definition applies, in fact, to a larger class of directed graphs: a rooted graph (G, i) is *weakly connected* if there is a directed path from i to each other vertex in the graph. Note that indeed a strongly connected graph is necessarily weakly connected, but not vice versa. Note also that a rooted graph (G, i) is weakly connected if and only if for every vertex j there exists an r such that j is in $B_r(G, i)$.

Let $[G', i']$ and $[G, i]$ be isomorphism classes of strongly connected rooted graphs. The distance $D([G, i], [G', i'])$ is defined by

$$D([G, i], [G', i']) = \inf\{2^{-r} : B_r(G, i) \cong B_r(G', i')\}. \quad (1)$$

That is, the larger the radius around the roots in which the graphs are isomorphic, the closer they are. In fact, the quantitative dependence of $D(\cdot, \cdot)$ on r (exponential in our definition) will not be of importance here, as we shall only be interested in the topology induced by this metric.

It is straightforward to show that $D(\cdot, \cdot)$ is well defined; a standard diagonalization argument (which we use repeatedly in this article) is needed to show that it is indeed a metric rather than a pseudometric (Claim A.1). The assumption of strong connectivity is crucial here, since $D(\cdot, \cdot)$ is otherwise a pseudometric.

Let \mathcal{SCG} be the set of isomorphism classes of *strongly connected* rooted graphs. This set is a topological space when equipped with the topology induced by the metric $D(\cdot, \cdot)$. Given a strongly connected graph G , let $\mathcal{R}(G) \subset \mathcal{SCG}$ be the set of all rooted graph isomorphism classes of the form $[G, i]$, for i a vertex in G . This can be thought of as the set of “points of view” in the graph G . The notion of (d, L) -egalitarianism now arises naturally, in the sense that the number of “approximate points of view” in G is finite if and only if G is egalitarian. This is formalized in the following lemma.

Lemma 6.1. *Let G be a strongly connected graph. Then the closure of $\mathcal{R}(G)$ is compact in \mathcal{SCG} if and only if G is (d, L) -egalitarian, for some d and L .*

We would like to suggest that Lemma 6.1, which we prove in Appendix A, may be of independent mathematical interest, as it extends the well understood notion of compactness in undirected graphs to directed, strongly connected graphs.

6.2 The space of rooted graph strategy profiles and its topology

In this section we use the above topology on rooted graphs to construct a topology on what we call *rooted graph strategy profiles*. This will be the main tool at our disposal in proving both the existence of equilibria, and our main result, Theorem 1. Intuitively, a *rooted graph strategy profile* will be a graph, together with a strategy profile, as seen from the point of view of the root. As in the case of rooted graphs, two points in this space will be close if they look alike from the points of view of the roots.

Let $G = (V, E)$ and $G' = (V', E')$ be strongly connected directed graphs, and let $(G, i), (G', i') \in \mathcal{SCG}$ be rooted graphs. Let \bar{Q} and \bar{R} be strategy profiles for the agents in V and V' , respectively. We say that the triplet (G, i, \bar{Q}) is equivalent to the triplet (G', i', \bar{R}) if there exists a rooted graph isomorphism h from (G, i) to (G', i') such that $\bar{Q}^j = \bar{R}^{h(j)}$ for all $j \in V$. The *rooted graph strategy profiles* \mathcal{GS} are the set of equivalence classes induced by this equivalence relation. We denote an element of \mathcal{GS} by $[G, i, \bar{Q}]$.

In Appendix C we apply the classical work of Milgrom and Weber [22] to define a metric d on a single agent’s strategy space, with the property that when the number of agents is finite then utilities are continuous in the induced topology.

We use this metric, and the metric of rooted graphs to define a metric on rooted graph strategy profiles. Intuitively, $[G, i, \bar{Q}]$ and $[G', i', \bar{R}]$ will be close in this metric if, in a large

radius around i and i' , it holds both that the graphs are isomorphic and that the strategies are similar.

Let d be a metric on a single agent's strategy space. Let i and i' be agents in graphs G and G' , respectively. We can use d as a metric between their strategies, as long as we uniquely identify each neighbor of one with a neighbor of the other. Let h be a bijection between $N(i')$ and $N(i)$. Then $d_h(Q^i, Q^{i'})$ will denote the distance thus defined between Q^i and $Q^{i'}$.

We next define $D_r(\cdot, \cdot)$, a pseudometric on graph strategy profiles which only takes into account the graph and the strategies at balls of radius r around the root. Two graph strategy profiles are close in D_r if (1) these balls are isomorphic, so that agents in these balls can be identified, and if (2) under some such identification, identified agents have similar strategies. This is a pseudometric rather than a metric since there could be two graph strategy profiles that are at distance 0 under D_r , but are not identical; differences will, however, occur only at distances that are larger than r from the roots.

Let $[G, i, \bar{Q}]$ and $[G', i', \bar{R}]$ be rooted graph strategies. For $r \in \mathbb{N}$, let $H(r)$ be the (perhaps empty) set of rooted graph isomorphisms between $B_r(G, i)$ and $B_r(G', i')$. Let

$$D_r\left([G, i, \bar{Q}], [G', i', \bar{R}]\right) = \min_{h \in H(r+1)} \max_{j \in B_r(G, i)} d_h(Q^j, R^{h(j)}),$$

when $H(r+1)$ is non-empty, and 1 otherwise. The choice of $h \in H(r+1)$ and then $j \in B_r(G, i)$ guarantees that h is a bijection from the set of neighbors of j to the set of neighbors of $h(j)$.

Finally, define the metric $D([G, i, \bar{Q}], [G', i', \bar{R}])$ by

$$D\left([G, i, \bar{Q}], [G', i', \bar{R}]\right) = \inf_{r \in \mathbb{N}} \left\{ \max \left\{ 2^{-r}, D_r([G, i, \bar{Q}], [G', i', \bar{R}]) \right\} \right\}. \quad (2)$$

Note that D will be small whenever D_r is small for large r . It is straightforward (if tedious) to show that $D(\cdot, \cdot)$ is indeed a well defined metric.

6.3 Properties of the space of rooted graph strategy profiles

Two rooted graph strategy profiles will be close in the topology induced by D if, in a large neighborhood of the roots, it holds both that the graphs are isomorphic, and also that the strategies are similar. This captures the root's "point of view" *of the entire strategy profile*.

While many possible topologies may have this property, this topology has some technical features that make it a useful analytical tool. First, expected utilities are continuous in this topology. Formally, let the *utility map* $u : \mathcal{GS} \rightarrow \mathbb{R}$ be given by

$$u([G, i, \bar{Q}]) = u_i(G, \bar{Q}).$$

This is a straightforward recasting of the previous definition of expected utility into the language of rooted graph strategy spaces. In Lemma D.1 we show that $u : \mathcal{GS} \rightarrow \mathbb{R}$ is continuous; this follows from the fact that payoff is discounted, and so the strategies of far away agents have only a small effect on an agent's utility. Another property of this topology that makes it applicable is that the set of *equilibrium* rooted graph strategy profiles

is closed (Lemma D.2). These properties are also instrumental in proving that equilibria exist (Theorem D.5).

Additionally, the probability of learning is lower semi-continuous in this topology. Let the *probability of learning* map $p : \mathcal{GS} \rightarrow \mathbb{R}$ be given by

$$p([G, i, \bar{Q}]) = \lim_{t \rightarrow \infty} \mathbb{P} [A_t^i(G, \bar{Q}) = S].$$

In Section G we prove that p is well defined and that it is lower semi-continuous (Theorem G.5). We also show that $p([G, i, \bar{Q}]) = 1$ if and only if the agents learn; i.e., if and only if $\lim_t A_t^j = S$ almost surely for all agents j in G (Claim G.3).

Finally, if G is an egalitarian graph, then the set of rooted graph strategy profiles on G is precompact (Claim D.3). Intuitively, this means that when G is egalitarian then not only are there finitely many approximate points of view of the graph (as discussed above), but also just finitely many approximate points of view of the strategy profile.

7 Learning

7.1 Learning on infinite egalitarian graphs

Let G be an infinite, connected, (d, L) -egalitarian graph, and let \bar{Q} be an equilibrium strategy profile. In this section we show that all agents learn S almost surely.

Recall that all agents converge to the same (random) action or set of actions. Denote by \hat{S}_∞ the random variable that is equal to 0 if all agents converge to 0 and is equal to 1 if all agents converge to 1, or if they all do not converge. Our choice of notation here follows from the fact that \hat{S}_∞ is a maximum a posteriori (MAP) estimator of any particular agent, given all that it learns: namely, the probability that an agent learns S is equal to the probability that \hat{S}_∞ equals S (Claim G.1). Since the private signals are informative, $\hat{S}_\infty = S$ with probability which is strictly greater than one half (Claim G.4), so \hat{S}_∞ is a non-trivial estimator of S .

Note that \hat{S}_∞ is measurable in the sequence of every agent's actions. Hence each agent eventually learns it, or something "close to it" at large finite times: formally, for every $\delta > 0$ there will be a time t and random variable $\hat{S}_\infty^{i,\delta}$ that can be calculated by i at time t , and such that $\mathbb{P} [\hat{S}_\infty^{i,\delta} = \hat{S}_\infty] > 1 - \delta$.

Now, \hat{S}_∞ is a deterministic function of the agents' private signals and pure strategies. Hence (e.g., by the martingale convergence theorem) \hat{S}_∞ is an *almost deterministic* function of the private signals and pure strategies of a large but *finite* group of agents. Formally, for every $\epsilon > 0$ there is a random variable \hat{S}_∞^ϵ that depends only on the private signals and pure strategies of some finite set of agents V^ϵ , and such that $\mathbb{P} [\hat{S}_\infty^\epsilon = \hat{S}_\infty] > 1 - \epsilon$.

Let i be an agent who is far away (in graph distance) from V^ϵ , so that the nearest member of V^ϵ is at *distance* at least t from i . Then everything that i observes up to *time* t is independent of \hat{S}_∞^ϵ , and hence "approximately independent" of \hat{S}_∞ (Claim H.4); we formalize a notion of "approximate independence" in Section F.

Now, as we note above, i eventually learns \hat{S}_∞ (or more precisely an estimator $\hat{S}_\infty^{i,\delta}$ that is equal to \hat{S}_∞ with high probability), gaining a new estimator of S which is (approximately)

independent of any estimators that it has learned up to time t . What we have so far outlined can thus be summarized informally as follows: the estimator \hat{S}_∞ is “decided upon” by a finite group of agents. When those far away eventually learn it they gain a new, approximately independent estimator of S .

We apply this argument inductively, relying crucially on the fact that the space of rooted graph strategies on G is precompact: Assume by induction that for every “point of view” $[H, j, \bar{R}]$ in the closure of this space there is an agent i in H with $k - 1$ approximately independent estimators of S by some time t . By compactness and the infinitude of G , there are infinitely many agents in G whose points of view are approximately equal to that of such an agent i . These will all also have $k - 1$ approximately independent estimators of S by time t . Some (in fact, almost all) of these agents will be sufficiently far from V^ϵ . These will then gain a new estimator when they eventually learn \hat{S}_∞ .

Hence in egalitarian graph, for any k and any degree of approximation, there will always be an agent who, given enough time, will accumulate k approximately independent estimators of S (Lemma H.1). A standard concentration of measure inequality then guarantees that the agent’s probability of learning will be approximately 1 (Theorem I.1). This proves that the agents learn on infinite graphs.

7.2 Learning on finite egalitarian graphs and the proof of Theorem 1

We reduce the case of finite graphs to that of infinite graphs, thus proving our main theorem.

Theorem 1. *Fix the distributions of the agents’ private signals, with non-atomic private beliefs. Fix also a discount factor $\lambda \in (0, 1)$, and positive integers L and d . Then in any connected, (d, L) -egalitarian, countably infinite network*

$$\mathbb{P}[\text{all agents learn } S] = 1$$

in any equilibrium. Furthermore, for every $\epsilon > 0$ there exists an n such that for any connected, (d, L) -egalitarian network with at least n agents

$$\mathbb{P}[\text{all agents learn } S] \geq 1 - \epsilon,$$

in any equilibrium.

Given a set of graphs \mathcal{K} , let $\mathcal{R}(\mathcal{K})$ be the set of rooted graphs $[G, i]$ such that $G \in \mathcal{K}$. Let $\mathcal{EQ}(\mathcal{K})$ be the set of equilibrium strategy profiles $[G, i, \bar{Q}]$ such that $G \in \mathcal{K}$.

Proof of Theorem 1. Let G be a (d, L) -egalitarian graph. The case that G is infinite is treated in Theorem I.1.

We hence consider finite graphs. Let \mathcal{K}_n be the set of L -locally-connected, degree d graphs with n vertices. Since \mathcal{K}_n is finite then $\mathcal{R}(\mathcal{K}_n)$ is finite and hence compact. It follows that $\mathcal{EQ}(\mathcal{K}_n)$ is also compact (Claim D.4). Since the map p is lower semi-continuous it attains a minimum on $\mathcal{EQ}(\mathcal{K}_n)$. Let $[G_n, i_n, \bar{Q}_n]$ be a minimum point, and denote $q(n) = p([G_n, i_n, \bar{Q}_n])$. We will prove the claim by showing that $\lim_n q(n) = 1$. Let $\{q(n_k)\}_{k=1}^\infty$ be a subsequence such that $\lim_k q(n_k) = \liminf_n q(n)$.

Since the set of (d, L) -egalitarian graphs is compact (Theorem A.2), by again invoking Claim D.4, we have that the sequence $\{[G_{n_k}, i_{n_k}, \bar{Q}_{n_k}]\}_{k=1}^\infty$ has a converging subsequence that must converge to some *infinite* L -locally-connected, degree d equilibrium graph strategy $[G, i, \bar{Q}]$. By the above, we have that $p([G, i, \bar{Q}]) = 1$, and so, by the lower semi-continuity of p , it follows that

$$\liminf_{n \rightarrow \infty} q(n) = \lim_k q(n_k) = \lim_{k \rightarrow \infty} p([G_{n_k}, i_{n_k}, \bar{Q}_{n_k}]) \geq p([G, i, \bar{Q}]) = 1.$$

□

8 Conclusion

8.1 Summary

Learning on social networks by observing the actions of others is a natural phenomenon that has been studied extensively in the literature. However, the question of how strategic agents fare has been largely ignored. We tackle this problem in a standard framework of a discounted game of incomplete information and conditionally independent private signals.

We show that on some networks agents learn in every equilibrium, and that they do not necessarily learn on others. The geometric condition of learning is one of egalitarianism, and is similar in spirit to conditions of learning identified in some boundedly-rational models.

8.2 Extensions and open problems

Our techniques, by their topological nature, give only asymptotic results: we show that the probability that agents learn on a (d, L) -egalitarian graph with n agents tends to one. It may be interesting to study the rate at which this happens, but our techniques do not seem to apply to this question.

Natural extensions of our model include those in which agents do not act synchronously, and those in which the agents do not know the structure of the graph, but have some prior regarding it. The latter is particularly compelling, since the assumption that the agents know exactly the structure of the graph is a strong one, especially in the case of large graphs.

We believe that our results should extend to these cases, but chose not to pursue their study, given the length and considerable complexity of the argument presented here.

Although we show that there exist non-egalitarian graphs with equilibria at which learning fails, we are far from characterizing those graphs. For example: is there a simple geometric characterization of the infinite graphs on which the agents learn with probability one?

References

- [1] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. Preprint, 2008.

- [2] D. Aldous and J. Steele. The objective method: Probabilistic combinatorial optimization and local weak convergence. *Probability on Discrete Structures (Volume 110 of Encyclopaedia of Mathematical Sciences)*, ed. H. Kesten, 110:1–72, 2003.
- [3] I. Arieli and M. Mueller-Frank. Inferring beliefs from actions. *Available at SSRN*, 2013.
- [4] R. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [5] V. Bala and S. Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, July 1998.
- [6] A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- [7] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Selected Works of Oded Schramm*, pages 533–545, 2011.
- [8] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, pages 992–1026, 1992.
- [9] P. Billingsley. *Convergence of probability measures*. John Wiley & Sons, Inc., New York, 1999.
- [10] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [11] P. DeMarzo, D. Vayanos, and J. Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118:909–968, 2003.
- [12] G. Ellison and D. Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 110(1):93–126, 1995.
- [13] D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, November 2003.
- [14] J. Glazer and A. Rubinstein. Motives and implementation: On the design of mechanisms to elicit opinions. *Journal of Economic Theory*, 79(2):157–173, 1998.
- [15] I. Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. In *Proc. Am. Math. Soc.*, volume 3, pages 170–174, 1952.
- [16] B. Golub and M. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- [17] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi. Information heterogeneity and the speed of learning in social networks. *Columbia Business School Research Paper*, (13-28), 2013.
- [18] J. Laslier and J. Weibull. Committee decisions: optimality and equilibrium. *Working Paper Series in Economics and Finance*, 2008.

- [19] I. Lobel and E. Sadler. Social learning and network uncertainty. Working paper, 2012.
- [20] R. McKelvey and T. Page. Common knowledge, consensus, and aggregate information. *Econometrica: Journal of the Econometric Society*, pages 109–127, 1986.
- [21] L. Ménager. Consensus, communication and knowledge: an extension with bayesian agents. *Mathematical Social Sciences*, 51(3):274–279, 2006.
- [22] P. Milgrom and R. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, pages 619–632, 1985.
- [23] E. Mossel, A. Sly, and O. Tamuz. Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, pages 1–31, 2013.
- [24] M. Mueller-Frank. A general framework for rational learning in social networks. Working paper, 2010.
- [25] R. Parikh and P. Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52(1):178–189, 1990.
- [26] D. Rosenberg, E. Solan, and N. Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009.
- [27] J. Sebenius and J. Geanakoplos. Don’t bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.
- [28] L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

A Rooted graphs

Claim A.1. *If $D([G, i], [G', i']) = 0$ then $(G, i) \cong (G', i')$.*

Proof. By the definition of $D(\cdot, \cdot)$, $D([G, i], [G', i']) = 0$ implies that $B_r(G, i) \cong B_r(G', i')$ for all $r \in \mathbb{N}$. Hence for each r there exists a (finite) graph isomorphism h_r from the vertices of $B_r(G, i)$ to the vertices of $B_r(G', i')$. The goal is to construct the (potentially infinite) graph isomorphism between (G, i) and (G', i') .

For each $r \in \mathbb{N}$, the isomorphism h_r can be restricted to an isomorphism between $B_1(G, i)$ and $B_1(G', i')$. Since $B_1(G, i)$ is finite, there are only finitely many possible isomorphisms between it and $B_1(G', i')$, and therefore at least one of them must appear infinitely often in $\{h_r\}_{r \geq 1}$. Hence let $\{h_{r,1}\}_{r \geq 1}$ be an infinite subsequence of $\{h_r\}_{r \geq 1}$ that consists of isomorphisms that are identical, when restricted to balls of radius one. By the same argument, there exists a sub-subsequence $\{h_{r,2}\}_{r \geq 2}$ that agrees on balls of radius two. Indeed, for any $n \in \mathbb{N}$ let $\{h_{r,n}\}_{r \geq n}$ be a subsequence of $\{h_{r,n-1}\}_{r \geq n-1}$ that agrees on balls of radius n .

In the diagonal sequence $\{h_{r,r}\}_{r \geq 1}$, $h_{r,r}$ agrees on balls of radius r with all $h_{s,s}$ such that $s \geq r$. We can therefore define an isomorphism $h : (G, i) \rightarrow (G', i')$ by specifying that $h(j) = h_{r,r}(j)$ for all $r \geq \Delta(i, j) + 1$. h is indeed an isomorphism, since if $(j, k) \in E$ then $(h(j), h(k)) = (h_r(j), h_r(k)) \in E'$, where $r = \max\{\Delta(i, j), \Delta(i, k)\}$. \square

Let $\mathcal{E}(d, L) \subset \mathcal{SCG}$ be the subspace of isomorphism classes of (d, L) -egalitarian strongly connected rooted graphs.

Theorem A.2. *$\mathcal{E}(d, L)$ is compact.*

Proof. Let $\{[G_n, i_n]\}_{n=1}^\infty$ be a sequence in $\mathcal{E}(d, L)$. Since the degrees are bounded, it follows that for fixed r , the number of possible balls $B_r(G_n, i_n)$ is finite, and therefore, by a standard diagonalization argument, there exists a subsequence that converges to some $[G, i]$. It remains to show that any such $[G, i]$ is L -locally-connected. This follows from the fact that for any edge (k, j) in (G, i) , $B_L(G, j) \cong B_L(G_n, j_n)$ for some n ; this ball must then include a path from k back to j of length at most L . \square

Rather than prove Lemma 6.1 directly, we prove the following more general theorem, which might be of independent interest, as it extends the well understood notion of compactness in undirected graphs to directed, strongly connected graphs. Lemma 6.1 is an immediate consequence.

Theorem A.3. *Let $\mathcal{S} \subseteq \mathcal{SCG}$ have the property that if $[G, i]$ is in \mathcal{S} , and if j is another vertex in G , then $[G, j]$ is also in \mathcal{S} . Then \mathcal{S} is precompact in \mathcal{SCG} if and only if $\mathcal{S} \subseteq \mathcal{E}(d, L)$ for some d and L .*

Proof. By Theorem A.2 $\mathcal{E}(d, L)$ is compact. Hence $\mathcal{S} \subseteq \mathcal{E}(d, L)$ implies that \mathcal{S} is precompact.

To prove the other direction, consider first a sequence $\{[G_n, i_n]\}$ in \mathcal{S} such that the degree of i_n is at least n . Then clearly $\{[G_n, i_n]\}$ has no converging subsequence, since the degree of i' in any limit $[G', i']$ would have to be larger than any n . It follows that \mathcal{S} is not precompact.

Finally, let there exist a sequence $\{[G_n, i_n]\}$ in \mathcal{S} with a sequence of edges $\{(i_n, j_n)\}$ in G_n where the shortest path from j_n back to i_n is at least of length n . Assume that $[G', i']$ is a limit of a subsequence of this sequence. It follows that there exists a $j' \in N(i')$ such

that the shortest path from j' back to i' is of length larger than any n , and so doesn't exist. Hence G' is not strongly connected, $[G', i'] \notin \mathcal{SCG}$, and \mathcal{S} is not compact in \mathcal{SCG} . \square

The following is a general claim that will be useful later.

Claim A.4. *Let $\{[G_n, i_n]\}_{n=1}^\infty$ be a sequence of rooted graph isomorphism classes such that*

$$\lim_{n \rightarrow \infty} [G_n, i_n] = [G, i].$$

Then for every $r > 0$ there exists an $N > 0$ such that for all $n > N$ it holds that $B_r(G_n, i_n) \cong B_r(G, i)$. Furthermore, there exists a subsequence $\{[G_{n_r}, i_{n_r}]\}_{r=1}^\infty$ such that $B_r(G_{n_r}, i_{n_r}) \cong B_r(G, i)$.

Proof. The first part of the claim follows directly from the definition of limits and Eq. (1). The second part holds for $n_r = \min\{n : B_r(G_n, i_n) \cong B_r(G, i)\}$, which is guaranteed to be finite by the first part. \square

B Locality

An important observation is that the actions and the utility of an agent, up to *time* t , depends only on the strategies of the agents that are at *distance* at most t from it. We formalize this notion in this section.

Claim B.1. *Let $\mathcal{G}_1 = (\mu_0, \mu_1, \lambda, G_1)$ and $\mathcal{G}_2 = (\mu_0, \mu_1, \lambda, G_2)$ be games. Let h be a rooted graph isomorphism between $B_{r+1}(G_1, i_1)$ and $B_{r+1}(G_2, i_2)$ for some $r > 0$, and let $Q_1^{j_1} = Q_2^{j_2}$ for all $j_1 \in B_r(G_1, i_1)$ and $j_2 = h(j_1)$.*

Then the games, as probability spaces, can be coupled so that $A_t^{j_1} = A_t^{j_2}$ for all $t \leq r$.

Some care needs to be taken with a statement such as “agent j_1 plays the same strategy as agent j_2 ”; it can only be meaningful in the context of a bijection that identifies each neighbor of j_1 with each neighbor of j_2 . We here naturally take this bijection to be h , and accordingly demand that it be an isomorphism between balls of radius $r + 1$ (rather than r), so that the neighbors of the agents on the surface of the ball are also mapped.

Proof. Couple the two processes by equating the states of nature and setting $W_{j_1} = W_{h(j_1)}$ for all $j_1 \in B_r(G_1, i_1)$, and furthermore coupling the choices of pure strategies of j_1 and $h(j_1)$.

We shall prove by induction a stronger statement, namely that under the claim hypothesis, $A_t^{j_1} = A_t^{j_2}$ for any $j_1 \in B_r(G_1, i_1)$, $j_2 = h(j_1)$ and $t \leq r - \Delta(i_1, j_1)$.

We prove the statement by induction on t . For $t = 0$, $A_0^{j_1}$ depends only on agent j_1 's private signal and choice of pure strategy, which are both equal to those of j_2 . Hence $A_0^{j_1} = A_0^{j_2}$ for all $j \in B_r(G_1, i_1)$.

Assume now that the claim holds up to some $t - 1 \leq r - 1$. Let j_1 be such that $t \leq r - \Delta(i_1, j_1)$. We would like to show that $A_t^{j_1} = A_t^{j_2}$. Let k_1 be a neighbor of j_1 . Then $t - 1 \leq r - \Delta(i_1, k_1)$, and so $A_{t'}^{k_1} = A_{t'}^{k_2}$ for all $t' \leq t - 1$, by the inductive assumption. Since $A_t^{j_1}$ depends only on j_1 's private signals, choice of pure strategy and the actions of her neighbors in previous time periods, and since these are all identical to those of j_2 , then it indeed follows that $A_t^{j_1} = A_t^{j_2}$. \square

Recalling the definition

$$u_{i,t} = \mathbb{P} [A_t^i = S],$$

the following corollary is a direct consequence of this claim.

Corollary B.2. *Let $\mathcal{G}_1 = (\mu_0, \mu_1, \lambda, G_1)$ and $\mathcal{G}_2(\mu_0, \mu_1, \lambda, G_2)$ be games. Let h be a rooted graph isomorphism between $B_{r+1}(G_1, i_1)$ and $B_{r+1}(G_2, i_2)$ for some $r > 0$, and let $Q_1^{j_1} = Q_2^{j_2}$ for all $j_1 \in B_r(G_1, i_1)$ and $j_2 = h(j_1)$.*

Then $u_{i_1,t} = u_{i_2,t}$ for all $t \leq r$.

C A topology on strategies and the existence of equilibria for finite graphs

In the following theorem we show that the agents' set of strategies admits a compact topology which preserves the continuity of the utilities. We use this topology to define our topology on equilibria, which is a crucial component of the proof of our main theorem. We also use it to infer the existence of equilibria for this game when the number of players is infinite.

For a fixed private belief I_i , a pure strategy is a function from the actions of neighbors to actions, which we call a response. Formally, let $G = (V, E)$ be a social network. A *response at time t* of an agent $i \in V$ is a function $r_{i,t} : \{0, 1\}^{|N(i)|t} \rightarrow \{0, 1\}$. A *response* of an agent i is the sequence of functions $r_i = (r_{i,0}, r_{i,1}, \dots)$. Let \mathcal{R}_i be the space of responses of agent i .

A (mixed) strategy of agent i can be thought of as a measure on the product space $[0, 1] \times \mathcal{R}_i$ of private beliefs and responses, with the marginal on the first coordinate equaling the distribution of I_i . Milgrom and Weber [22] call this representation a *distributional strategy*. In the proof of their Theorem 1, they show that for a game with incomplete information and a finite number of players, and given some conditions, the weak topology on distributional strategies is compact and keeps the utilities continuous. Then, using Glicksberg's theorem [15] they infer that the game has an equilibrium. The next theorem shows that these conditions apply in our case, when the number of agents is finite.

Lemma C.1. *Fix $G = (V, E)$, with V finite. Then for each agent i there exists a topology \mathcal{T}_i on her strategy space such that the strategy space is compact and the utilities u_j are continuous in the product of the strategy spaces. Furthermore, there exists an equilibrium strategy profile.*

Proof. We prove by showing that the conditions of Theorem 1 in [22] are met.

1. The set of private beliefs (*types* in the language of [22]) is $[0, 1]$, a complete separable metric space, as required. Furthermore, the distribution of private beliefs is absolutely continuous with respect to the product of their marginal distributions. This fulfills condition R2 of [22].
2. The utilities u_j are bounded, measurable functions of the private beliefs and the responses.

3. Define a metric D on i 's responses \mathcal{R}_i by

$$D(r_i, r'_i) = \exp(-\min\{t : r_{i,t} \neq r'_{i,t}\}).$$

This can be easily verified to indeed be a metric. By a standard diagonalization argument it follows that \mathcal{R}_i is compact in the topology induced by this metric, as required.

Furthermore, for fixed private beliefs, the utilities u_j are equicontinuous in the responses: if a response is changed by at most $\delta = e^{-T}$ (in terms of the metric D) then it remains unchanged in the first T time periods, and so the utilities are changed by at most $\epsilon = (1 - \lambda) \sum_{t=T}^{\infty} \lambda^t = \lambda^T$. This fulfills condition R1 of [22].

Since these conditions are met, it follows by the proof of Theorem 1 in [22] that the mixed strategies of agent i are compact in the weak topology \mathcal{T}_i , and that the utilities u_j are, under \mathcal{T}_i , a continuous function of the strategies. Furthermore, and again by Milgrom and Weber's Theorem 1, this game also has an equilibrium. \square

Note that under the above defined topology on \mathcal{R}_i the set of pure strategies is separable, and so the topology \mathcal{T}_i on (mixed) strategies is metrizable, e.g. with the Lévy-Prokhorov metric [9].

The following variant of Lemma C.1 will be useful later.

Lemma C.2. *Fix $G = (V, E)$, with V finite. Then for each agent i there exists a topology \mathcal{T}_i on her strategy space such that the strategy space is compact and the utilities in each time t , $u_{j,t}$, are continuous in the product of the strategy spaces.*

Proof. The proof is identical to the proof of Lemma C.1 above, except that we let each agent's expected utility be given by $u'_j = u_{j,t}$; that is, we set the discount factor to be one at time t and 0 otherwise. Since in the proof above we required of the discount factors nothing more than to have a finite sum, the proof still applies, and the utilities (in this case $u_{j,t}$), are continuous in the product of the strategy spaces. \square

D A topology on rooted graph strategy profiles

Let the *utility map at time t* $u_t : \mathcal{GS} \rightarrow \mathbb{R}$ be given by

$$u_t([G, i, \bar{Q}]) = u_{i,t}(G, \bar{Q}).$$

Lemma D.1. *The utility map $u : \mathcal{GS} \rightarrow \mathbb{R}$ is continuous.*

Proof. We will prove the claim by showing that u_t is continuous. The claim will follow because, by the bounded convergence theorem, if f is a linear combination of the uniformly bounded maps $\{f_t\}_{t=0}^{\infty}$, with summable positive coefficients, then the continuity of all the maps f_t implies the continuity of f .

Let $[G_n, i_n, \bar{Q}_n] \rightarrow_n [G, i, \bar{Q}]$. We will show that $u_t([G_n, i_n, \bar{Q}_n]) \rightarrow_n u_t(G, i, \bar{Q})$.

Consider a sequence of games \mathcal{G}_n which are all played on the finite graph $B = B_{t+1}(G, i)$. Since $[G_n, i_n, \bar{Q}_n] \rightarrow_n [G, i, \bar{Q}]$ then there exists some N such that, for $n > N$, it holds

that $D([G_n, i_n, \bar{Q}_n], [G, i, \bar{Q}]) < 2^{-(t+1)}$. Hence, by the definition of $D(\cdot, \cdot)$, it holds that $B \cong B_{t+1}(G_n, i_n)$ for $n > N$. Denote by h_n an isomorphism between the two balls that minimizes $\max_{j \in B_t(G, i)} d_{h_n}(Q^j, Q_n^{h_n(j)})$, as appears in the definition of $D(\cdot, \cdot)$.

Let each agent j in $B_t(G, i)$ play $Q_n^{h_n(j)}$ in \mathcal{G}_n , and let the rest of the agents in B (i.e., those at distance $t + 1$ from i) play arbitrary strategies. Denote by \bar{R}_n the strategy profile played by the agents at game \mathcal{G}_n , and denote by \bar{R} the restriction of \bar{Q} to B . By Corollary B.2, $u_t([G_n, i_n, \bar{Q}_n]) = u_t([B, i, \bar{R}_n])$ and $u_t([G, i, \bar{Q}]) = u_t([B, i, \bar{R}])$. Therefore it is left to show that $u_t([B, i, \bar{R}_n]) \rightarrow_n u_t([B, i, \bar{R}])$.

Now, $D([G_n, i_n, \bar{Q}_n], [G, i, \bar{Q}]) \rightarrow 0$. Hence $d(R^j, R_n^j) \rightarrow 0$, and so the strategies of each agent j in $B_t(G, i)$ converge to the strategy $R^j = Q^j$. Furthermore, the strategies in $B_t(G, i)$ converge uniformly, since there is only a finite number of them, and so the strategy profile converges in the product topology. It follows that $u_{i,t}$, which by Lemma C.2 is a continuous function of the strategies in $B_t(G, i)$, converges to $u_{i,t}$. \square

Lemma D.2. *The set of equilibrium rooted graph strategy profiles is closed.*

Proof. Let $\lim_n [G_n, i_n, \bar{Q}_n] = [G, i, \bar{Q}]$, and let each \bar{Q}_n be an equilibrium.

Let \bar{R} be a strategy profile for the agents in G such that $\bar{R}^j = \bar{Q}^j$ for all $j \neq i$. We will show that $u_i(G, \bar{Q}) \geq u_i(G, \bar{R})$.

Let \bar{R}_n be the strategy profile for agents on G_n defined by $\bar{R}_n^j = \bar{Q}_n^{j_n}$ for $j_n \neq i_n$, and let $\bar{R}_n^{i_n} = \bar{R}^i$. Note that $[G_n, i_n, \bar{R}_n] \rightarrow_n [G, i, \bar{R}]$.

Since \bar{Q}_n is an equilibrium profile of \mathcal{G}_n ,

$$u_{i_n}(G, \bar{Q}_n) \geq u_{i_n}(G, \bar{R}_n).$$

Taking the limit of both sides and substituting the definition of the utility map we get that

$$\lim_{n \rightarrow \infty} u([G_n, i_n, \bar{Q}_n]) \geq \lim_{n \rightarrow \infty} u(G_n, i_n, \bar{R}_n).$$

Finally, since by Lemma D.1 above the utility map is continuous, we have that

$$u([G, i, \bar{Q}]) \geq u([G, i, \bar{R}]).$$

\square

Claim D.3. *Let $R \subseteq \mathcal{SCG}$ be a compact set of rooted graphs, and let $\mathcal{SP}(R)$ be the set of rooted graph strategy profiles $[G, i, \bar{Q}]$ with $[G, i] \in R$. Then $\mathcal{SP}(R)$ is compact.*

Proof. Let $\{[G_n, i_n, \bar{Q}_n]\}_{n=1}^\infty$ be a sequence of rooted graph strategy profiles in $\mathcal{SP}(R)$. We will prove the claim by showing that it has a converging subsequence.

Let $[G, i]$ be the limit of some subsequence of $\{[G_n, i_n]\}_{n=1}^\infty$; this exists because R is compact.

By Claim A.4 there exists a sub-subsequence $\{[G_{n_r}, i_{n_r}]\}_{r=1}^\infty$ such that $B_r(G_{n_r}, i_{n_r}) \cong B_r(G, i)$. We will therefore assume without loss of generality that $n_r = n$, i.e., limit ourselves to this sub-subsequence. Accordingly, let $h_n : V \rightarrow V_n$ be a sequence rooted graph isomorphisms between $B_n(G, i)$ and $B_n(G_n, i_n)$. Note that since out-degrees are finite then $B_n(G, i)$ is finite for all n .

Let j be a vertex of G , and let r_j be the graph distance between i and j . For $n \geq r_j + 1$, denote $j_n = h_n(j)$. Note that h_n also maps the neighbors of j_n to the neighbors of j .

We will now construct \bar{Q} , the strategy profile of the agents in G such that $[G_{n_k}, i_{n_k}, \bar{Q}_{n_k}] \rightarrow_k [G, i, \bar{Q}]$, for some subsequence $\{n_k\}$. We start with agent 1 of G . Since \mathcal{T}_1 is compact, the sequence $\{Q_n^1\}_{n=r_1+1}^\infty$ has a converging subsequence, i.e., one along which $d_{h_n}(Q_n^1, Q^1) \rightarrow_n 0$ for some strategy Q^1 , which we will assign to agent 1 in G . Likewise, this subsequence has a subsequence along which $d_{h_n}(Q_n^2, Q^2) \rightarrow_n 0$, etc. Thus, by a standard diagonalization argument, we have that there exists a subsequence $\{[G_{n_k}, i_{n_k}, \bar{Q}_{n_k}]\}_{k=1}^\infty$ with isomorphisms h_{n_k} such that $d_{h_{n_k}}(Q_{n_k}^{j_{n_k}}, Q^j) \rightarrow_k 0$ for all j . It is now straightforward to verify that $D([G_{n_k}, i_{n_k}, \bar{Q}_{n_k}], [G, i, \bar{Q}]) \rightarrow_k 0$: pick some $r > 0$ and then k large enough so that h_{n_k} is an isomorphism between $B_{r+1}(G_{n_k}, i_{n_k})$ and $B_{r+1}(G, i)$. Then by definition (Eq. (2))

$$D([G_{n_k}, i_{n_k}, \bar{Q}_{n_k}], [G, i, \bar{Q}]) \leq \max \left\{ 2^{-r}, \max_{j \in B_r(G, i)} d_{h_{n_k}}(Q_{n_k}^{j_{n_k}}, Q^j) \right\}.$$

If we now further increase k then $d_{h_{n_k}}(Q_{n_k}^{j_{n_k}}, Q^j) \rightarrow_k 0$, and since $B_r(G, i)$ is finite then we have that $D([G_{n_k}, i_{n_k}, \bar{Q}_{n_k}], [G, i, \bar{Q}]) \leq 2^{-r}$, for k large enough. Since this holds for all r then

$$D([G_{n_k}, i_{n_k}, \bar{Q}_{n_k}], [G, i, \bar{Q}]) \rightarrow_k 0$$

and

$$\lim_{k \rightarrow \infty} [G_{n_k}, i_{n_k}, \bar{Q}_{n_k}] = [G, i, \bar{Q}].$$

□

An easy consequence of Claim D.3 and Lemma D.2 is the following claim. Let \mathcal{R} be a set of rooted graph isomorphism classes. Denote by $\mathcal{EQ}(\mathcal{R})$ the set of rooted graph strategy profiles $[G, i, \bar{Q}]$ such that $[G, i] \in \mathcal{R}$ and \bar{Q} is an equilibrium strategy profile. For a set of graphs \mathcal{K} , let $\mathcal{EQ}(\mathcal{K})$ be a shortened notation for $\mathcal{EQ}(\mathcal{R}(\mathcal{K}))$.

Claim D.4. *Let $\mathcal{R} \in \mathcal{SCG}$ be a precompact set of strongly connected rooted graphs. Then $\mathcal{EQ}(\mathcal{R})$ is a compact set of equilibrium rooted graph strategy profiles.*

Proof. We will prove the claim by showing that any sequence in $\mathcal{EQ}(\mathcal{R})$ has a converging subsequence with a limit that is an equilibrium.

Let $\{[G_n, i_n, \bar{Q}_n]\}_{n=1}^\infty$ be a sequence of points in $\mathcal{EQ}(\mathcal{R})$. Since \mathcal{R} is precompact in \mathcal{SCG} , the sequence $\{[G_n, i_n]\}_{n=1}^\infty$ has a converging subsequence $\{[G_{n_k}, i_{n_k}]\}_{k=1}^\infty$ that converges to some $[G, i] \in \mathcal{SCG}$. Hence, by Claim D.3, the sequence $\{[G_n, i_n, \bar{Q}_n]\}_{k=1}^\infty$ has a converging subsequence that, for some \bar{Q} , converges to $[G, i, \bar{Q}]$. Finally, by Lemma D.2 \bar{Q} is an equilibrium strategy profile for G , and so $[G, i, \bar{Q}]$ is an equilibrium rooted graph strategy. □

Using these claims, we are now ready to easily prove that every game has an equilibrium; the one additional ingredient will be Lemma C.1, which shows that finite games have equilibria.

Theorem D.5. *Every game \mathcal{G} has an equilibrium.*

Proof. Let $\mathcal{G} = (\mu_0, \mu_1, \lambda, G)$. Let i be a vertex in G , denote $G_n = B_n(G, i)$, and denote its root by i_n . Let $\{\mathcal{G}_n = (\mu_0, \mu_1, \lambda, G_n)\}_{n=1}^\infty$ be a sequence of finite games with equilibria strategy profiles \bar{Q}_n ; finite games have equilibria by Lemma C.1. Then $[G_n, i_n] \rightarrow_n [G, i]$, and so by Claim D.3 we have that there exists a strategy profile \bar{Q} and a subsequence $\{[G_{n_k}, i_{n_k}]\}_{k=1}^\infty$ such that

$$\lim_{k \rightarrow \infty} [G_{n_k}, i_{n_k}, \bar{Q}_{n_k}] = [G, i, \bar{Q}].$$

Finally, by Lemma D.2, \bar{Q} is an equilibrium profile of \mathcal{G} . □

E Agreement

Denote by Z_t^i the log-likelihood ratio of the events $S = 1$ and $S = 0$, conditioned on \mathcal{F}_t^i , the information available to agent i at time t

$$Z_t^i = \log \frac{\mathbb{P}[S = 1 | \mathcal{F}_t^i]}{\mathbb{P}[S = 0 | \mathcal{F}_t^i]},$$

and let

$$Z_\infty^i = \log \frac{\mathbb{P}[S = 1 | \mathcal{F}_\infty^i]}{\mathbb{P}[S = 0 | \mathcal{F}_\infty^i]}.$$

Let Y_t^i be defined as follows:

$$Y_t^i = \log \frac{\mathbb{P}\left[A_{[0,t]}^{N(i)} \mid A_{[0,t]}^i, Q^i, S = 1\right]}{\mathbb{P}\left[A_{[0,t]}^{N(i)} \mid A_{[0,t]}^i, Q^i, S = 0\right]},$$

where $A_{[0,t]}^i$ is the sequence of actions of i up to time $t - 1$, and $A_{[0,t]}^{N(i)}$ is the sequence of actions of i 's neighbors up to time $t - 1$. Finally, let

$$Y_\infty^i = \lim_{t \rightarrow \infty} Y_t^i.$$

Note that it is not clear that the limit $\lim_t Y_t^i$ exists. We show this in the following claim.

Claim E.1. *Denote by I_{-i} the private beliefs of all agents but i . Then*

1. $\lim_t Z_t^i = Z_\infty^i$ almost surely.
2. $Z_t^i = Y_t^i + Z_0^i$.
3. $\lim_t Y_t^i = Y_\infty^i$ almost surely, and Y_∞^i is measurable in $\sigma(A_{[0,\infty]}^i, I_{-i}, \bar{Q})$.

Proof. 1. Recall that

$$Z_t^i = \log \frac{\mathbb{P}[S = 1 | \mathcal{F}_t^i]}{\mathbb{P}[S = 0 | \mathcal{F}_t^i]}.$$

Since $\{\mathcal{F}_t^i\}_{t=0}^\infty$ is a filtration then $\mathbb{P}[S = 1 | \mathcal{F}_t^i]$ is a martingale, which converges a.s. since it is bounded. Hence Z_t^i also converges, and in particular

$$\lim_{t \rightarrow \infty} Z_t^i = \log \frac{\mathbb{P}[S = 1 | \mathcal{F}_\infty^i]}{\mathbb{P}[S = 0 | \mathcal{F}_\infty^i]} = Z_\infty^i.$$

2. By the definition of \mathcal{F}_t^i

$$\begin{aligned} Z_t^i &= \log \frac{\mathbb{P}[S = 1 | I_i, Q^i, A_{[0,t]}^{N(i)}]}{\mathbb{P}[S = 0 | I_i, Q^i, A_{[0,t]}^{N(i)}]} \\ &= \log \frac{\mathbb{P}[A_{[0,t]}^{N(i)} | I_i, Q^i, S = 1]}{\mathbb{P}[A_{[0,t]}^{N(i)} | I_i, Q^i, S = 0]} \frac{\mathbb{P}[S = 1 | I_i, Q^i]}{\mathbb{P}[S = 0 | I_i, Q^i]}, \end{aligned}$$

where the second equality follows from Bayes' law. Now, conditioned on S and i 's pure strategy Q^i , the probability for a sequence of actions $A_{[0,t]}^{N(i)}$ of i 's neighbors depends on I_i only in as much as I_i affects i 's actions up to time $t - 1$, $A_{[0,t]}^i$. Hence $\mathbb{P}[A_{[0,t]}^{N(i)} | I_i, Q^i, S] = \mathbb{P}[A_{[0,t]}^{N(i)} | A_{[0,t]}^i, Q^i, S]$. Note also that

$$\frac{\mathbb{P}[S = 1 | I_i, Q^i]}{\mathbb{P}[S = 0 | I_i, Q^i]} = Z_0^i.$$

Therefore

$$Z_t^i = Y_t^i + Z_0^i.$$

3. Since Z_t^i converges almost surely and $Z_t^i = Y_t^i + Z_0^i$ then Y_t^i also converges almost surely. Since each Y_t^i is a function of $A_{[0,t]}^{N(i)}$ and Q^i , it follows that their limit, Y_∞^i , is measurable in $\sigma(A_{[0,\infty)}^{N(i)}, Q^i)$. However, given \bar{Q} , $A_{[0,\infty)}^{N(i)}$ is a function of I_{-i} and $A_{[0,\infty)}^i$: for a choice of pure strategies the actions of all agents but i can be determined given their private signals and the actions of i . Hence Y_∞^i is also measurable in $\sigma(A_{[0,\infty)}^i, I_{-i}, \bar{Q})$. \square

Claim E.2. *The distribution of Z_0^i is non-atomic, as is the distribution of Z_0^i conditioned on S .*

Proof. By definition,

$$Z_0^i = \log \frac{\mathbb{P}[S = 1 | I_i, Q^i]}{\mathbb{P}[S = 0 | I_i, Q^i]}.$$

However, the choice of strategy Q^i is independent of both I_i and S , and so

$$Z_0^i = \log \frac{\mathbb{P}[S = 1|I_i]}{\mathbb{P}[S = 0|I_i]} = \log \frac{I_i}{1 - I_i}.$$

Since the distribution of I_i is non-atomic (see the definition of I_i and the comment after it) then so is the distribution of Z_0^i . Since S takes only two values then the same holds when conditioned on S . \square

Theorem E.3. *For any agent i it holds that $B_\infty^i = C_i$ almost surely, in any equilibrium.*

Proof. By its definition, C_i takes values in $\{0, 1, \{0, 1\}\}$, and by Theorem 5.2 we have that $C_i \subseteq B_\infty^i$. Therefore the claim holds when $B_\infty^i = 0$ or $B_\infty^i = 1$, and it remains to show that $C_i = \{0, 1\}$ when $B_\infty^i = \{0, 1\}$, or that $\mathbb{P}[C_i \neq \{0, 1\}, B_\infty^i = \{0, 1\}] = 0$.

Let $a = (a_0, a_1, \dots)$ be a sequence of actions, each in $\{0, 1\}$, in which only one action appears infinitely often. Since there are only countably many such sequences, then if $\mathbb{P}[C_i \neq \{0, 1\}, B_\infty^i = \{0, 1\}] > 0$, then there exists such a sequence a for which $\mathbb{P}[A_{[0, \infty)}^i = a, B_\infty^i = \{0, 1\}] >$

0. We shall prove the claim by showing that $\mathbb{P}[A_{[0, \infty)}^i = a, B_\infty^i = \{0, 1\}] = 0$.

Recall that by Claim E.1, the event that $B_\infty^i = \{0, 1\}$ is equal to the event that $Z_0^i = -Y_\infty^i$. Recall also that by the same claim, Y_∞^i is measurable in $\sigma(A_{[0, \infty)}^i, I_{-i}, \bar{Q})$. Hence

$$\begin{aligned} \mathbb{P}[A_{[0, \infty)}^i = a, B_\infty^i = \{0, 1\} | S, I_{-i}, \bar{Q}] &= \mathbb{P}[A_{[0, \infty)}^i = a, Z_0^i = -Y_\infty^i(a, I_{-i}, \bar{Q}) | S, I_{-i}, \bar{Q}] \\ &\leq \mathbb{P}[Z_0^i = -Y_\infty^i(a, I_{-i}, \bar{Q}) | S, I_{-i}, \bar{Q}] \end{aligned}$$

Now, by Claim E.2, Z_0^i conditioned on S has a non-atomic distribution. Further conditioning on \bar{Q} and I_{-i} leaves its distribution unchanged, since it is independent of the former, and independent of the latter conditioned on S . Hence the probability that it equals $-Y_\infty^i(a, S, I_{-i}, \bar{Q})$ is 0. Hence

$$\mathbb{P}[A_{[0, \infty)}^i = a, Z_0^i] = \mathbb{E}[\mathbb{P}[A_{[0, \infty)}^i = a, Z_0^i = -Y_\infty^i | S, I_{-i}, \bar{Q}]] = 0.$$

\square

Proof of Theorem 5.1. Let i and j be agents. Since G is strongly connected, there exists a path from i to j . By Theorem 5.3 we have, by induction along this path, that $C_j \subseteq B_\infty^i$ almost surely. But $C_i = B_\infty^i$ by Theorem E.3 above, and so we have that $C_j \subseteq C_i$. However, there also exists a path from j back to i , and so $C_i \subseteq C_j$, and the two are equal. This holds for any pair of agents, and so it follows that there exists a random variable C such that $C_i = C$ for all i , almost surely. \square

F δ -independence

In this section we introduce a technical notion of δ -independent random variables and (p, δ) -good estimators⁷. These definitions will be useful in the proof of our main theorem. Informally, δ -independent random variables are almost independent. The random variables

⁷ These definitions are taken (almost) verbatim from [23].

(X_1, \dots, X_n) are (p, δ) -good estimators of S if each is equal to S with probability at least p , and if they are δ -independent, conditioned on both $S = 0$ and $S = 1$.

Let $d_{TV}(\cdot, \cdot)$ denote the total variation distance between two measures defined on the same measurable space, and let (X_1, X_2, \dots, X_k) be random variables. We refer to them as δ -independent if

$$d_{TV}(\mu_{(X_1, \dots, X_k)}, \mu_{X_1} \times \dots \times \mu_{X_k}) \leq \delta,$$

where $\mu_{(X_1, \dots, X_k)}$ is their joint distribution, and $\mu_{X_1} \times \dots \times \mu_{X_k}$ is the product of their marginal distributions. I.e., the joint distribution $\mu_{(X_1, \dots, X_k)}$ has total variation distance of at most δ from the product of the marginal distributions $\mu_{X_1} \times \dots \times \mu_{X_k}$. Likewise, (X_1, \dots, X_l) are δ -dependent if the total variation distance between these distributions is more than δ .

Let S be a binary random variable such that $\mathbb{P}[S = 1] = 1/2$. We say that the binary random variables (X_1, \dots, X_k) are (p, δ) -good estimators of S if they are δ -independent conditioned both on $S = 0$ and on $S = 1$, and if $\mathbb{P}[X_\ell = S] \geq p$, for $\ell = 1, \dots, k$.

The following standard concentration of measure lemma captures the idea that the aggregation of sufficiently many (p, δ) -good estimators gives an arbitrarily good estimate, for any $p > \frac{1}{2}$ and for δ small enough.

Claim F.1. *Let S be a binary random variable such that $\mathbb{P}[S = 1] = 1/2$, and let (X_1, \dots, X_k) be $(\frac{1}{2} + \epsilon, \delta)$ -good estimators of S . Then there exists a function $a : \{0, 1\}^k \rightarrow \{0, 1\}$ such that*

$$\mathbb{P}[a(X_1, \dots, X_k) = S] > 1 - e^{-2\epsilon^2 k} - \delta.$$

Proof. Let (Y_1, \dots, Y_k) be random variables such that the distribution of (S, Y_i) is equal to the distribution of (S, X_i) for all i , and let (Y_1, \dots, Y_k) be independent, conditioned on S . Then (X_1, \dots, X_k) can be coupled to (Y_1, \dots, Y_k) in such a way that they differ only with probability δ . Therefore, if we show that $\mathbb{P}[a(Y_1, \dots, Y_k) = S] > q + \delta$ for some a then it will follow that $\mathbb{P}[a(X_1, \dots, X_k) = S] > q$.

Denote $\hat{Y} = \frac{1}{k} \sum_{i=1}^k Y_i$, and denote $\alpha_0 = \mathbb{E}[\hat{Y} | S = 0]$ and $\alpha_1 = \mathbb{E}[\hat{Y} | S = 1]$. It follows that

$$\alpha_1 - \alpha_0 = \frac{1}{k} \sum_{i=1}^k (2\mathbb{P}[Y_i = S] - 1) > 2\epsilon.$$

By the Hoeffding bound

$$\mathbb{P}[\hat{Y} \leq \alpha_1 - \epsilon | S = 1] < e^{-2\epsilon^2 k}$$

and

$$\mathbb{P}[\hat{Y} \geq \alpha_0 + \epsilon | S = 0] < e^{-2\epsilon^2 k}.$$

Let $a(Y_1, \dots, Y_k) = \mathbf{1}_{\hat{Y} > \alpha_1 - \epsilon}$. Then by the above we have that $\mathbb{P}[a(Y_1, \dots, Y_k) \neq S] < e^{-2\epsilon^2 k}$, and so

$$\mathbb{P}[a(X_1, \dots, X_k) = S] > 1 - e^{-2\epsilon^2 k} - \delta.$$

□

G The probability of learning

In this section we start to explore the probability of learning, with the ultimate goal of proving that it equals 1 under the appropriate conditions.

Recall that the *probability of learning* map $p : \mathcal{GS} \rightarrow \mathbb{R}$ is given by

$$p([G, i, \bar{Q}]) = \lim_{t \rightarrow \infty} \mathbb{P} [A_t^i(G, \bar{Q}) = S].$$

Before showing that p is well defined (i.e., the limit exists), and proving that it is lower semi-continuous, we make the following additional definition. Let

$$\hat{S}_\infty = \hat{S}_\infty([G, i, \bar{Q}]) = \begin{cases} 0 & B_\infty^i(G, \bar{Q}) = 0 \\ 1 & B_\infty^i(G, \bar{Q}) = 1 \text{ or } B_\infty^i(G, \bar{Q}) = \{0, 1\} \end{cases}$$

be a maximum a posteriori (MAP) estimator of S given \mathcal{F}_∞^i , for some agent i in G .

Note that $\hat{S}_\infty([G, i, \bar{Q}])$ is independent of i , since, by Theorem 5.1, $B_\infty^i = C$ for all agents i, j in G . Note also that \hat{S}_∞ is indeed a MAP estimator of S given \mathcal{F}_∞^i , since by definition B_∞^i is the set of most probable estimates of S , given \mathcal{F}_∞^i .

Claim G.1.

$$p([G, i, \bar{Q}]) = \mathbb{P} [\hat{S}_\infty([G, i, \bar{Q}]) = S].$$

Proof. We first condition on the event⁸ that $C = \{0, 1\}$. Then

$$\lim_{t \rightarrow \infty} \mathbb{P} [A_t^i = S | C = \{0, 1\}] = \lim_{t \rightarrow \infty} \mathbb{P} [A_t^i = S | B_\infty^i = \{0, 1\}].$$

Since the event that $B_\infty^i = \{0, 1\}$ is identical to the event that $\mathbb{P}[S = 1 | \mathcal{F}_\infty^i] = \frac{1}{2}$, and since A_t^i is \mathcal{F}_∞^i -measurable for all t , then it follows that

$$\lim_{t \rightarrow \infty} \mathbb{P} [A_t^i = S | C = \{0, 1\}] = \frac{1}{2}.$$

and also that

$$\lim_{t \rightarrow \infty} \mathbb{P} [\hat{S}_\infty = S | C = \{0, 1\}] = \frac{1}{2}.$$

When $C = 0$ or $C = 1$ then $\lim_t A_{i,t} = \hat{S}_\infty$, and so

$$\lim_{t \rightarrow \infty} \mathbb{P} [A_t^i = S | C \neq \{0, 1\}] = \mathbb{P} [\hat{S}_\infty = S | C \neq \{0, 1\}].$$

Since we have equality when conditioning on both $C \neq \{0, 1\}$ and $C = \{0, 1\}$ then we also have unconditioned equality and

$$p([G, i, \bar{Q}]) = \lim_{t \rightarrow \infty} \mathbb{P} [A_t^i = S] = \mathbb{P} [\hat{S}_\infty = S].$$

□

⁸Note that this might be a zero probability event (which we are not able to prove or disprove), in which case the following argument is not needed.

It follows that p is well defined. Since $\hat{S}_\infty([G, i, \bar{Q}])$ is independent of i then the following is a direct consequence of Claim G.1.

Corollary G.2. $p([G, i, \bar{Q}]) = p([G, j, \bar{Q}])$ for all i, j .

Another consequence is that if $p([G, i, \bar{Q}]) = 1$ then $\hat{S}_\infty = S$ almost surely. Since we could have also defined \hat{S}_∞ to equal 0 when $C = \{0, 1\}$, it follows that also $C = S$ almost surely. Hence $p([G, i, \bar{Q}]) = 1$ if and only if the agents learn S :

Claim G.3. $p([G, i, \bar{Q}]) = 1$ if and only if $\lim_t A_t^j = S$ almost surely for all agents j in G .

We next show that, not surprisingly, if the private signals are informative then $p > 1/2$. Given μ_0 and μ_1 , denote $p^*(\mu_0, \mu_1) = \frac{1}{2} + \frac{1}{2}d_{TV}(\mu_0, \mu_1)$.

Claim G.4. Given μ_0 and μ_1

$$p([G, i, \bar{Q}]) \geq p^*(\mu_0, \mu_1) > \frac{1}{2}$$

for any G, i and equilibrium strategy profile \bar{Q} .

Proof. $p^*(\mu_0, \mu_1) > \frac{1}{2}$, since $\mu_0 \neq \mu_1$. Let $\hat{S}_{i,0}$ be the maximum a posteriori (MAP) estimator of S given i 's private signal, W_i . Then (see Claim 3.30 in [23])

$$\mathbb{P} \left[\hat{S}_{i,0} = S \right] = p^*(\mu_0, \mu_1).$$

Now, \hat{S}_∞ is a MAP estimator of S given \mathcal{F}_∞^i . Since \mathcal{F}_∞^i includes W_i then

$$\mathbb{P} \left[\hat{S}_\infty = S \right] \geq \mathbb{P} \left[\hat{S}_{i,0} = S \right] = p^*(\mu_0, \mu_1),$$

and the claim follows by Claim G.1. □

We end this section with our main claim regarding p .

Theorem G.5. p is lower semi-continuous, i.e., if $[G_n, i_n, \bar{Q}_n] \rightarrow_n [G, i, \bar{Q}]$ then

$$\liminf_n p([G_n, i_n, \bar{Q}_n]) \geq p([G, i, \bar{Q}]).$$

Proof. Recall that the expected utility of agent i at time t is given by the utility map at time t :

$$u_t([G, i, \bar{Q}]) = \mathbb{P} \left[A_t^i(G, \bar{Q}) = S \right].$$

Hence an alternative definition of p is that

$$p([G, i, \bar{Q}]) = \lim_{t \rightarrow \infty} u_t([G, i, \bar{Q}]). \tag{3}$$

Now, A_t^i is \mathcal{F}_∞^i -measurable. Hence, since \hat{S}_∞ is a MAP estimator of S given \mathcal{F}_∞^i , it follows that

$$\mathbb{P} \left[\hat{S}_\infty = S \right] \geq \mathbb{P} \left[A_t^i = S \right],$$

or that

$$p([G, i, \bar{Q}]) \geq u_t([G, i, \bar{Q}]).$$

This, combined with (3), yields

$$p([G, i, \bar{Q}]) = \sup_t u_t([G, i, \bar{Q}]),$$

and since u_t is continuous (see the proof of Lemma D.1), it follows that p is lower semi-continuous. \square

H Finding good estimators

The following lemma is the technical core of the proof of the main result of this article. Before stating it we would like to remind the reader that if \mathcal{K} is a set of graphs then $\mathcal{R}(\mathcal{K})$ is the set of rooted \mathcal{K} graphs, and $\mathcal{EQ}(\mathcal{R}(\mathcal{K}))$ is the set of $\mathcal{R}(\mathcal{K})$ equilibrium graph strategy profiles. Recall also that $p^*(\mu_0, \mu_1) = \frac{1}{2} + \frac{1}{2}d_{TV}(\mu_0, \mu_1)$.

Lemma H.1. *Let G be an infinite, strongly connected graph such that $\overline{\mathcal{EQ}(\mathcal{R}(\{G\}))}$ is a compact set of equilibrium rooted graph strategy profiles. Then for all equilibrium strategy profiles on G , $k \in \mathbb{N}$, $\epsilon > 0$ and $\delta > 0$ there exists an agent i in G , a time t and \mathcal{F}_t^i -measurable binary random variables (X_1, \dots, X_k) that are $(p^*(\mu_0, \mu_1) - \epsilon, \delta)$ -good estimators of S .*

Before proving this lemma we will need some additional definitions and claims.

We shall make use of the following notation: Let X_1, \dots, X_k be random variables, and let S be a binary random variable. We say that (X_1, \dots, X_k) are δ -independent conditioned on S if they are δ -independent conditioned on both $S = 0$ and $S = 1$. Denote

$$\text{dep}_S(X_1, \dots, X_k) = \min\{\delta : (X_1, \dots, X_k) \text{ are } \delta\text{-independent conditioned on } S\}$$

Note that this minimum is indeed attained, by the definition of δ -independence.

The proofs of the following two general claims are elementary and fairly straightforward. They appear in [23].

Claim H.2. *Let A, B and C be random variables such that $\mathbb{P}[A \neq B] \leq \delta$ and (B, C) are δ' -independent. Then (A, C) are $2\delta + \delta'$ -independent.*

Claim H.3. *Let $A = (A_1, \dots, A_k)$, and X be random variables. Let (A_1, \dots, A_k) be δ_1 -independent and let (A, X) be δ_2 -independent. Then (A_1, \dots, A_k, X) are $(\delta_1 + \delta_2)$ -independent.*

Claim H.4. *Let $[G, i_0, \bar{Q}]$ be an equilibrium graph strategy. Let $\{i_n\}_{n=1}^\infty$ be a sequence of vertices such that the graph distance $\Delta(i_0, i_n)$ diverges with n . Fix t , and for each n let X^{i_n} be $\mathcal{F}_t^{i_n}$ -measurable. Then*

$$\lim_{n \rightarrow \infty} \text{dep}_S(X^{i_n}, \hat{S}_\infty) = 0.$$

Proof. Let $\mathcal{B}_r^i = \sigma(\{W_j, Q^j : j \in B_r(G, i)\})$. We first show, by induction on r , that $\mathcal{F}_r^i \subseteq \mathcal{B}_r^i$: any \mathcal{F}_r^i -measurable random variable is also \mathcal{B}_r^i -measurable. It will follow that X^{i_n} is $\mathcal{B}_t^{i_n}$ -measurable.

Clearly $\mathcal{F}_0^i \subseteq \mathcal{B}_0^i$. Assume now that $\mathcal{F}_{r'}^j \subseteq \mathcal{B}_{r'}^j$ for all j and $r' < r$. By definition, $\mathcal{F}_r^i = \sigma(\mathcal{F}_{r-1}^i, A_{r-1}^{N(i)})$. For $j \in N(i)$ we have that A_{r-1}^j is \mathcal{B}_{r-1}^j -measurable. Finally, $\mathcal{B}_{r-1}^j \subseteq \mathcal{B}_r^i$, and so $\mathcal{F}_r^i \subseteq \mathcal{B}_r^i$.

Note that for i, j and r_1, r_2 such that $B_{r_1}(G, i)$ and $B_{r_2}(G, j)$ are disjoint it holds that $\mathcal{B}_{r_1}^i$ and $\mathcal{B}_{r_2}^j$ are independent conditioned on S , since the choices of pure strategies are independent and private beliefs are independent conditioned on S .

Let R_r^i be a MAP estimator of \hat{S}_∞ given \mathcal{B}_r^i . Since $\Delta(i_0, i_n) \rightarrow_n \infty$, it follows that for any r and n large enough $B_t(G, i_n)$ and $B_r(G, i_0)$ are disjoint, and so X^{i_n} and $R_r^{i_0}$ are independent, conditioned on S . For such n , by Claim H.2, we have that $(X^{i_n}, \hat{S}_\infty)$ are 2δ -independent, for $\delta = \mathbb{P}\left[R_r^i \neq \hat{S}_\infty\right]$.

Finally, since \hat{S}_∞ is \mathcal{B}_∞^i -measurable, it follows that

$$\lim_{r \rightarrow \infty} \mathbb{P}\left[R_r^i \neq \hat{S}_\infty\right] = 0,$$

and so

$$\lim_{n \rightarrow \infty} \text{dep}_S(X^{i_n}, \hat{S}_\infty) = 0.$$

□

We are now ready to prove Lemma H.1.

Proof of Lemma H.1. Denote by \mathcal{C} the closure of $\mathcal{EQ}(\mathcal{R}(\{G\}))$. Note that by Lemma D.2 any graph strategy in \mathcal{C} is an equilibrium.

We shall prove by induction a stronger claim, namely that under the claim hypothesis, for all $[H, j, \bar{Q}] \in \mathcal{C}$, $k \in \mathbb{N}$, $\epsilon > 0$ and $\delta > 0$ there exists an agent i in H , a time t and \mathcal{F}_t^i -measurable binary random variables (X_1, \dots, X_k) that are $(p^*(\mu_0, \mu_1) - \epsilon, \delta)$ -good estimators of S .

We prove the claim by induction on k . The claim holds trivially for $k = 0$. Assume that the claim holds up to k .

Let $[H, j, \bar{Q}] \in \mathcal{C}$. Let $\{j_n\}_{n=1}^\infty$ be a sequence of vertices in H such that $\lim_n \Delta(j, j_n) = \infty$. Since \mathcal{C} is compact then there exists a converging sequence $[H, j_n, \bar{Q}] \rightarrow_n [F, i', \bar{R}] \in \mathcal{C}$. By the inductive assumption, there exists an agent i in F , a time t and random variables (X_1^i, \dots, X_k^i) which are \mathcal{F}_t^i -measurable and are $(p^*(\mu_0, \mu_1) - \epsilon', \delta')$ -good estimators of S , for some $0 < \epsilon' < \epsilon$ and $0 < \delta' < \delta$. Denote

$$\bar{X}_k^i = (X_1^i, \dots, X_k^i).$$

Let $r = \Delta(i', i)$. Since for n large enough $B_r(H, j_n) \cong B_r(F, i')$ then, if we let $i_n \in B_r(H, j_n)$ be the vertex that corresponds to $i \in B_r(F, i')$ then $[H, i_n, \bar{R}] \rightarrow_n [F, i, \bar{R}]$, and $\lim_n \Delta(j, i_n) = \infty$.

Since \bar{X}_k^i is \mathcal{F}_t^i -measurable, there exists a function x_k^i such that $\bar{X}_k^i = x_k^i(I_i, A_{[0,t]}^{N(i)})$. Denote $\bar{X}_k^{i_n} = x_k^i(I_{i_n}, A_{[0,t]}^{N(i_n)})$.

Now, since the strategies of agents in the neighborhood of i_n in H converge in the weak topology to those of i in F , then the random variables $\{(S, \bar{X}_k^{i_n})\}_{n=1}^\infty$ converge in the weak topology to (S, \bar{X}_k^i) . Moreover, the measures of these random variables are over the finite space $\{0, 1\}^{k+1}$, and so we also have convergence in total variation. In particular, $(X_1^{i_n}, \dots, X_k^{i_n})$ approach δ' -independence:

$$\lim_{n \rightarrow \infty} \text{dep}_S(X_1^{i_n}, \dots, X_k^{i_n}) = \text{dep}_S(X_1^i, \dots, X_k^i) \leq \delta'. \quad (4)$$

Likewise,

$$\lim_{n \rightarrow \infty} \mathbb{P}[X_\ell^{i_n} = S] = \mathbb{P}[X_\ell^i = S] > p^*(\mu_0, \mu_1) - \epsilon'. \quad (5)$$

for $\ell = 1, \dots, k$. Additionally, since $\Delta(j, i_n) \rightarrow_n \infty$, it follows by Claim H.4 that

$$\lim_{n \rightarrow \infty} \text{dep}_S(\bar{X}_k^{i_n}, \hat{S}_\infty) = 0, \quad (6)$$

that is, $\bar{X}_k^{i_n}$ and \hat{S}_∞ are practically independent, for large n .

Now, recall that \hat{S}_∞ is \mathcal{F}_∞^i -measurable. Therefore, if we let $R_{t'}^{i_n}$ be a MAP estimator of \hat{S}_∞ given $\mathcal{F}_{t'}^i$ then for any n it holds that

$$\lim_{t' \rightarrow \infty} \mathbb{P}[R_{t'}^{i_n} = \hat{S}_\infty] = 1. \quad (7)$$

By Claim H.2, a consequence of Eqs. (6) and (7) is that

$$\lim_{n \rightarrow \infty} \lim_{t' \rightarrow \infty} \text{dep}_S(\bar{X}_k^{i_n}, R_{t'}^{i_n}) = 0.$$

That is, $\bar{X}_k^{i_n}$ and $R_{t'}^{i_n}$ are practically independent, for large enough n and t' . It follows by Claim H.3 that

$$\lim_{n \rightarrow \infty} \lim_{t' \rightarrow \infty} \text{dep}_S(X_1^{i_n}, \dots, X_k^{i_n}, R_{t'}^{i_n}) \leq \delta'. \quad (8)$$

It follows from Eq. (7) that

$$\lim_{t' \rightarrow \infty} \mathbb{P}[R_{t'}^{i_n} = S] = \mathbb{P}[\hat{S}_\infty = S] \geq p^*(\mu_0, \mu_1). \quad (9)$$

Gathering the above results, we have that for n and t' large enough,

1. $\mathbb{P}[X_\ell^{i_n} = S] \geq p^*(\mu_0, \mu_1) - \epsilon$, by Eq. (5). Likewise $\mathbb{P}[R_{t'}^{i_n} = S] = \mathbb{P}[\hat{S}_\infty = S] \geq p^*(\mu_0, \mu_1) - \epsilon$ by Eq. (9).
2. $(X_1^{i_n}, \dots, X_k^{i_n}, R_{t'}^{i_n})$ are δ -independent, by Eq. (8).

We therefore have that $(X_1^{i_n}, \dots, X_k^{i_n}, R_{t'}^{i_n})$ are $\mathcal{F}_{t'}^{i_n}$ -measurable $(p^*(\mu_0, \mu_1) - \epsilon, \delta)$ -good estimators of S . \square

I Proof of main theorem

The next theorem is a more general version of our main theorem, for infinite graphs.

Theorem I.1. *Let G be an infinite, strongly connected graph such that $\overline{\mathcal{EQ}(\mathcal{R}(\{G\}))}$ is a compact set of equilibrium rooted graph strategy profiles, and let \bar{Q} be any equilibrium strategy profile of \mathcal{G} . Then*

$$\mathbb{P} \left[\lim_t A_t^i = S \right] = 1$$

for all agents i .

It follows that $\mathbb{P} [\lim_t A_t^i = S] = 1$ whenever G is an infinite, strongly connected, (d, L) -egalitarian graph, by Theorem A.3 and Claim D.4.

Proof. We first show that $p = p([G, i, \bar{Q}]) = 1$. Assume by way of contradiction that $p < 1$. By Claim G.4, we have that $p > 1/2$.

By Lemma H.1, for every $k \in \mathbb{N}$, and $\delta > 0$ there exist (X_1, \dots, X_k) that are \mathcal{F}_t^i -measurable for some i and t , are δ -independent conditioned on S , and are each equal to S with probability bounded away from one half, since $p^*(\mu_0, \mu_1) > \frac{1}{2}$.

By Claim F.1 it follows that for k large enough and δ small enough, there exists an estimator \hat{S} of S that is a function of (X_1, \dots, X_k) , and is equal to S with probability strictly greater than p .

This \hat{S} is \mathcal{F}_∞^i -measurable, and so a MAP estimator of S given \mathcal{F}_∞^i must also equal S with probability greater than p . However, \hat{S}_∞ is a MAP estimator of S given \mathcal{F}_∞^i , and it equals S with probability p (Claim G.1), and so we have a contradiction. Hence $p([G, i, \bar{Q}]) = 1$.

Now, by Claim G.1 we have that

$$\mathbb{P} \left[\hat{S}_\infty([G, i, \bar{Q}]) = S \right] = p([G, i, \bar{Q}]) = 1.$$

By the definition of \hat{S}_∞ we have that $\hat{S}_\infty = S$ if and only if $B_\infty^i = S$ for some (equivalently all) i . Since, by Theorem 5.1, $C = C_i = B_\infty^i$, it follows that $\mathbb{P} [C = S] = 1$, and that therefore $\mathbb{P} [\lim_t A_t^i = S] = 1$. \square

Online Appendix: Examples

In this appendix we give two examples showing that the assumptions of bounded out-degree and L -connectedness are crucial. Our approach in constructing equilibria will be to prescribe the initial moves of the agents and then extend this to an equilibrium strategy profile.

Define the set of times and histories agents have to respond to as $\mathcal{H} = \{(i, t, a) : i \in V, t \in \mathbb{N}_0, a \in [0, 1] \times \{0, 1\}^{|N(i)| \cdot t}\}$. The set $[0, 1] \times \{0, 1\}^{|N(i)| \cdot t}$ is interpreted as the pair of the private belief of i and the history of actions observed by agent i up to time t . If $a \in [0, 1] \times \{0, 1\}^{|N(i)| \cdot t}$ then for $0 \leq t' \leq t$ we let $a_{t'} \in [0, 1] \times \{0, 1\}^{|N(i)| \cdot t'}$ denote the history restricted to times up to t' . We say that a subset $\mathcal{H} \subseteq \mathcal{H}$ is *history-closed* if for every $(i, t, a) \in \mathcal{H}$ we have that for all $0 \leq t' \leq t$ that $(i, t', a_{t'}) \in \mathcal{H}$.

For a strategy profile \bar{Q} denote the optimal expected utility for i under any response as $u_i^*(\bar{Q}) = \sup_{\bar{R}} u_i(\bar{R})$ where the supremum is over strategy profiles \bar{R} such that $R^j = Q^j$ for all $j \neq i$ in V .

Definition I.2. *On a history-closed subset $\mathcal{H} \in \mathcal{H}$ a forced response $q_{\mathcal{H}}$ is a map $q_{\mathcal{H}} : \mathcal{H} \rightarrow \{0, 1\}$ denoting a set of actions we force the agents to make. A strategy profile \bar{Q} is $q_{\mathcal{H}}$ -forced if for every $(i, t, a) \in \mathcal{H}$ if agent i at time t has seen history a from her neighbors then she selects action $q_{\mathcal{H}}(i, t, a)$. A strategy profile \bar{Q} is a $q_{\mathcal{H}}$ -equilibrium if it is $q_{\mathcal{H}}$ -forced and for every agent $i \in V$ it holds that $u_i(\bar{Q}) \geq u_i(\bar{R})$ for any $q_{\mathcal{H}}$ -forced strategy profile \bar{R} such that $R^j = Q^j$ for all $j \neq i$ in V .*

The following lemma can be proved by a minor modification of Theorem D.5 and so we omit the proof.

Lemma I.3. *Let $\mathcal{H} \in \mathcal{H}$ be history-closed and let $q_{\mathcal{H}}$ be a forced response. There exists a $q_{\mathcal{H}}$ -equilibrium.*

Having constructed $q_{\mathcal{H}}$ -equilibria we then will want to show that they are equilibria. In order to do that we appeal to the following lemma.

Lemma I.4. *Let \bar{Q} be a $q_{\mathcal{H}}$ -equilibrium. Suppose that for every agent i , any strategy profile \bar{R} that attains $u_i^*(\bar{Q})$ has that for all t ,*

$$\mathbb{P} \left[\bar{Q}_i^i(I_i, A_{[0,t]}^{N(i)}) \neq \bar{R}_i^i(I_i, A_{[0,t]}^{N(i)}), (i, t, (I_i, A_{[0,t]}^{N(i)})) \in \mathcal{H} \right] = 0. \quad (10)$$

Then \bar{Q} is an equilibrium.

Proof. If \bar{Q} is not an equilibrium then by compactness there exists a strategy profile for \bar{R} that attains u_i^* and differs from \bar{Q} only for agent i . By equation (10) this implies that agent i following \bar{R} must take the same actions almost surely as if they were following \bar{Q} until the end of the forced moves. Hence it is $q_{\mathcal{H}}$ -forced and so \bar{R} is a $q_{\mathcal{H}}$ -equilibrium. It follows that i cannot increase the expected utility of \bar{Q} , which is therefore an equilibrium. \square

In order to show that every agent follows the forced moves almost surely we now give a lemma which gives a sufficient condition for an agent to act myopically, according to her posterior distribution. For an equilibrium strategy profile \bar{Q} let $\bar{Q}_{i,t,a}^\dagger$ be the strategy profile

where the agents follow \bar{Q} except that if agent i has $a = (I_i, A_{[0,t]}^{N(i)})$ then from time t onwards agent i acts myopically, taking action $B_{t'}^i(G, \bar{Q}_{i,t,a}^\dagger)$ for time $t' \geq t$. We denote

$$Y_\ell = Y_\ell(i, t, a) := \mathbb{E} \left[\left| \mathbb{P} \left[S = 1 \mid \mathcal{F}_{t+\ell}^i(G, \bar{Q}_{i,t,a}^\dagger) \right] - 1/2 \right| \mid \mathcal{F}_t^i, a = (I_i, A_{[0,t]}^{N(i)}) \right].$$

We will show that the following are sufficient conditions for agent i to act myopically. For $\ell \in \{1, 2, 3\}$ we set $\mathcal{B}_\ell = \left\{ 2Y_0 > \frac{\lambda^2(\frac{1}{2} - Y_{\ell-1})}{1-\lambda} \right\}$ and we set

$$\mathcal{B}_4 = \left\{ 2Y_0 > \lambda^2\left(\frac{1}{2} - Y_2\right) + \frac{\lambda^3\left(\frac{1}{2} - Y_3\right)}{1-\lambda} \right\}.$$

Since \bar{Q} and $\bar{Q}_{i,t,a}^\dagger$ are the same up to time $t-1$ we have that $\mathcal{F}_t^i(G, \bar{Q})$ is equal to $\mathcal{F}_t^i(G, \bar{Q}_{i,t,a}^\dagger)$. As Y_ℓ is the expectation of a submartingale it is increasing. Hence, after rearranging we see that $\mathcal{B}_1 \subseteq \mathcal{B}_2 \subseteq \mathcal{B}_3 \subseteq \mathcal{B}_4$.

Lemma I.5. *Suppose that for strategy profile \bar{Q} agent i has an optimal response, such that for any \bar{R} such that $R^j = \bar{Q}^j$ for all $j \neq i$ in V then $u_i(\bar{Q}) \geq u_i(\bar{R})$. Then for any t ,*

$$\mathbb{P} \left[A_t^i(G, \bar{Q}) \neq B_t^i, \mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4 \right] = 0,$$

that is, agent i acts myopically at time t under \bar{Q} almost surely, on the event $\mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4$.

Proof. If agent i acts under $\bar{Q}_{i,t,a}^\dagger$ then her expected utility from time t onwards given a is

$$\begin{aligned} u_{i,t,a}(\bar{Q}_{i,t,a}^\dagger) &:= (1-\lambda) \sum_{t'=t}^{\infty} \lambda^{t'} \mathbb{E} \left[\mathbb{P} \left[A_{t'}^i(G, \bar{Q}_{i,t,a}^\dagger) = S \mid \mathcal{F}_{t'}^i, a = (I_i, A_{[0,t]}^{N(i)}) \right] \right] \\ &\geq (1-\lambda) \lambda^t \left(\frac{1}{2} + Y_0 + \lambda \left(\frac{1}{2} + Y_1 \right) + \lambda^2 \left(\frac{1}{2} + Y_2 \right) + \frac{\lambda^3}{1-\lambda} \left(\frac{1}{2} + Y_3 \right) \right) \end{aligned}$$

under $\bar{Q}_{i,t,a}^\dagger$. Now assume that the action of agent i at time t under \bar{Q} is not the myopic choice. Then her expected utility is at most

$$\begin{aligned} u_{i,t,a}(\bar{Q}) &\leq (1-\lambda) \lambda^t \left(\frac{1}{2} - \left| \mathbb{P} \left[S = 1 \mid \mathcal{F}_t^i, a = (I_i, A_{[0,t]}^{N(i)}) \right] - \frac{1}{2} \right| \right. \\ &\quad \left. + \lambda \mathbb{E} \left[\mathbb{P} \left[A_{t+1}^i(G, \bar{Q}) = S \mid \mathcal{F}_t^i, a = (I_i, A_{[0,t]}^{N(i)}) \right] + \frac{\lambda^2}{1-\lambda} \right] \right). \end{aligned}$$

We note that at time $t+1$ the information available about S is the same under both strategies since the only difference is the choice of action by agent i at time t , hence as i takes the optimal action under \bar{Q}^\dagger ,

$$\frac{1}{2} + Y_1 = \mathbb{E} \left[\mathbb{P} \left[A_{t+1}^i(G, \bar{Q}_{i,t,a}^\dagger) = S \mid \mathcal{F}_t^i, a = (I_i, A_{[0,t]}^{N(i)}) \right] \right] \geq \mathbb{E} \left[\mathbb{P} \left[A_{t+1}^i(G, \bar{Q}) = S \mid \mathcal{F}_t^i, a = (I_i, A_{[0,t]}^{N(i)}) \right] \right].$$

Since \bar{Q} is optimal for i we have that

$$0 \geq u_{i,t,a}(\bar{Q}_{i,t,a}^\dagger) - u_{i,t}(\bar{Q}) \geq (1-\lambda) \lambda^t \left(2Y_0 - \lambda^2 \left(\frac{1}{2} - Y_2 \right) - \frac{\lambda^3}{1-\lambda} \left(\frac{1}{2} - Y_3 \right) \right). \quad (11)$$

Condition (11) does not hold under \mathcal{B}_4 so $\mathbb{P} \left[A_t^i(G, \bar{Q}) \neq B_t^i, \mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4 \right] = 0$. \square

I.1 The royal family

In the main theorem we require that the graph G not only be strongly connected, but also L -connected and have bounded out-degrees, which are local conditions. In the following example the graph is strongly connected, has bounded out-degrees, but is not L -connected. We show that for bounded private beliefs asymptotic learning does not occur in all equilibria⁹.




Figure 4: The Royal Family. Each member of the public (on the left), observes each royal (on the right), as well as her next door neighbors. The royals observe each other, and one royal observes one member of the public.

Consider the following graph (Figure 4). The vertex set is comprised of two groups of agents: a “royal family” clique of R agents who all observe each other, and $n \in \mathbb{N} \cup \{\infty\}$ agents - the “public” - who are connected in an undirected chain, and in addition can all observe all the agents in the royal family. Finally, a single member of the royal family observes one of the public, so that the graph is strongly connected.

We choose that μ_0 and μ_1 so that $\mathbb{P}[Z_0^i \in (1, 2) \cup (-2, -1)] = 1$ and set the forced moves so that all agents act myopically at time 1. By Lemma I.3 we can extend this to a forced equilibrium \bar{Q} . By Lemma I.4 it is sufficient to show that no agent can achieve their optimum without choosing the myopic action in the first round. By our choice of μ_0 and μ_1 we have that

$$\left| \mathbb{P}[S = 1 | \mathcal{F}_0^i] - \frac{1}{2} \right| = \frac{e^{|Z_0^i|}}{1 + e^{|Z_0^i|}} - \frac{1}{2} \geq \frac{e}{1 + e} - \frac{1}{2} \geq \frac{1}{5}.$$

Hence in the notation of Lemma I.5 we have that $Y_0 \geq \frac{1}{5}$ when $t = 0$ for all i and a almost surely. Moreover, after the first round all agents see the royal family and can combine their information. Since the signals are bounded it follows that for some $c = c(\mu_0, \mu_1) > 0$, independent of R and n

$$\mathbb{E} \left[\frac{1}{2} - \left| \mathbb{P}[S = 1 | \mathcal{F}_1^i] - \frac{1}{2} \right| \middle| \mathcal{F}_0^i \right] \leq e^{-cR}.$$

Hence if R is a large constant \mathcal{B}_2 holds so by Lemma I.5 if an agent is to attain her maximal expected utility given the actions of the other agents, she must act myopically almost surely at time 0. Thus \bar{Q} is an equilibrium.

Let \mathcal{J} denote the event that all agents in the royal family have a signal favoring state 1. On this event under \bar{Q} all agents in the royal family choose action 1 at time 0 and this is

⁹We draw on Bala and Goyal’s [5] *royal family* graph.

observed by all the agents so $\mathcal{J} \in \mathcal{F}_1^i$ for all i . Since agents observe at most one other agent this signal overwhelms their other information and so

$$\mathbb{P}[S = 1 | \mathcal{F}_1^i, \mathcal{J}] \geq 1 - e^{-cR},$$

for all $i \in V$. Thus if R is a large constant \mathcal{B}_1 holds for all the agents at time 1 so by Lemma I.5 they all act myopically and choose action 1 at time 1. Since $\mathcal{J} \in \mathcal{F}_1^i$ they also all knew this was what would happen so gain no extra information. Iterating this argument we see that all agents choose 1 in all subsequent rounds. However, $\mathbb{P}[\mathcal{J}, S = 0] \geq e^{-c'R}$ where c' is independent of R and n . Hence as we let n tend to infinity the probability of learning does not tend to 1, and when n equals infinity the probability of learning does not equal 1.

I.2 The mad king

More surprising is that there exist *undirected* (i.e., 1-connected) graphs with equilibria where asymptotic learning fails; These graphs have unbounded out-degrees. Note that in the myopic case learning is achieved on these graphs [23], and so this is an example in which strategic behavior impedes learning.

In this example we consider a finite graph which includes 5 classes of agents. There is a king, labeled u , and a regent labeled v . The court consists of R_C agents and the bureaucracy of R_B agents. The remaining n are the people. Note again that the graph is undirected.

- The king is connected to the regent, the court and the people.
- The regent is connected to the king and to the bureaucracy.
- The members of the court are each connected only to the king.
- The members of the people are each connected only to the king.
- The members of the bureaucracy are each connected only to the regent.

See Figure 5.

As in the previous example we will describe some initial forced equilibrium and then appeal to existence results to extend it to an equilibrium. We suppose that μ_0 and μ_1 are such that $\mathbb{P}[Z_0^i \in (1, 1 + \epsilon) \cup (-\sqrt{7}, -\sqrt{7} + \epsilon)] = 1$ where ϵ is some very small positive constant, and will choose R_C, λ and R_B so that e^{R_C} is much smaller than $\frac{1}{1-\lambda}$ which in turn will be much smaller than R_B :

$$e^{R_C} \ll \frac{1}{1-\lambda} \ll R_B.$$

The equilibrium we describe will involve the people being forced to choose action 0 in rounds 0 and 1, as otherwise the king “punishes” them by withholding his information. As an incentive to comply he offers them the opinion of his court and, later, of his bureaucracy. While the opinion of the bureaucracy is correct with high probability, it is still bounded, and so, even as the size of the public tends to infinity, the probability of learning stays bounded away from one.

We now describe a series of forced moves for the agents, fixing $\delta > 0$ to be some small constant.

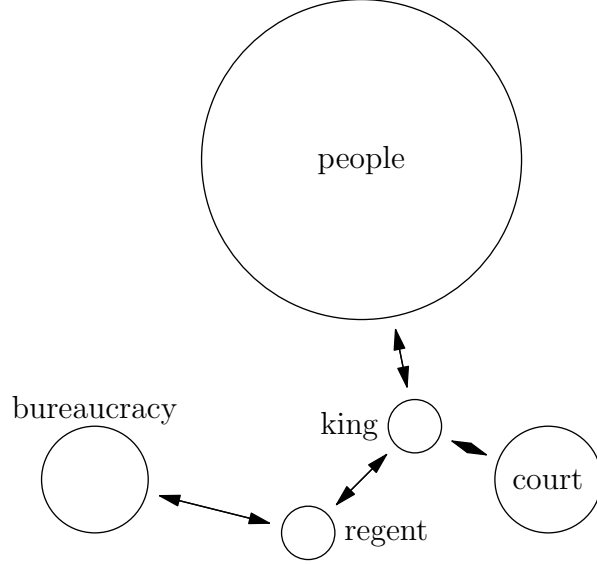


Figure 5: The mad king.

- The regent acts myopically at time 0. If for some state s $\mathbb{P}[S = s | \mathcal{F}_1^v] \geq 1 - e^{-\delta R_B}$ then the regent chooses states s in round 1 and all future rounds, otherwise his moves are not forced.
- The king acts myopically in rounds 0 and 1 unless one or more of the people chooses action 1 in round 0 or 1, in which case he chooses action 1 in all future rounds. Otherwise if s is the action of the regent at time 1 then from time 2 the king takes action s until the regent deviates and chooses another action.
- The members of the bureaucracy act myopically in round 0 and 1. If s is the action of the regent at time 1 then from time 2 the members of the bureaucracy take action s until the regent deviates and chooses another action.
- The members of the court act myopically in round 0 and 1. At time 2 they copy the action of the king from time 1. If s is the action of the king at time 2 then from time 3 the members of the bureaucracy take action s until the king deviates and chooses another action.
- The people choose action 0 in rounds 1 and 2. At time 2 they copy the action of the king from time 1. If s is the action of the king at time 2 then from time 3 the people take action s until the king deviates and chooses another action.

By Lemma I.3 this can be extended to a forced equilibrium strategy \bar{Q} . We will show that this is also an equilibrium strategy in the unrestricted game by establishing equation (10). In what follows when we say acts optimally or in an optimal strategy we mean for an agent with respect to the actions of the other agents under \bar{Q} .

First consider the regent. By our choice of μ_0, μ_1 we have that $Y_0 > \frac{1}{5}$. Let $\mathcal{J} = \mathcal{J}_0 \cup \mathcal{J}_1$ where \mathcal{J}_s denotes the event that $\mathbb{P}[S = s | \mathcal{F}_1^v] \geq 1 - e^{-\delta R_B}$. Since the regent views all the

myopic actions of the bureaucracy he knows the correct value of S except with probability exponentially small in R_B so for $s \in \{0, 1\}$, if $\delta > 0$ is small enough,

$$\mathbb{P}[\mathcal{J}_s | S = s] \geq 1 - e^{-\delta R_B}$$

and hence for large enough R_B we have that $Y_1 \geq \frac{1}{2} - 2e^{-\delta R_B}$ which implies that \mathcal{B}_2 holds at time 1. By Lemma I.5 in any optimal strategy the regent acts myopically in round 0, and so follows the forced move. On the event \mathcal{J}_s the regent follows s in all future steps. At time 1 condition \mathcal{B}_1 holds so again the regent follows the forced move in any optimal strategy. We next claim that for large enough R_B

$$\mathbb{P}[\mathbb{P}[S = s | \mathcal{F}_2^v] \geq 1 - e^{-\delta R_B/2} | \mathcal{J}_s] = 1 \quad (12)$$

Assuming (12) holds then condition \mathcal{B}_1 again holds so the regent must choose s at time 2 in any optimal strategy. By construction of the forced moves from time 2 onwards the king and bureaucracy simply imitate the regent and so he receives no further information from time 2 onwards. Thus again using Lemma I.5 we see that under any optimal strategy the regent must follow his forced moves.

To establish that the regent follows the forced moves in any optimal strategy it remains to show that Condition (12) holds. The information available to the regent at time 2 includes the actions of the king and the bureaucracy at times 0 and 1. Consider the actions of the bureaucracy at times 0 and 1. At time 0 they follow their initial signal. At time 1 they also learn the initial action of the regent who acts myopically. By our assumption on μ_0 and μ_1 that $\mathbb{P}[Z_0^i \in (1, 1 + \epsilon) \cup (-\sqrt{7}, -\sqrt{7} + \epsilon)] = 1$, an initial signal towards 0 is much stronger than an initial signal towards 1, since whenever Z is negative it is at most $-\sqrt{7} + \epsilon$. For i , a member of the bureaucracy, we have that $Z_1^i \geq 2$ if both i and the regent choose action 1 at time 1. However, if either i or the regent choose action 0 at time 1 then $Z_t^i \leq -\sqrt{7} + \epsilon + 1 + \epsilon < -1$. Since the actions of i and the regent at time 0 are known to the regent at time 1, he gains no extra information at time 2 from his observation of i at time 1 since he can correctly predict his action.

The information the regent has available at time 2 is thus his information from time 1 together with the information from observing the king. The information available to the king is a function of his initial signal and that of the regent and the court. Since this is only $R_C + 1$ members and we choose R_B to be much larger than R_C it is insignificant compared to the information the regent observed from the court at time 0 and hence (12) holds. Thus, there is no optimal strategy for the regent that deviates from the forced moves.

As we noted above the members of the bureaucracy have $|Z_0^i|, |Z_1^i| \geq 1$ almost surely. For $t \geq 1$ let $\mathcal{M}_{s,t}$ denote the event that the regent chose action s for times 1 up to t . As argued above, $\mathcal{J}_s \subset \mathcal{M}_{s,t}$ for all t under \bar{Q} . This analysis holds even if a single member of the bureaucracy adopts a different strategy as we have taken R_B to be large so this change is insignificant. Given that $\mathcal{M}_{s,t}$ holds, the only additional information available to agent i , a member of the bureaucracy, is their original signal and the action at time 1 of the regent. Thus

$$\mathbb{P}[S = s | \mathcal{F}_t^i, \mathcal{M}_{s,t}] \geq 1 - e^{-\delta R_B/2}.$$

It follows then by Lemma I.5 that acting myopically at times 0 and 1 and then imitating the regent until he changes his action is the sole optimal strategy for a member of the bureaucracy.

Next consider the forced responses of the king. Since under \bar{Q} the people always choose action 0 at times 0 and 1, the rule forcing the king to choose action 1 after seeing a 1 from the people is never invoked. We claim that, provided R_B is taken to be sufficiently large, that the king acts myopically at times 0 and 1. At time 0 the posterior probability of $S = 1$ is bounded away from $1/2$ so Y_0 is bounded away from 0 while $\frac{1}{2} - Y_2 \leq 2e^{-\delta R_B/2}$ so by Lemma I.5 the king must act myopically. Similarly at time 1 since our choice of μ_0 and μ_1 to have their log-likelihood ratio concentrated around either 1 or $-\sqrt{7}$ a posterior calculation gives that,

$$|Z_1^u - \#\{i \in N(u) : A_0^i(\bar{Q}) = 1\} + \sqrt{7}\#\{i \in N(u) : A_0^i = 0\}| \leq \epsilon(2 + R_C)$$

and thus for some $\epsilon(R_C) > 0$ sufficiently small we can find an $\epsilon'(\epsilon, R_C) > 0$ such that $Y_0 = \left| \frac{e^{Z_1^u}}{1+e^{Z_1^u}} - \frac{1}{2} \right| > \epsilon'$. Since we again have that $\frac{1}{2} - Y_1 \leq 2e^{-\delta R_B/2}$ taking $R_B = R_B(\epsilon, R_C)$ to be sufficiently large \mathcal{B}_2 holds and so the king must act myopically. It remains to see that the king should imitate the regent from time 2 onwards unless the regent subsequently changes his action in any optimal strategy. This follows from a similar analysis to the case of the members of the bureaucracy so we omit it.

We next move to an agent i , a member of the court. At time 0 the agent has $Y_0 > \frac{e}{1+e} - \frac{1}{2} > \frac{1}{5}$. Agent i at time 1 views the action of the king who has in turn viewed the actions of the whole court at time 0 so $\frac{1}{2} - Y_2 \leq e^{-cR_C}$. At time 2 the agent sees the action of the king who has imitated the action of the regent at time 1 so $\frac{1}{2} - Y_3 \leq e^{-\delta R_B/2}$. Hence provided that R_C is sufficiently large and $R_B(R_C, \lambda)$ is sufficiently large then \mathcal{B}_4 holds and i must act myopically at time 0. The information of a member of the court at time 1 is a combination of their initial signal and the action of the king at time 1. Similarly to a member of the bureaucracy, by the choice of μ_0 and μ_1 we have that $|Z_1^i| \geq 1$ and so $Y_0 > \frac{1}{5}$. Also $\frac{1}{2} - Y_2 \leq e^{-\delta R_B/2}$ since this includes the information from the action of the regent at time 1. Thus \mathcal{B}_3 holds and i must act myopically at time 1. At time 2 agents i knows the action of the king from round 2 so $Y_0 \geq \frac{1}{2} - e^{-cR_C}$ and $\frac{1}{2} - Y_1 \leq e^{-\delta R_B/2}$ so \mathcal{B}_2 holds and i must act myopically at time 2. Finally from time 3 onwards agent i knows the action of the regent at time 1. As with the king and bureaucracy this will not be changed unless i receives new information, that is the king changes his action sometime after time 2. Thus any optimal strategy of i follows the forced moves.

This finally leaves the people. Let agent i be one of the people. We first check that it is always better for them to wait and just say 0 in rounds 0 and 1 in order to get more information from the king, their only source. If agent i chooses action 1 at time 0 then the total information it receives is a function of the initial signals of i and the king. Thus, since the signals are uniformly bounded, even if the agent knew the signals exactly we would have that for some $c'(\mu_0, \mu_1)$ that the expected utility from such a strategy is at most $1 - e^{-2c'}$. If an agent acts with 0 at time 0 but 1 at time 1, she can potentially receive information from the initial signals of the king, court and regent as well as her own; still, the optimal expected utility even using all of this information is at most $1 - e^{-c'(R_C+3)}$. Consider instead the expected utility following the forced moves. On the event \mathcal{J} agent i will have expected utility at least $\lambda^3(1 - e^{-\delta R_B})$ which is greater than $1 - e^{-c'(R_C+3)}$ provided that λ is sufficiently close to 1 and R_B is sufficiently large. Thus agent i must choose action 0 at times 0 and 1 in any optimal strategy. The analysis of rounds 2 and onwards follows similarly to the court and thus any optimal strategy of i follows all the forced moves.

This exhaustively shows that there is no alternative optimal strategy for any of the agents which differs from the forced moves. Thus \bar{Q} is an equilibrium. However, on the event \mathcal{J}_1 all the agents actions converge to 1. However, $\mathbb{P}[\mathcal{J}, S = 0] \geq e^{-c''R_B} > 0$ where c'' is independent of R_C, R_B, λ and n . Hence, as we let n tend to infinity the probability of learning does not tend to 1.