

A Theory of Non-Bayesian Social Learning*

Pooya Molavi[†]

Alireza Tahbaz-Salehi[‡]

Ali Jadbabaie[§]

December 2017

Abstract

This paper studies the behavioral foundations of non-Bayesian models of learning over social networks and develops a taxonomy of conditions for information aggregation in a general framework. As our main behavioral assumption, we postulate that agents follow social learning rules that satisfy “imperfect recall,” according to which they treat the current beliefs of their neighbors as sufficient statistics for the entire history of their observations. We augment this assumption with various restrictions on how agents process the information provided by their neighbors and obtain representation theorems for the corresponding learning rules (including the canonical model of DeGroot). We then obtain general long-run learning results that are not tied to the learning rules’ specific functional forms, thus identifying the fundamental forces that lead to learning, non-learning, and mislearning in social networks. Our results illustrate that, in the presence of imperfect recall, long-run aggregation of information is closely linked to (i) the rate at which agents discount their neighbors’ information over time, (ii) the curvature of agents’ social learning rules, and (iii) whether their initial tendencies are amplified or moderated as a result of social interactions.

Keywords: social networks, non-Bayesian learning, imperfect recall.

JEL Classification: D83, D85.

*An earlier draft of this paper was circulated under the title “Foundations of non-Bayesian Social Learning.” We thank the co-editor and four anonymous referees for helpful remarks and suggestions. We also thank Daron Acemoglu, Kostas Bimpikis, Emily Breza, Christophe Chamley, Arun Chandrasekhar, Matt Elliott, Drew Fudenberg, Ben Golub, Sanjeev Goyal, Matt Jackson, Ilan Lobel, Mohamed Mostagir, Pietro Ortoleva, Andy Postlewaite, Amin Rahimian, Matthew Rabin, Kamiar Rahnama-Rad, Ronny Razin, Johannes Ruf, Evan Sadler, Alvaro Sandroni, Adam Szeidl, Juuso Toikka, Andrea Vedolin, Ricky Vohra, Alex Wolitzky, Jeff Zwiebel, and seminar participants at Bocconi, Boston University, Cambridge, Cleveland Fed, Columbia, Duke, IESE, MIT, Penn, Pompeu Fabra, Stanford, Fourth European Meeting on Networks (Aix-Marseille School of Economics), the Second Annual Conference on Network Science and Economics (Stanford), and the 2016 Workshop on Information and Social Economics (Caltech) for useful feedback and suggestions. Jadbabaie acknowledges financial support from the Army Research Office (MURI Award No. W911NF-12-1-0509).

[†]Department of Economics, Massachusetts Institute of Technology, mpooya@mit.edu.

[‡]Kellogg School of Management, Northwestern University, alirezat@kellogg.northwestern.edu.

[§]Institute for Data, Systems, and Society, Massachusetts Institute of Technology, jadbabai@mit.edu.

1 Introduction

The standard model of rational learning maintains that individuals use Bayes' rule to incorporate any new piece of information into their beliefs. In addition to its normative appeal, this Bayesian paradigm serves as a highly useful benchmark by providing a well-grounded model of learning. Despite these advantages, a growing body of evidence has scrutinized this framework on the basis of its unrealistic cognitive demand on individuals, especially when they make inferences in complex environments consisting of a large number of other decision-makers. Indeed, the complexity involved in Bayesian learning becomes particularly prohibitive in real-world social networks where people have to make inferences about a wide range of parameters while only observing the actions of a handful of individuals.

To address these issues, a growing literature has adopted an alternative paradigm by assuming non-Bayesian behavior on the part of the agents. These models, which for the most part build on the linear model of DeGroot (1974), impose relatively simple functional forms on agents' learning rules, thus capturing the richness of the network interactions while maintaining analytical and computational tractability. Such heuristic non-Bayesian models, however, can in turn be challenged on several grounds. First, in many instances, the suggested heuristics are at best only loosely connected to the behavioral assumptions that are used to motivate them. Second, although Bayesian learning is a well-defined concept, deviations from the Bayesian benchmark are bound to be ad hoc and arbitrary. Third, it is often unclear whether the predictions of such heuristic models rely on ancillary behavioral assumptions baked into their specific functional forms or are illustrative of more robust and fundamental forces.

In this paper, we address these challenges by taking an axiomatic approach towards social learning: rather than assuming a specific functional form for agents' social learning rules, we use a general framework to uncover the structure of social learning rules under a variety of behavioral assumptions. This approach not only enables us to provide a systematic way of capturing deviations from Bayesian inference but also reveals fundamental forces that are central to information aggregation and may be obscured by the restrictions built into the functional forms commonly used in the literature. In particular, we obtain general long-run learning results that are not tied to the specific functional form of the learning rules and identify the forces that lead to learning, non-learning, and mislearning in social networks.

We consider an environment in which agents obtain information about an underlying state through private signals and communication with other agents in their social clique. As our main behavioral assumption, we postulate that agents follow social learning rules that satisfy *imperfect recall*, according to which they treat the current beliefs of their neighbors as sufficient statistics for all the information available to them while ignoring how or why these opinions were formed. Besides being a prevalent assumption in the models of non-Bayesian learning such as DeGroot's, imperfect recall is a formalization of the idea that real-world individuals do not fully account for the information buried in the entire past history of actions or the complex dynamics of beliefs over social networks. We then supplement this assumption by a variety of additional assumptions on agents' behavior to obtain sharp characterizations of its implications.

We start our analysis by imposing three natural restrictions on how agents process their neighbors' information to obtain a simple learning rule that will serve as a benchmark for the rest of our results. First, we assume that agents' social learning rules are *label neutral*, which means that relabeling the underlying states has no bearing on how agents process information. Second, we assume that individuals do not discard their neighbors' most recent observations by requiring their social learning rules to be increasing in their neighbors' last period beliefs, a property we refer to as *monotonicity*. Third, we require agents' learning rules to satisfy *independence of irrelevant alternatives* (IIA): each agent treats her neighbors' beliefs about any subset of states as sufficient statistics for their collective information regarding those states. Besides their simplicity and intuitive appeal, these three restrictions are satisfied by Bayesian agents when the social network satisfies certain structural properties.¹

As our first result, we show that, in conjunction with imperfect recall, these three restrictions lead to a unique representation of agents' social learning rules up to a set of constants: at any given time period, each agent linearly combines the log-likelihood ratios of her private signal with the log-likelihood ratios of her and her neighbors' beliefs in the previous period. Given its functional form, we refer to this learning rule as *log-linear learning*.

Our representation theorem reveals that all other non-Bayesian models of social learning that satisfy imperfect recall (such as DeGroot's model) must violate at least one of the other three restrictions. To further clarify this point, we then shift our focus to DeGroot's model and show that this learning rule indeed violates the IIA assumption. In fact, we provide a second representation theorem by establishing that DeGroot's model is the unique learning rule that satisfies imperfect recall, label neutrality, monotonicity, and a fourth, alternative restriction, which we refer to as *separability*. This assumption, which serves as an alternative notion of independence to IIA, requires the posterior belief that each agent assigns to any given state to be independent of her neighbors' opinions about any other state.

Given the different functional forms and distinct foundations of the log-linear and DeGroot learning models, it is not surprising that agents who follow these two rules process information differently and, as a result, have distinct beliefs at any given (finite) time. Nonetheless, we show that the two learning rules have analogous implications for agents' long-run beliefs. In particular, we show that agents asymptotically learn the underlying state in either model as long as social learning rules satisfy two key conditions. The first condition requires the weights that agents use to take their neighbors' information into account to decay to zero at a sufficiently slow rate, if at all. The role of this condition is to ensure a continuous flow of information from agents who observe informative signals to those who may not. The second condition, which we refer to as *unanimity*, requires each agent to adopt the common belief of her neighbors whenever they all agree with one another. Unanimity guarantees that independent pieces of information that are revealed to agents over time are incorporated into their beliefs using roughly equal weights, thus enabling agents to accumulate information over time.

Motivated by the parallels between the log-linear and DeGroot learning rules' long-run properties,

¹We provide a formal representation of these structural properties in Section 3.

we then develop a taxonomy of conditions for asymptotic learning, non-learning, and mislearning that is not tied to the specific functional form of agents' learning rules, thus identifying the key underlying forces that shape long-run beliefs in the presence of imperfect recall.

We achieve this by replacing IIA and separability with a weaker notion of independence and obtaining a general class of learning rules that encompasses the log-linear and DeGroot models as special cases. According to this notion, which we refer to as *weak separability*, each agent's posterior likelihood ratio over a pair of states can be expressed as the ratio of some homogeneous mapping ψ applied to her neighbors' beliefs on the same two states at the previous time period, with each choice of ψ leading to a distinct learning rule. We then identify two characteristics of weakly-separable learning rules that — once the continuous flow of information among agents is ensured — determine the long-run outcomes of social learning: (i) the degree of homogeneity of ψ , which determines whether the underlying learning rule is unanimous, and (ii) its logarithmic curvature (defined as the curvature in the log-log scale), which measures the learning rule's departure from the benchmark of log-linear learning.

We underscore the role of the degree of homogeneity of ψ in shaping agents' long-run beliefs by establishing that learning is guaranteed only if ψ is homogeneous of degree 1 — a condition that is equivalent to unanimity for weakly-separable learning rules. More specifically, we show that individuals asymptotically mislearn the underlying state (i.e., become confident that a false state is true) with positive probability if ψ is a homogeneous function of a degree greater than 1. Agents who follow such learning rules overweigh evidence they encounter early on at the expense of the more recent pieces of information; consequently, every round of social interaction reinforces individuals' initial tendencies, hence potentially driving them to asymptotic mislearning. We also obtain a diametrically opposite result by establishing that whenever ψ is homogeneous of a degree less than 1, agents remain uncertain forever about the underlying state, as they downplay the already accumulated information in favor of their more recent signals. These results thus demonstrate the importance of unanimity as a condition for successful aggregation of information in the long run.

But, as our next result shows, the unanimity of learning rules is not sufficient for the asymptotic aggregation of information over the social network; rather, information aggregation also requires a restriction on the curvatures of agents' social learning rules. This restriction is formally captured by the requirement that the logarithmic curvature of ψ falls within the interval $[-1, 1]$, with the log-linear learning rule serving as the benchmark with logarithmic curvature equal to 0. We show that when this condition is violated, the distortion in how each agent aggregates her neighbors' information (relative to the benchmark of log-linear learning) is so large that she may mislearn the underlying state or remain uncertain forever, even if agents' social learning rules are unanimous. Taken together, our results provide a set of conditions that lead to asymptotic learning, non-learning, and mislearning in the presence of imperfect recall.

We conclude the paper by characterizing the speed of information aggregation, defined as the rate at which individuals rule out the incorrect states. Our result illustrates that this rate is determined by (i) the detailed structure of the social network and (ii) how information is dispersed among different agents in the society. More specifically, we show that the rate of learning is given by the long-run

average of agents' relative entropies, which measure the informativeness of their private observations, weighted by a novel notion of network centrality that measures each agent's importance as a source of information at a given time. A key consequence of this result is that, in the presence of non-trivial social interactions, reallocating signals across various agents may have a first-order effect on the speed of information aggregation, even if the total information content of agents' observations is kept constant.

Our paper belongs to the literature that studies non-Bayesian learning over social networks, such as [DeMarzo, Vayanos, and Zwiebel \(2003\)](#) and [Golub and Jackson \(2010, 2012\)](#). The standard approach in this literature is to analyze belief dynamics while imposing a specific functional form on agents' social learning rules. We part ways with these papers by taking an axiomatic approach and studying the broader class of learning rules that are subject to imperfect recall. This approach enables us to not only determine the restrictions that give rise to various non-Bayesian models of social learning within this class but also obtain general and novel long-run learning results.

In parallel to the non-Bayesian literature, a large body of work has focused on Bayesian learning over social networks. Going back to the works of [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#), this literature explores the implications of Bayesian inference in an environment where individuals can only observe the actions and/or beliefs of a subset of other agents.² Our work is related to a recent stream of papers that study how specific departures from the Bayesian paradigm alter the predictions of these models. For example, [Eyster and Rabin \(2010, 2014\)](#) study the long-run aggregation of information when people fail to appreciate redundancies in the information content of others' actions. Similarly, [Rahimian, Molavi, and Jadbabaie \(2014\)](#) consider a model in which an individual does not account for the fact that her neighbors' beliefs are in turn affected by their own social interactions, whereas [Li and Tan \(2017\)](#) assume that each agent updates her belief as if her local neighborhood is the entire society. We contribute to this literature by focusing on imperfect recall as our main behavioral assumption and studying its implications for agents' long-run beliefs.

Our result on the rate of learning is related to the work of [Harel, Mossel, Strack, and Tamuz \(2017\)](#), who characterize the speed of learning for Bayesian agents located on a social network. It is also reminiscent of the work of [Golub and Jackson \(2010\)](#), who relate agents' eigenvector centralities to their asymptotic beliefs in the DeGroot model. We complement this result by showing that when agents use non-stationary learning rules, the rate of learning depends on a novel notion of network centrality that is defined recursively over both time and space.

The rest of the paper is organized as follows. The formal setup is presented in Section 2, where we also introduce the notion of imperfect recall as our main behavioral assumption. In Section 3, we present our first representation theorem for the log-linear learning rule, followed by the foundations of the DeGroot model. We then focus on the long-run implications of various learning rules in Section 4 and provide a taxonomy of conditions for information aggregation. The rate of learning is characterized in Section 5. Proofs are provided in the Appendix, while an Online Appendix contains several omitted proofs and some additional results.

²See [Golub and Sadler \(2016\)](#) for a survey of the social learning literature, covering both Bayesian and non-Bayesian paradigms.

2 Setup

Consider a collection of n individuals, denoted by $N = \{1, 2, \dots, n\}$, who are attempting to learn an underlying state of the world θ . The underlying state is drawn at $t = 0$ from some finite set Θ according to the uniform distribution.

Even though the realized state is unobservable, individuals make repeated noisy observations about θ in discrete time. At each time period $t \in \mathbb{N}$ and conditional on the realization of state θ , agent i observes a private signal ω_{it} that is drawn from a finite set S according to the potentially time-dependent distribution $\ell_{it}^\theta \in \Delta S$. The realized signals are conditionally independent across individuals and over time. We assume that ℓ_{it}^θ has full support over S for all $\theta \in \Theta$, but we do not require an individual's private signals to be informative about the state. In particular, agent i may face an identification problem in the sense of not being able to distinguish between two or more states.³

In addition to her private signals, each agent observes the beliefs of a subset of other agents, whom we refer to as her *neighbors*. More specifically, at the beginning of time period t and before observing the realization of her private signal ω_{it} , agent i observes the beliefs held by her neighbors at the previous time period.⁴ This form of social interactions can be represented by a directed graph on n vertices, which we refer to as the *social network*. Each vertex of this graph corresponds to an agent and a directed edge is present from vertex j to vertex i if agent i can observe the beliefs of agent j . Throughout the paper, we use N_i to denote the set consisting of agent i and her neighbors.

We assume that the underlying social network is strongly connected; that is, we assume that there exists a directed path from each vertex to any other vertex. This assumption ensures that the information available to any given agent can potentially flow to other individuals in the social network. We define the social network's *diameter* as the length of the shortest (directed) path between the two individuals who are farthest from one another.

2.1 Social Learning Rules

At any given period, agents use their private observations and the information provided to them by their neighbors to update their beliefs about the underlying state. In particular, each agent first combines her prior belief with the information provided to her by her neighbors to obtain an interim belief. Following the observation of her private signal, she updates this interim belief in a Bayesian fashion to form her posterior belief. The belief of agent i at the end of period t is thus given by

$$\mu_{it+1} = \text{BU}(f_{it}(\mu_i^t); \omega_{it+1}), \quad (1)$$

where $\mu_i^t = (\mu_{j\tau})_{j \in N_i, 0 \leq \tau \leq t}$ is the history of beliefs of i and her neighbors up to period t and $\text{BU}(\mu; \omega)$ denotes the Bayesian update of μ conditional on the observation of signal ω . The function

³For instance, there may exist a pair of states $\theta \neq \hat{\theta}$ such that $\ell_{it}^\theta = \ell_{it}^{\hat{\theta}}$ for all t . When studying conditions for long-run learning in Section 4, we require that agents' observations are collectively informative by assuming that, for any distinct pair of states, there exists an agent with access to private signals that enable her to distinguish the two states in the long run.

⁴The observational learning literature for the most part assumes that agents can observe their neighbors' actions, as opposed to their beliefs. We abstract from actions and simply assume that individuals have access to their neighbors' beliefs. The observability of beliefs is equivalent to that of actions whenever the action space is "rich" enough that an individual's actions fully reveal her beliefs.

$f_{it} : \Delta\Theta^{|N_i|(t+1)} \rightarrow \Delta\Theta$, which we refer to as the *social learning rule* of agent i , is a continuous mapping that captures how she incorporates the information provided by her neighbors into her belief.⁵ Throughout the paper, we assume that agents share a (common) uniform prior belief over Θ at $t = 0$.

Although each agent incorporates her private signals into her beliefs in a Bayesian fashion, our flexible specification of social learning rules allows agents to follow alternative (and hence potentially non-Bayesian) updating rules for processing their neighbors' information. The disparity between the ways agents process their private and social information in (1) is imposed for two reasons. First, it is natural to expect that agents find it easier to rationally process their private signals compared to the information provided by other individuals: whereas each agent's private signals are distributed according to a distribution known to her, her neighbors' beliefs may encompass multiple pieces of potentially redundant information, which she may find hard to disentangle without complete knowledge of the social network or other agents' signal structures. Second and more importantly, the assumption that agents incorporate their private signals into their beliefs in a Bayesian fashion serves as a natural benchmark for our forthcoming results; it guarantees that any deviation from the predictions of Bayesian learning is driven by the nature of agents' social learning rules, as opposed to how they process their private signals.

2.2 Imperfect Recall

We now introduce our main behavioral assumption on agents' social learning rules by assuming that agents take the current beliefs of their neighbors as sufficient statistics for all the information available to them while ignoring how or why those opinions were formed. Formally:

Imperfect Recall (IR). $f_{it}(\mu_i^t)$ is independent of $\mu_{j\tau}$ for all j and all $\tau \leq t - 1$.

The restriction imposed by imperfect recall represents a departure from Bayesian rationality. For instance, in a social network consisting of two Bayesian agents, agent i can make inferences about j 's latest private signal only by comparing j 's current belief to her belief in the previous period; yet such a comparison is ruled out by the imperfect recall assumption. More generally, Bayesian inference requires agents to (i) keep track of the entire history of their neighbors' beliefs, (ii) determine the source of all the information they have observed so far, and (iii) extract any piece of new information not already incorporated into their beliefs in the previous time periods, while only observing the evolution of their neighbors' opinions. Such complicated inference problems — which are only intensified if agents are also uncertain about the social network structure — require a high level of sophistication on the part of the agents. In contrast, under IR, even though agent i may use a different learning rule at any given time instance t , she simply treats her neighbors' most recent opinions as sufficient statistics for all the information available to her while ignoring the rest of the history.

We remark that the deviation from Bayesian rationality captured by imperfect recall is a fairly standard notion of bounded rationality that is (implicitly or explicitly) imposed in a wide range of

⁵With some abuse of notation, we treat the social learning function f_{it} as if its domain is $\Delta\Theta^{n(t+1)}$ (as opposed to $\Delta\Theta^{|N_i|(t+1)}$) with the understanding that f_{it} does not depend on the beliefs of agents who are not i 's neighbors.

non-Bayesian learning models in the literature. Most notably, the DeGroot model and its different variations (e.g., [Golub and Jackson \(2010\)](#) and [Chandrasekhar, Larreguy, and Xandri \(2016\)](#)) rely on imperfect recall by assuming that agents only use their neighbors' reports from the previous period.

3 Foundations

In this section, we supplement the assumption of imperfect recall by a variety of other restrictions on agents' behavior and obtain representation theorems for agents' social learning rules. This approach not only enables us to identify the behavioral assumptions that underpin various learning rules but also provides us with a systematic way of uncovering the forces that shape the agents' long-run beliefs in the presence of imperfect recall, an issue that we will study in detail in Section 4.

3.1 Log-Linear Learning

We start our analysis by imposing three simple restrictions — other than imperfect recall — on how agents process their neighbors' information to obtain a learning rule that will serve as a benchmark for the rest of our results. Besides their simplicity and intuitive appeal, these three restrictions are satisfied by Bayesian agents when the social network satisfies certain structural properties.⁶ As our first result, we establish that there exist a unique social learning rule (up to a set of constants) that satisfies all four restrictions. According to this learning rule, which we refer to as *log-linear learning*, each agent linearly combines the log-likelihood ratios of her private signal with the log-likelihood ratios of her and her neighbors' beliefs in the previous period.

As the first restriction, we require that relabeling the underlying states has no bearing on how agents process information. For any permutation $\sigma : \Theta \rightarrow \Theta$ on the set of states, let $\text{perm}_\sigma : \Delta\Theta \rightarrow \Delta\Theta$ denote the operator that maps a belief to the corresponding belief after relabeling the states according to σ ; that is, $\text{perm}_\sigma(\mu)(\theta) = \mu(\sigma(\theta))$ for all θ .

Label Neutrality (LN). For any permutation $\sigma : \Theta \rightarrow \Theta$ and all histories μ_i^t ,

$$\text{perm}_\sigma(f_{it}(\mu_i^t)) = f_{it}(\text{perm}_\sigma(\mu_i^t)),$$

where $\text{perm}_\sigma(\mu_i^t) = (\text{perm}_\sigma(\mu_{j\tau}))_{j \in N_i, 0 \leq \tau \leq t}$.

Under label-neutral learning rules, any asymmetry in how an individual updates her opinion about different states is only due to asymmetries in her, or her neighbors' subjective beliefs about those states as opposed to how different states are labeled.

The next restriction requires agents to respond to an increase in their neighbors' beliefs on a given state by increasing their own posterior beliefs on that state in the next period. Formally:

Monotonicity. If for some $j \in N_i$ observation histories μ_i^t and ν_i^t satisfy (i) $\mu_i^{t-1} = \nu_i^{t-1}$, (ii) $\mu_{kt} = \nu_{kt}$ for all $k \in N_i \setminus \{j\}$, (iii) $\mu_{jt}(\theta) > \nu_{jt}(\theta)$, and (iv) $\mu_{jt}(\hat{\theta}) \leq \nu_{jt}(\hat{\theta})$ for all $\hat{\theta} \neq \theta$, then $f_{it}(\mu_i^t)(\theta) > f_{it}(\nu_i^t)(\theta)$.

⁶See Subsection 3.3 for a sufficient condition on the structure of the social network under which the learning rules employed by Bayesian agents satisfy these three restrictions.

Monotonicity captures the idea that keeping the history of observations μ_i^{t-1} and the time t reports of agents $k \neq j$ fixed, agent i interprets an increase in $\mu_{jt}(\theta)$ as evidence that either agent j has observed a private signal in favor of θ at period t or that j 's neighbors whose beliefs are unobservable to i have provided j with such information. Under either interpretation, agent i finds an increase in $\mu_{jt}(\theta)$ as more evidence in favor of θ and hence increases the belief she assigns to that state.

To state our next restriction on agents' social learning rules, let $\text{cond}_{\bar{\Theta}} : \Delta\Theta \rightarrow \Delta\Theta$ denote the operator that maps a belief to the corresponding belief conditioned on the subset of states $\bar{\Theta} \subseteq \Theta$; that is, $\text{cond}_{\bar{\Theta}}(\mu)(\theta) = \mu(\theta|\bar{\Theta})$.

Independence of Irrelevant Alternatives (IIA). For any subset of states $\bar{\Theta} \subseteq \Theta$ and all histories μ_i^t ,

$$\text{cond}_{\bar{\Theta}}(f_{it}(\mu_i^t)) = f_{it}(\text{cond}_{\bar{\Theta}}(\mu_i^t)), \quad (2)$$

where $\text{cond}_{\bar{\Theta}}(\mu_i^t) = (\text{cond}_{\bar{\Theta}}(\mu_{j\tau}))_{j \in N_i, 0 \leq \tau \leq t}$.

The above restriction requires the conditional belief of agent i after aggregating her neighbors' opinions to be identical to the belief obtained by aggregating her neighbors' conditional beliefs using the same social learning rule. Thus, under IIA, i 's conditional posterior belief exclusively depends on the history of her and her neighbors' beliefs on the states in the conditioning set $\bar{\Theta}$ and is independent of beliefs assigned by any individual to $\theta \notin \bar{\Theta}$ in any of the previous time periods. Put differently, as far as agent i is concerned, her neighbors' beliefs about states in $\bar{\Theta}$ are sufficient statistics for their collective information regarding all $\theta \in \bar{\Theta}$.

With the above three restrictions in hand, we now provide a characterization of agents' social learning rules:

Theorem 1. *Suppose $|\Theta| \geq 3$. If agents' social learning rules satisfy IR, LN, monotonicity, and IIA, there exist constants $a_{ijt} > 0$ such that*

$$\log \frac{f_{it}(\mu_i^t)(\theta)}{f_{it}(\mu_i^t)(\hat{\theta})} = \sum_{j \in N_i} a_{ijt} \log \frac{\mu_{jt}(\theta)}{\mu_{jt}(\hat{\theta})} \quad (3)$$

for all $\theta, \hat{\theta} \in \Theta$.

The significance of this characterization is twofold. First, it shows that the restrictions imposed by IR, LN, monotonicity, and IIA yield a unique representation of agents' social learning rules up to a set of constants. Second, Theorem 1 establishes that all other non-Bayesian models of social learning in which agents interact with one another repeatedly and satisfy imperfect recall (such as DeGroot's model) violate at least one of the other three restrictions.

It is instructive to elaborate on the role of each assumption in determining the functional form of the social learning rule in (3). Imperfect recall requires i 's posterior belief at time $t+1$ to solely depend on other agents' beliefs at time t . The log-linear nature of the learning rule, on the other hand, is a consequence of LN and IIA. In particular, IIA guarantees that the ratio of i 's posterior beliefs on any two states only depends on her and her neighbors' likelihood ratios for those two states.⁷ Given that

⁷Note that when $|\Theta| = 2$, IIA is trivially satisfied and hence does not impose any restrictions on agents' social learning rules. Consequently, our representation theorem requires Θ to consist of at least three elements.

such independence holds for any pair of states, LN implies that the only possible functional form has to be linear in agents' log-likelihood ratios. In addition, label neutrality guarantees that constants a_{ijt} do not depend on the pair of states θ and $\hat{\theta}$ under consideration. Finally, the positivity of these constants is an immediate implication of the monotonicity assumption.

A key implication of the representation in Theorem 1 is that under IR, LN, monotonicity, and IIA agents' belief dynamics are given by

$$\log \frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \log \frac{\ell_{it+1}^{\theta}(\omega_{it+1})}{\ell_{it+1}^{\hat{\theta}}(\omega_{it+1})} + \sum_{j \in N_i} a_{ijt} \log \frac{\mu_{jt}(\theta)}{\mu_{jt}(\hat{\theta})} \quad (4)$$

for all $\theta, \hat{\theta} \in \Theta$. Thus, at every period, agent i linearly combines the log-likelihood ratios of her private signal with the log-likelihood ratios of her and her neighbors' beliefs in the previous time period, with a_{ijt} representing the weight that i assigns to the belief of agent j in her neighborhood at time t .

Note that the representation in (4) does not impose any restrictions on constants a_{ijt} besides positivity. However, as we show in Section 4, whether the above learning rule results in the long-run aggregation of information depends on the weights that agents assign to their neighbors' beliefs.

We also emphasize that the assumption that agents incorporate their private signals into their beliefs in a Bayesian fashion does not play a crucial role in our characterization. More specifically, altering the way agents process the information content of their private signals only impacts the way the first term on the right-hand side of (4) interacts with the rest of the expression while preserving the log-linear structure of the learning rule f_{it} .

As a final remark, we note that the log-linear learning rule characterized in Theorem 1 coincides with the ‘‘Bayesian Peer Influence’’ heuristic of [Levy and Razin \(2016\)](#), who show that individuals who treat their marginal information sources as (conditionally) independent follow a learning rule similar to (3) with corresponding weights given by $a_{ijt} = 1$. Our result provides a distinct foundation for log-linear updating by demonstrating the importance of the imperfect recall and IIA assumptions.

3.2 DeGroot Learning

A key implication of Theorem 1 is that any learning rule that satisfies imperfect recall but is distinct from (3) has to violate either LN, IIA, or monotonicity. One such model is the learning model of [DeGroot \(1974\)](#), which serves as the canonical model of non-Bayesian social learning in the literature. Under DeGroot learning and its many variants, agents update their beliefs by linearly combining their viewpoints with their neighbors' opinions in the previous time period. As such, it is immediate to see that DeGroot learning satisfies imperfect recall. Furthermore, as long as the linear weights used by the agents to incorporate their neighbors' beliefs are positive and independent of the underlying state θ , monotonicity and label neutrality are trivially satisfied. Consequently, by Theorem 1, DeGroot learning has to violate IIA. In fact, this can be easily verified by noting that no linear function f_{it} can satisfy condition (2) when Θ contains more than two elements. To further clarify this point, we propose the following new restriction on agents' social learning rules as an alternative to IIA:

Separability. $f_{it}(\mu_i^t)(\theta)$ does not depend on $\mu_{j\tau}(\hat{\theta})$ for any $\hat{\theta} \neq \theta$ and all j and $\tau \leq t$.

According to separability, the posterior belief that agent i assigns to any given state θ only depends on her and her neighbors' beliefs about θ and is independent of their opinions about any other state. Thus, separability imposes a different form of "independence" on agents' social learning rules than IIA, which requires the ratio of beliefs assigned to states θ and $\hat{\theta}$ to be a function of other agents' likelihood ratios of the same pair of states. We have the following representation theorem:

Theorem 2. *Suppose $|\Theta| \geq 3$. If agents' learning rules satisfy IR, LN, monotonicity, and separability, there exists a set of constants $a_{ijt} > 0$ and $c_{it} \geq 0$ such that*

$$f_{it}(\mu_i^t)(\theta) = c_{it} + \sum_{j \in N_i} a_{ijt} \mu_j(\theta) \quad (5)$$

for all $\theta \in \Theta$.

Therefore, replacing IIA with separability results in a learning rule according to which each agent's belief depends linearly on her neighbors' opinions in the previous time period, in line with DeGroot's model.⁸ As in the log-linear learning rule, LN guarantees that the weights that each agent assigns to her neighbors' beliefs are independent of the underlying state θ , while monotonicity implies that these weights are strictly positive for all $j \in N_i$.⁹

Just like Theorem 1, Theorem 2 imposes few restrictions, other than positivity, on the weights a_{ijt} used by the agents to incorporate their neighbors' beliefs.¹⁰ We return to these weights in Section 4, where we study how their various specifications determine the extent of information aggregation in the long run.

3.3 Bayesian Learning

We conclude this section by clarifying the relationship between the above-mentioned restrictions and Bayesian learning. More specifically, we provide a sufficient condition on the structure of the social network under which learning rules employed by Bayesian agents satisfy LN, IIA, and monotonicity. We also show that, under the same condition, Bayesian learning takes a log-linear functional form similar to the non-Bayesian learning rule studied in Subsection 3.1. As already argued, however, the learning rules employed by Bayesian agents do not satisfy imperfect recall.

Recall that the underlying social network is assumed to be strongly connected with at least one directed path from each agent to any other agent. We say two directed paths from agent j to agent i are *vertex independent* if they do not have any vertices (other than i and j) in common.

Assumption 1. For any distinct ordered pair of vertices (i, j) , if there are two or more vertex-independent paths from j to i , then j is a neighbor of i .

⁸The specification of the DeGroot model frequently used in the literature assumes that agents set their beliefs as a convex combination of their neighbors' beliefs. In our setting, this translates to assuming that $\sum_{j \in N_i} a_{ijt} = 1$ and $c_{it} = 0$ for all agents i . Also see Lehrer and Wagner (1981) for a characterization similar to ours, albeit under a different set of restrictions.

⁹Like IIA, separability is trivially satisfied whenever the state space consists of only two elements. The assumption that $|\Theta| \geq 3$ in Theorem 2 thus ensures that separability imposes a meaningful restriction on the learning rules.

¹⁰Sharper characterizations of these weights can be obtained by imposing more stringent restrictions on agents' social learning rules. For instance, one can tighten the definition of imperfect recall by requiring agent i 's learning rule f_{it} to be independent of not only her neighbors' reports from previous periods but also the time index t . This would lead to learning rules in which weights in (3) and (5) are time-invariant. See an earlier version of our work (Molavi, Tahbaz-Salehi, and Jadbabaie, 2016) for formal statements of our representation theorems under this more stringent restriction.




Figure 1. This figure depicts various strongly connected social networks that satisfy Assumption 1. Arrows indicate the direction over which information can flow from one agent to another, whereas undirected edges indicate a bidirectional flow of information between two agents.

Figure 1 depicts a few examples of social networks that satisfy this assumption. Intuitively, Assumption 1 guarantees that if the information available to agent j can reach agent i via multiple paths, then either (i) agent i can directly observe j 's beliefs and hence is able to discern any potential correlation in the reports of i 's neighbors that are attributable to j ; or (ii) the various paths from j to i are not independent, which implies that i can rely on some other agent $k \neq i, j$ to perform such a deduction on her behalf.¹¹ We can now state the following result:

Theorem 3. *Suppose the social network structure is common knowledge and satisfies Assumption 1. Under the common knowledge of Bayesian rationality,*

- (a) *there exists a collection of learning rules $(f_{1t}, \dots, f_{nt})_{t=0}^{\infty}$ independent of agents' beliefs and signal structures such that $\mu_{it+1} = \text{BU}(f_{it}(\mu_i^t); \omega_{it+1})$ for all i and all t ;*
- (b) *the learning rule f_{it} satisfies LN, IIA, and monotonicity for all agents i and all times t ;*
- (c) *there exist coefficients $a_{it,j\tau}$ independent of the beliefs and signal structures such that*

$$\log \frac{f_{it}(\mu_i^t)(\theta)}{f_{it}(\mu_i^t)(\hat{\theta})} = \sum_{\tau=1}^t \sum_{j \in N_i} a_{it+1,j\tau} \log \frac{\mu_{j\tau}(\theta)}{\mu_{j\tau}(\hat{\theta})} \quad (6)$$

for all agents i , all times t , and all pairs of states $\theta, \hat{\theta} \in \Theta$.

Statement (a) of the theorem establishes that whenever the social network satisfies Assumption 1, Bayesian updating has a representation in the form of equation (1): each agent combines her prior belief with the information provided to her by her neighbors using a mapping that is independent of agents' beliefs and signal structures (though in general this mapping depends on the structure of the underlying social network). More importantly for our purposes, statements (b) and (c) of Theorem 3 establish a sufficient condition under which the learning rule employed by a Bayesian agent satisfies the three key restrictions of LN, IIA, and monotonicity and takes a log-linear functional form akin to the (non-Bayesian) log-linear learning rule in equation (3). The key distinction between the two learning rules is that while Bayesian agents may use their neighbors' entire history of reports, non-Bayesian agents who are subject to imperfect recall only rely on their neighbors' reports in the very last period.

¹¹The existence of such an agent k is a consequence of what is known as Menger's Theorem in graph theory (McCuaig, 1984). See the proof of Theorem 3 for details.

Note that monotonicity — as defined in Subsection 3.1 and established in part (b) of the theorem — requires agent i to respond to an increase in her neighbors' beliefs by increasing her own posterior belief on the corresponding state in the *next* period. In the context of equation (6), this requirement translates to $a_{it+1,jt} > 0$ for all t and all $j \in N_i$. This, however, does not mean that i 's posterior belief is also increasing in her neighbors' earlier reports from $\tau < t$. In fact, as our characterization of weights $a_{it+1,j\tau}$ in the proof of Theorem 3 illustrates, Bayesian agents may assign negative weights to reports from $\tau < t$ in order to account for potential informational redundancies in their neighborhoods. This observation is in line with some of the results in the observational learning literature, such as Acemoglu et al. (2011) and Eyster and Rabin (2014), that show that Bayesian agents may revise their beliefs downwards in response to an increase in their predecessors' reports.

One consequence of the restriction imposed on the social network by Assumption 1 is that Bayesian learning takes a log-linear form for *all* information structures. That is, if the social network satisfies Assumption 1, agents' belief dynamics are given by

$$\log \frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \log \frac{\ell_{it+1}^{\theta}(\omega_{it+1})}{\ell_{it+1}^{\hat{\theta}}(\omega_{it+1})} + \sum_{\tau=1}^t \sum_{j \in N_i} a_{it+1,j\tau} \log \frac{\mu_{j\tau}(\theta)}{\mu_{j\tau}(\hat{\theta})}, \quad (7)$$

with weights $a_{it+1,j\tau}$ that are independent of agents' signal structures. We note that Bayesian learning takes a similar representation if instead one restricts the distributions from which agents' private signals are drawn to be normal while allowing for a general social network. The key difference between the two cases is that whereas under Assumption 1 the weights in (7) are independent of the distribution of private signals, the corresponding weights when the signals are normally distributed may depend on the agents' signal structures.¹² Still, without any restriction on either the signal structure or the social network, Bayesian belief dynamic may take a considerably more complex form that does not lend itself to a log-linear representation in the form of (7).¹³

4 Information Aggregation

In this section, we study whether the social network can serve as a mechanism for the propagation and aggregation of information dispersed among the agents when they rely on learning rules that satisfy imperfect recall. As a first step, we determine the conditions under which the log-linear and DeGroot learning rules characterized in Theorems 1 and 2 lead to the long-run learning of the underlying state. We then develop a taxonomy of results on agents' asymptotic beliefs in a more general setting.

Throughout this section, we restrict our attention to environments in which agents' signal structures do not change with time; that is, $\ell_{it}^{\theta} = \ell_i^{\theta}$ for all states θ and all time t . But more importantly, we require that for any pair of states $\theta \neq \hat{\theta}$ there exists an agent i such that $\ell_i^{\theta} \neq \ell_i^{\hat{\theta}}$. This assumption guarantees that agents' observations are *collectively* informative about the underlying state, even if they face identification problems in isolation.

¹²In particular, if agents have normal priors and all signals are normally distributed, beliefs remain normal at all times irrespective of the social network structure (Mossel, Olshman, and Tamuz, 2016). Consequently, agents' belief dynamics can be cast in the form of (7), but with weights that depend on signal precisions. See Online Appendix D for details.

¹³For a discussion of the complexity of Bayesian learning, see Jadbabaie, Mossel, and Rahimian (2017).

We also define the following property of agents' learning rules, which, as our subsequent results illustrate, plays a crucial role in the successful aggregation of agents' information in the long run:

Definition 1. Social learning rule $f_{it} : \Delta\Theta^n \rightarrow \Delta\Theta$ satisfying imperfect recall is *unanimous* if $f_{it}(\mu, \dots, \mu) = \mu$ for all $\mu \in \Delta\Theta$.

Under a unanimous learning rule f_{it} , agent i adopts the common belief of her neighbors whenever they all agree with one another. Under the log-linear or DeGroot learning rules, unanimity of f_{it} is equivalent to assuming that the weights in (3) and (5), respectively, satisfy $\sum_{j \in N_i} a_{ijt} = 1$.¹⁴

4.1 Log-Linear Learning

We start our analysis by studying the long-run implications of the log-linear learning rule (3), which serves as the benchmark rule for the rest of our results. Recall from our earlier discussion that each agent i may face an identification problem in isolation (for instance, when $\ell_i^\theta = \ell_i^{\hat{\theta}}$ for a pair of states $\theta \neq \hat{\theta}$). In such an environment, agent i can successfully learn the state only if she has access to a continuous flow of information from agents with no such identification problems. This suggests that the long-run aggregation of information hinges on how weights a_{ijt} used by the agents to incorporate their neighbors' reports into their beliefs vary over time.

We discipline the long-run behavior of these weights by assuming that there exist a sequence $\lambda_t \in (0, 1)$ and constants $\underline{a}, \bar{a} \in (0, 1)$ such that

$$a_{ijt} \geq \lambda_t \underline{a}, \quad (8)$$

$$\sum_{k \neq i} a_{ikt} \leq \lambda_t \bar{a} \quad (9)$$

for all t and all pairs of agents $i \neq j$ such that $j \in N_i$. By imposing upper and lower bounds on a_{ijt} , the above inequalities guarantee that if the weights that agents assign to their neighbors' beliefs decay to zero, they do so at a common rate λ_t . Such a restriction enables us to express our results in terms of a simple sequence that governs the long-run behavior of agents' learning rules. Note that inequalities (8) and (9) are trivially satisfied (with a constant sequence $\lambda_t = \lambda$) whenever weights a_{ijt} remain uniformly bounded away from zero for all $j \in N_i$ and all t . Nonetheless, with some abuse of terminology, we refer to λ_t as the rate of decay of weights a_{ijt} . We can now state the following result:

Theorem 4. *Suppose agents follow the log-linear learning rule (3) with weights that decay at rate λ_t . If learning rules are unanimous and $\lim_{t \rightarrow \infty} t\lambda_t = \infty$, then all agents learn the state almost surely.*

This result thus illustrates that, under log-linear learning, the information dispersed throughout the social network is fully aggregated as long as (i) agents rely on unanimous social learning rules and (ii) the rate of decay of the weights they assign to their neighbors' beliefs is slower than $1/t$. Learning is complete despite the fact that individuals may face identification problems in isolation, do not make any deductions about how their neighbors obtained their opinions, do not account for potential

¹⁴In the case of the DeGroot learning rule in (5), since $\sum_{\theta \in \Theta} f_{it}(\mu_i^t)(\theta) = 1$, the restriction $\sum_{j \in N_i} a_{ijt} = 1$ also implies that $c_{it} = 0$.

redundancies in different information sources, and may be unaware of the intricate details of the social network.

To see the intuition underlying Theorem 4, it is instructive to first consider an environment in which agents do not face identification problems individually and compare their beliefs to a Bayesian agent with access to all the realized signals throughout the social network. At any given time, such a Bayesian agent assigns an equal weight to any independent piece of information by simply adding up the log-likelihood ratios corresponding to the signals. Such an updating is clearly impossible when individuals cannot observe each others' private signals and have to rely on the non-Bayesian learning rule in (4) to update their beliefs. In particular, under (4), the log-likelihood ratio of an agent's belief at any given time is a *weighted* sum of the log-likelihood ratios corresponding to the signals that reach her up to that time, either via direct observation or through the reports of her neighbors. Nonetheless, unanimity guarantees that the effective weights that each agent assigns to various independent pieces of information satisfy two key long-run properties. First, it ensures that the sum of the weights assigned to independent pieces of information diverges as $t \rightarrow \infty$. Second, it implies that the share of this total weight that goes to any single piece of information converges to zero. While the first property guarantees that each agent keeps accumulating more and more information over time, the second property ensures that no specific piece of information is assigned an outsized weight in the long run. Together, these properties guarantee that all agents eventually uncover the state, despite imperfect recall and partial observability of signals.¹⁵

In addition to unanimity, Theorem 4 also imposes a restriction on how fast the weights a_{ijt} in (3) can decay to zero. The need for such a requirement is fairly straightforward: individuals who face identification problems can learn the underlying state only by relying on the information provided by their neighbors, yet a rapid dismissal of their neighbors' reports would mean that such individuals would not accumulate enough information to resolve their identification problems. The following partial converse to Theorem 4 formalizes this intuition:

Proposition 1. *Suppose agents follow unanimous log-linear learning rules with weights that decay at rate $\lambda_t = 1/t^\alpha$ for some $\alpha > 1 + 1/\delta$, where δ denotes the diameter of the social network. There exist signal structures under which a subset of agents remain uncertain forever almost surely.*

Thus, when the weights that agents assign to their neighbors' reports decay to zero at a rate that is faster than $1/t^{(1+1/\delta)}$, information aggregation may fail in the long run despite the fact that agents have access to enough information to collectively uncover the underlying state. Such a failure is due to the fact that a too rapid dismissal of other agents' reports hinders the flow of information from individuals with informative signals to those who face identification problems. Importantly, Proposition 1 illustrates that the critical rate of decay beyond which learning may fail is closely tied to how far information may have to travel from one agent to another, formally captured via the social network's diameter.

¹⁵The key role played by unanimity in asymptotic learning is not specific to the log-linear learning rule (3). Rather, as we show in our subsequent results, the predictions of Theorem 4 generalize to a larger class of learning rules. We also discuss how the absence of unanimity may result in mislearning or non-learning in the long run. Also see Online Appendix C for a generalization of Theorem 4 to a collection of learning rules f_{it} that, even though not unanimous for all t , are asymptotically unanimous.

Proposition 1 also demonstrates the sharpness of the assumption imposed in Theorem 4 on how fast the weights that agents assign to their neighbors' beliefs can decay to zero. Recall from Theorem 4 that learning is successful for any network and information structure as long as the weights that agents assign to their neighbors' beliefs decay to zero at a rate that is slower than $1/t$. Proposition 1 illustrates that the rate $1/t$ in this statement cannot be replaced by $1/t^{1+\epsilon}$ for any $\epsilon > 0$: given any $\epsilon > 0$, there exist a social network (with a large enough diameter) and signal structures given which learning fails when the weights decay to zero at rate $1/t^{1+\epsilon}$.

We conclude this discussion by noting that the requirement $\lim_{t \rightarrow \infty} t\lambda_t = \infty$ in Theorem 4 is automatically satisfied whenever agents use time-invariant learning rules (that is, when $a_{ijt} = a_{ij}$ for all $j \in N_i$ and all t), in which case unanimity is all that is required for the long-run aggregation of information. We thus have the following corollary to Theorem 4.

Corollary 1. *Suppose agents follow the log-linear learning rule (3) with weights that are time-invariant. If learning rules are unanimous, then all agents learn the state almost surely.*

4.2 DeGroot Learning

We next turn to studying the long-run implications of the DeGroot learning rule characterized in Theorem 2. Recall from the previous section that replacing IIA with separability results in a learning rule according to which agent i linearly combines the beliefs of her neighbors $j \in N_i$ with weights $a_{ijt} > 0$. Our next result, which generalizes the main result of Jadbabaie et al. (2012) and proves a conjecture of Liu et al. (2014), characterizes the conditions under which such a learning rule leads to the long-run aggregation of information.

Theorem 5. *Suppose agents follow the DeGroot learning rule (5) with weights that decay at rate λ_t . If learning rules are unanimous and $\lim_{t \rightarrow \infty} t\lambda_t = \infty$, then all agents learn the state almost surely.*

Contrasting this result with Theorem 4 highlights that, despite their different behavioral foundations which may lead to different sets of beliefs at any given finite time, the DeGroot and log-linear learning rules result in asymptotic learning under analogous sets of conditions. In particular, aside from the restriction on the rate of decay λ_t , the unanimity of agents' learning rules serves as a sufficient condition for the successful aggregation of information. In the next subsection, we show that the long-run convergence of beliefs for agents who follow these two learning rules is no coincidence and is a more general phenomenon.

The above result also illustrates that, as in Theorem 4 for the log-linear learning rule, asymptotic learning is independent of the specific values of the weights that each individual assigns to her neighbors' beliefs: as long as weights a_{ijt} in (5) do not decay too rapidly, the continuous flow of information over the network guarantees that all agents will eventually uncover the underlying state.

As a final remark, we note that if agents' learning rules do not change over time, Theorem 5 reduces to the following counterpart to Corollary 1:

Corollary 2. *Suppose agents follow the DeGroot learning rule (5) with weights that are time-invariant. If learning rules are unanimous, then all agents learn the state almost surely.*

4.3 Beyond Functional Forms

In the remainder of this section, we present a set of long-run learning results that generalize our earlier results and illustrate the limits to the successful aggregation of information in the presence of imperfect recall. We achieve this by replacing IIA and separability with a weaker notion of independence and obtaining a general class of learning rules that encompasses the log-linear and DeGroot rules as special cases. To simplify the analysis and emphasize the features that are unrelated to whether information can flow from one agent to others, we restrict our attention to environments in which agents' learning rules are time-invariant, in which case we can drop the time index t .

Definition 2. Agent i 's social learning rule satisfying IR is *weakly separable* if there exists a smooth, homogeneous, and increasing function $\psi_i : [0, 1]^n \rightarrow \mathbb{R}_+$ such that

$$\frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} = \frac{\psi_i(\mu(\theta))}{\psi_i(\mu(\hat{\theta}))} \quad (10)$$

for all belief profiles $\mu \in \Delta\Theta^n$ and all $\theta, \hat{\theta} \in \Theta$.

In other words, in determining the relative likelihoods of two given states, agent i relies solely on the ratio of the aggregates of her neighbors' beliefs about those two states, where beliefs about each state are aggregated using the same homogeneous function ψ_i . Thus, any learning rule in this class has a representation of the form

$$f_i(\mu)(\theta) = \frac{\psi_i(\mu(\theta))}{\sum_{\hat{\theta} \in \Theta} \psi_i(\mu(\hat{\theta}))} \quad \text{for all } \theta \in \Theta, \quad (11)$$

with the log-linear and DeGroot learning rules as two special cases. In fact, our representation Theorems 1 and 2 establish that time-invariant log-linear and DeGroot learning rules belong to the class of weakly-separable learning rules with $\psi_i(x) = \prod_{j \in N_i} x_j^{a_{ij}}$ and $\psi_i(x) = \sum_{j \in N_i} a_{ij} x_j$, respectively. The unanimity of a weakly-separable learning rule f_i is in turn determined by the degree of homogeneity of the corresponding ψ_i : the former is unanimous if and only if the degree of homogeneity of the latter is equal to 1.

Weak separability therefore imposes a less stringent requirement on agents' social learning rules than either IIA or separability; yet the weak-separability assumption has enough bite to allow us to provide a taxonomy of how different features of agents' social learning rules determine their long-run beliefs. In our next result, we focus on the role of the degree of homogeneity of ψ_i in the aggregation of information by showing that the violation of the unanimity assumption imposed in the previous subsections can result in asymptotic mislearning or non-learning.

Theorem 6. *Suppose agents' social learning rules belong to the class of weakly-separable learning rules.*

- (a) *If ψ_i 's are homogeneous of degree $\rho > 1$, all agents mislearn the state with positive probability.*
- (b) *If ψ_i 's are homogeneous of degree $\rho < 1$, all agents remain uncertain forever with probability one.*

Statement (a) establishes that when individuals rely on weakly-separable learning rules with corresponding ψ_i 's that are homogeneous of a degree greater than 1, they may assign probability one

to a false state as $t \rightarrow \infty$ in spite of having access to enough information to (collectively) uncover the underlying state. This result is a consequence of the fact that agents end up assigning progressively larger weights to their earlier observations at the expense of the signals they receive later on. Thus, they may mislearn the state by herding on misleading signals they observe early on. In contrast, statement (b) of the theorem establishes that when the degree of homogeneity of ψ_i is strictly smaller than 1, agents never reach certainty, even in the long run. This is due to the fact that when $\rho < 1$, agents downplay their past observations in favor of their most recent ones and as a result never accumulate enough information to uncover the underlying state.

Taken together, the two parts of Theorem 6 illustrate that the long-run aggregation of information fails with positive probability as long as $\rho \neq 1$, as agents do not properly combine their most recent signals with the information they have already accumulated. This result thus illustrates the importance of the unanimity assumption, imposed in Theorems 4 and 5 and Corollaries 1 and 2, for asymptotic learning.

In what follows, we study the dynamics of beliefs when agents employ weakly-separable learning rules with $\rho = 1$. The outcome of social learning in such an environment turns out to depend on a measure of the curvatures of agents' social learning rules. We define the *logarithmic curvature* of agent i 's social learning rule (in the direction of j and l) as

$$\kappa_i^{(jl)}(x) = - \left(\frac{\partial^2 \log \psi_i(x)}{\partial \log x_j \partial \log x_l} \right) / \left(\frac{\partial \log \psi_i(x)}{\partial \log x_j} \frac{\partial \log \psi_i(x)}{\partial \log x_l} \right), \quad (12)$$

where $j \neq l$ are agents in i 's neighborhood. This quantity measures the extent to which i 's learning rule departs from the benchmark of log-linear learning. In particular, it is easy to verify that $\kappa_i^{(jl)}(x) = 0$ for all $j, l \in N_i$ and all $x \in [0, 1]^n$ whenever agent i follows the log-linear learning rule in (3). As κ_i deviates from 0 in either direction, the functional form used by agent i to aggregate her neighbors' beliefs moves further away from the log-linear benchmark. The following example further clarifies this point:

Example 1. Suppose agent i employs a weakly-separable learning rule whose corresponding ψ_i is homogeneous of degree $\rho = 1$ and takes a CES functional form given by

$$\psi_i(x) = \left[\sum_{j \in N_i} a_{ij} x_j^\xi \right]^{1/\xi}, \quad (13)$$

where a_{ij} 's are some positive constants and $\xi \in \mathbb{R}$. Note that the (time-invariant) log-linear and DeGroot models belong to this class of learning rules, with $\xi \rightarrow 0$ and $\xi = 1$, respectively. It is also easy to verify that the logarithmic curvature of agent i 's learning rule is $\kappa_i^{(jl)}(x) = \xi$ for all pairs of agents $j, l \in N_i$ and all $x \in [0, 1]^n$. Therefore, any learning rule in this class for which $\xi > 0$ (such as the DeGroot learning rule) exhibits more logarithmic curvature compared to the benchmark of log-linear learning. In contrast, learning rules with a negative ξ exhibit less logarithmic curvature compared to the log-linear learning rule. One important special case is the case of the *harmonic learning rule* with $\xi = -1$. Under such a learning rule, agent i sets her posterior likelihood ratio of two states as the ratio of harmonic means of her neighbors' beliefs on those states.

With the above notion in hand, we can state the following result:

Theorem 7. *Suppose agents' social learning rules belong to the class of weakly-separable learning rules and are unanimous.*

- (a) *All agents learn the underlying state almost surely if the logarithmic curvatures of their learning rules fall within the interval $[-1, 1]$.*
- (b) *There exist social networks, signal structures, and learning rules with logarithmic curvatures less than -1 for which all agents mislearn the state with strictly positive probability.*
- (c) *There exist social networks, signal structures, and learning rules with logarithmic curvatures greater than 1 for which agents learn the underlying state almost never.*

Part (a) of this theorem illustrates that for agents to learn the underlying state, it is sufficient that their learning rules satisfy two key conditions. First, unanimity ensures that the effective weights that any given agent assigns to each independent piece of information are of the same order of magnitude, thus guaranteeing that information from early periods is neither discarded nor given an outsized significance as $t \rightarrow \infty$. Second, the restriction on the learning rules' logarithmic curvatures ensures that the learning rules' functional forms do not deviate significantly from the benchmark of log-linear learning (which exhibits a logarithmic curvature equal to zero throughout its domain).

The latter two parts of Theorem 7 establish that the restriction on the learning rules' logarithmic curvature cannot be dispensed with, thus illustrating that unanimity is not, in and of itself, sufficient for the long-run aggregation of information. More specifically, they illustrate that agents may not learn the state or may mislearn it if the logarithmic curvature falls outside of the $[-1, 1]$ interval. Parts (b) and (c) of the theorem also highlight that the DeGroot and harmonic learning rules correspond to the two extremes in the class of weakly-separable learning rules for which long-run aggregation of information is guaranteed to be successful.

The intuition for the role of curvature in asymptotic learning can be best understood by considering the special case of a symmetric environment with two states and focusing on the dynamics of agents' interim beliefs $\zeta_{it} = f_i(\mu_t)$ — i.e., beliefs that are formed after observing other agents' reports but before observing one's own private signal. In particular, suppose that agents interact over the complete social network depicted in the left panel of Figure 1 and rely on identical unanimous weakly-separable learning rules. Assuming that agent j is the only agent with informative signals, the dynamics of agent i 's (interim) beliefs can be expressed as

$$\log \frac{\zeta_{it+1}(\theta)}{\zeta_{it+1}(\hat{\theta})} = \varphi \left(\log \frac{\ell_j^\theta(\omega_{jt+1})}{\ell_j^{\hat{\theta}}(\omega_{jt+1})}; \zeta_{it}(\theta) \right) + \log \frac{\zeta_{it}(\theta)}{\zeta_{it}(\hat{\theta})},$$

where φ is a transformation of ψ . In particular, the former is a non-linear function of its first argument whenever the latter has a non-zero logarithmic curvature.¹⁶ Contrasting the above equation with the dynamics of log-linear learning in (4) clarifies how the curvature of ψ (acting via that of φ) can

¹⁶See the proof of Theorem 7 for details. Also note that due to the symmetric nature of the environment, the interim beliefs of all agents coincide at all times; that is, $\zeta_{it} = \zeta_{jt}$ for all i and j .

function as an impediment to learning: when ψ exhibits large (positive or negative) logarithmic curvature, it may be the case that the expected value of $\varphi(\log(\ell_j^\theta(\omega_{jt})/\ell_j^{\hat{\theta}}(\omega_{jt})))$ is negative despite the fact that the expected value of $\log(\ell_j^\theta(\omega_{jt})/\ell_j^{\hat{\theta}}(\omega_{jt}))$ is necessarily positive when the realized state is θ . Consequently, significant departures from the benchmark of log-linear learning can transform a drift towards learning into a drift in the opposite direction.

To summarize, Theorems 6 and 7 provide a fairly complete picture of the forces that underpin learning, non-learning, and mislearning in the class of weakly-separable learning rules. In particular, they clearly indicate that the long-run aggregation of information under the log-linear and DeGroot learning rules (established in Corollaries 1 and 2, respectively) is not due to these rules' specific functional forms. Rather, what matters for asymptotic learning is that (i) agents' learning rules satisfy some weak notion of independence across different states (as captured by (10)), (ii) the learning rules are unanimous, and (iii) the logarithmic curvatures of agents' learning rules are confined to the interval $[-1, 1]$. Crucially, our results also establish that if either the unanimity or curvature conditions are violated, agents may remain uncertain forever about the underlying state or assign probability one to a false state as $t \rightarrow \infty$.

5 Rate of Information Aggregation

Our results in Section 4 illustrate the conditions that lead the agents to eventually uncover the underlying state. These results, however, are silent about the precision of agents' beliefs away from the limit. In this section, we provide a refinement of our learning theorems and characterize the rate at which information is aggregated throughout the society. For concreteness, we restrict our attention to the benchmark rule of log-linear learning characterized in (3). In view of Theorem 4, we also assume that agents' learning rules are unanimous at all times, with corresponding weights that decay to zero at a rate that is slower than $1/t$. We explore the implications of relaxing the latter assumption at the end of this section.

5.1 Information and Centrality

We start by defining a measure for the expected informativeness of each agent's private signals. For any given pair of states $\theta, \hat{\theta} \in \Theta$, let

$$h_i(\theta, \hat{\theta}) = \sum_{\omega \in S} \ell_i^\theta(\omega) \log \frac{\ell_i^\theta(\omega)}{\ell_i^{\hat{\theta}}(\omega)}$$

denote the *relative entropy* of θ with respect to $\hat{\theta}$ in agent i 's signal structure. Loosely speaking, $h_i(\theta, \hat{\theta})$, which is always non-negative, measures the expected informativeness of agent i 's private signals about states θ and $\hat{\theta}$ when the underlying state is θ . In particular, if $h_i(\theta, \hat{\theta})$ is strictly positive, observing a sufficiently long sequence of signals generated by ℓ_i^θ enables i to rule out $\hat{\theta}$ with an arbitrarily large confidence. In contrast, if $h_i(\theta, \hat{\theta}) = 0$, agent i 's private observations provide her with no information that is conducive to disentangling $\hat{\theta}$ from θ .

While relative entropy serves as a measure for the informativeness of agents' signals at any given period, the fact that information has to be relayed over the social network means that the collection

$\{h_i(\theta, \hat{\theta}) : (\theta, \hat{\theta}) \in \Theta^2, i \in N\}$ may not be a sufficient statistic for the speed at which agents learn the underlying state. Rather, the speed of learning may also depend on the extent to which information travels from one agent to another.

To account for the differential roles of various agents in relaying information, we define a novel notion of network centrality that generalizes the well-known concept of eigenvector centrality to a sequence of (potentially time-varying) unanimous learning rules:

Definition 3. The *Kolmogorov centrality* of agent i at time t is

$$v_{it} = \sum_{j=1}^n a_{jit} v_{jt+1}, \quad (14)$$

where a_{jit} is the weight that agent j assigns to agent i 's belief at time t and $\sum_{i=1}^n v_{it} = 1$ for all t .

The Kolmogorov centrality, which is defined recursively over both time and space, is a measure of an agent's importance as a source of information: agent i is more central in the social network at time t if agents who are more central at time $t + 1$ assign a higher weight to her time- t belief.¹⁷ In this sense, v_{it} captures the extent to which agent i 's time- t information propagates throughout the social network over time. Furthermore, as is evident from the recursive definition in (14), the Kolmogorov centrality of an agent at any given time t depends not only on the weights that individuals assign to their neighbors at that time but also on the weights used in all subsequent time periods $\tau > t$.

We note that Kolmogorov centrality reduces to the well-known concept of eigenvector centrality when agents rely on learning rules that do not change over time. In particular, when the weights in learning rule (4) are time invariant, that is, $a_{ijt} = a_{ij}$ for all t , equation (14) reduces to $v_i = \sum_{j=1}^n a_{ji} v_j$, the solution to which coincides with the agents' eigenvector centralities.¹⁸

Proposition 2. *Suppose agents' learning rules are unanimous with corresponding weights that decay at rate λ_t . If $\lim_{t \rightarrow \infty} t \lambda_t = \infty$, there exists a unique collection of Kolmogorov centralities v_{it} satisfying (14). Furthermore, $\liminf_{t \rightarrow \infty} v_{it} > 0$ for all i .*

In other words, as long as the weights that agents assign to their neighbors' beliefs do not decay too rapidly over time, Kolmogorov centrality is a well-defined concept, is uniquely determined for any given agent at all times, and remains bounded away from zero even as $t \rightarrow \infty$.

5.2 Learning Rate

With the above notions of signal informativeness and network centrality in hand, we now proceed to study the rate of information aggregation.

Let $e_{it}^\theta = \sum_{\hat{\theta} \neq \theta} \mu_{it}(\hat{\theta})$ denote the belief that agent i assigns at time t to states other than the underlying state θ . Recall from Theorem 4 that, for any realization θ of the state, $e_{it}^\theta \rightarrow 0$ almost

¹⁷Our notion of Kolmogorov centrality is closely related to what is known as a sequence of *absolute probability vectors* corresponding to a sequence of stochastic matrices, a notion that first appeared in [Kolmogorov \(1936\)](#).

¹⁸As in the definition of eigenvector centrality, the requirement $\sum_{i=1}^n v_{it} = 1$ for all t in Definition 3 is a simple normalization, as equation (14) only defines the centralities up to a constant.

surely if the learning rules are unanimous and weights a_{ijt} decay to zero at a rate that is slower than $1/t$. Define

$$\gamma_i^\theta = \liminf_{t \rightarrow \infty} \frac{1}{t} |\log e_{it}^\theta|.$$

This quantity, which is inversely proportional to the number of time periods it takes for agent i 's beliefs on the false states to fall below some given threshold, is finite and positive if agent i learns the underlying state exponentially fast. We have the following result:

Theorem 8. *Let θ denote the underlying state, and suppose agents follow the log-linear learning rule (3) with weights that decay at rate λ_t . If learning rules are unanimous and $\lim_{t \rightarrow \infty} t\lambda_t = \infty$, then*

$$\gamma_i^\theta = \min_{\hat{\theta} \neq \theta} \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{j=1}^n v_{j\tau} h_j(\theta, \hat{\theta}) \quad (15)$$

almost surely, where $v_{j\tau}$ is agent j 's Kolmogorov centrality at time τ .

As a first observation, note that the above characterization guarantees that γ_i^θ is non-zero and finite, thus establishing that agent i learns the underlying state exponentially fast. This is a consequence of the facts that agents' centralities remain bounded away from zero (shown in Proposition 2) and that there exists an agent j who can distinguish $\hat{\theta}$ from θ in the long run. The significance of Theorem 8, however, lies in characterizing the dependence of the rate of learning on the information available to agents — as measured by the relative entropies of their signals — and how this information is distributed among them — as summarized by their centralities.

The characterization in (15) has an intuitive interpretation. The relative entropy of each agent's signal structure is weighted by her Kolmogorov centralities. This captures the fact that the informativeness of each agent's signal structure impacts the rate of learning via the effective (direct and indirect) attention she receives from other agents in the social network in all future periods. On the other hand, the minimization over $\hat{\theta} \neq \theta$ illustrates that the speed of learning is determined by the rate at which agents rule out the state $\hat{\theta}$ that is hardest to distinguish from the realized state θ . This is due to the fact that learning is complete only if agents can rule out all incorrect states.

5.3 Sub-Exponential Learning

We conclude this section by illustrating how relaxing the restriction on the rate of decay of the weights that agents assign to their neighbors' reports (captured via the sequence λ_t) may result in learning at a sub-exponential rate.

Recall from our discussion above that Theorems 4 and 8 imply that if λ_t decays to zero at a rate that is slower than $1/t$, agents not only learn the state but also do so exponentially fast. At the same time, in Proposition 1 we established that when λ_t decays to zero at a rate that is faster than $1/t^{1+1/\delta}$, information may never get aggregated in a social network with diameter δ . Our next result explores the gap between these two critical rates. For tractability, we focus our attention on the complete social network, in which each agent has access to the beliefs of all other agents in the society and follows the learning rule (3) with weights $a_{ijt} = \lambda_t/n$ for all $i \neq j$, where $\lambda_t = 1/t^\alpha$ and α is some non-negative

constant. Note that, in view of our previous results, it is sufficient to consider the case of $\alpha \in (1, 2)$, as the complete social network has diameter $\delta = 1$.

Proposition 3. *Suppose agents use unanimous log-linear learning rules and that $a_{ijt} = \lambda_t/n$ for all $i \neq j$, where $\lambda_t = 1/t^\alpha$ and $\alpha \in (1, 2)$. If agent i faces an identification problem for the realized state θ , then $\liminf_{t \rightarrow \infty} \frac{1}{t^{2-\alpha}} |\log e_{it}^\theta| = c_\alpha$ almost surely for some positive constant c_α .*

This result thus establishes that when the weights that agents assign to their neighbors' reports decay to zero at a rate that is faster than $1/t$ but slower than $1/t^2$, information is eventually aggregated but at a significantly slower pace compared to the case in which $t\lambda_t \rightarrow \infty$. The above result also indicates that the belief that agent i assigns to the false states converges to zero more slowly, the faster agents discount their neighbors' beliefs over time. In particular, learning occurs at an increasingly slow pace as α increases from 1 to 2 and fails altogether once α exceeds the critical value of $1 + 1/\delta = 2$ (as established in Proposition 1).

A Proofs

Notation. Throughout the proofs, we use \mathbb{P} to denote the joint probability distribution over the set of states and the set of all signals observed by agents in all time periods. We use $\mathbb{P}^\theta(\{\omega_{it}\}_{t \geq 1, i \in N})$ to denote the probability of observing $\{\omega_{it}\}_{t \geq 1, i \in N}$ given the event that the realized state is θ and use \mathbb{E}^θ to denote the expectation operator corresponding to \mathbb{P}^θ . We write \mathbb{E}_t^θ as shorthand for $\mathbb{E}^\theta[\cdot | \mathcal{F}_t]$, where \mathcal{F}_t denotes the σ -field generated by $(\{\omega_{i\tau}\}_{1 \leq \tau \leq t, i \in N})$. Finally, we use $A_t = [a_{ijt}]$ to denote the matrix of weights that agents assign to their neighbors' log-likelihood ratios (as in learning rule (3)) or beliefs (as in learning rule (5)) at time t , with the convention that $a_{ijt} = 0$ if $j \notin N_i$.

Lemma A.1. *Let A_t denote a sequence of stochastic weight matrices, satisfying inequalities (8) and (9), where $\lim_{t \rightarrow \infty} t\lambda_t = \infty$. There exists a unique sequence of probability vectors v_t with a uniform positive lower bound such that*

$$v'_{t+1} A_t A_{t-1} \dots A_\tau = v'_\tau \quad (16)$$

for all pairs of integers $t \geq \tau$.

Proof. We first show that there exists a sequence of probability vectors v_τ uniformly bounded away from zero such that

$$\lim_{t \rightarrow \infty} A_t A_{t-1} \dots A_\tau = \mathbf{1} v'_\tau \quad (17)$$

for all τ , where $\mathbf{1}$ denotes the vector of all ones. By assumption, there exist constants $\underline{a}, \bar{a} \in (0, 1)$ and a sequence $\lambda_t \in (0, 1)$ such that $a_{ijt} \geq \lambda_t \underline{a}$ and $\sum_{j \neq i} a_{ijt} \leq \lambda_t \bar{a}$ for all t . Thus, $A_t = (1 - \lambda_t)I + \lambda_t B_t$, where B_t is a stochastic matrix whose non-zero elements are uniformly lower bounded by $\eta = \min\{\underline{a}, 1 - \bar{a}\} \in (0, 1)$, i.e., $b_{ijt} \geq \eta$ for all t and all i and j such that $j \in N_i$.

Following DeMarzo et al. (2003), let $\{\Lambda_t\}$ denote a sequence of independent Bernoulli random variables with Λ_t taking value 1 with probability λ_t and zero otherwise. Consequently, $A_t = \mathbb{E}^*[(1 -$

$\Lambda_t)I + \Lambda_t B_t]$, where \mathbb{E}^* denotes the expectation operator corresponding to the product measure on the sequence of random variables Λ_t . By the dominated convergence theorem,

$$\lim_{t \rightarrow \infty} A_t A_{t-1} \dots A_\tau = \mathbb{E}^* \left[\lim_{t \rightarrow \infty} \prod_{r=\tau}^t ((1 - \Lambda_r)I + \Lambda_r B_r) \right] = \mathbb{E}^* \left[\prod_{\substack{t: \Lambda_t=1 \\ t \geq \tau}} B_t \right]. \quad (18)$$

On the other hand, the assumption that $\lim_{t \rightarrow \infty} t\lambda_t = \infty$ guarantees that $\sum_{t=\tau}^{\infty} \lambda_t = \infty$. Hence, by the Borel-Cantelli lemma, $\Lambda_t = 1$ infinitely often with probability one, which implies that $\prod_{\substack{t: \Lambda_t=1 \\ t \geq \tau}} B_t$ is almost surely an infinite product of irreducible matrices whose non-zero elements are uniformly lower bounded by η .¹⁹ Thus, by Theorem 4.19 of [Seneta \(1981\)](#), the matrix sequence $\{B_t\}_{\Lambda_t=1, t \geq \tau}$ is strongly ergodic with probability one. Hence, there exists a random probability vector w_τ such that,

$$\prod_{\substack{t: \Lambda_t=1 \\ t \geq \tau}} B_t = \mathbf{1} w_\tau'. \quad (19)$$

Combining the above with (18) results in (17), with $v_\tau = \mathbb{E}^*[w_\tau]$. To show that the vectors in the sequence v_τ are uniformly bounded away from zero, recall that the matrices in the sequence $\{B_t\}$ are all irreducible, with non-zero elements that are uniformly lower bounded by η for all t . Therefore, any product of length n of these matrices is element-wise strictly positive, with all elements lower bounded by constant η^n . This observation, coupled with the fact that B_t is stochastic, implies that the infinite product in (19) is lower bounded by matrix $\eta^n \mathbf{1}\mathbf{1}'$, thus guaranteeing that $v_{i\tau} = \mathbb{E}^*[w_{i\tau}] > \eta^n$.

With (17) in hand, Theorem 4.20 of [Seneta \(1981\)](#) implies that the sequence of vectors v_τ in (17) are the unique sequence of probability vectors that satisfy (16) for all pairs of integers $t \geq \tau$. \square

Lemma A.2. For a stochastic matrix A , define $\pi(A) = 1 - \min_{i,j} \sum_{k=1}^n \min\{a_{ik}, a_{jk}\}$.

(a) $\pi(A) \in [0, 1]$ and is convex over its domain.

(b) $\pi(A) \leq 1 - \max_j \min_i a_{ij}$.

(c) π is submultiplicative, i.e., $\pi(A_2 A_1) \leq \pi(A_2) \pi(A_1)$ for any pair of stochastic matrices A_1 and A_2 .

(d) Let w denote an arbitrary vector and A denote a stochastic matrix. If $z = Aw$, then $\max_i z_i - \min_i z_i \leq \pi(A)(\max_i w_i - \min_i w_i)$.

Proof. Statements (a) and (b) are immediate to verify. See Lemma 4.3 of [Seneta \(1981\)](#) for a proof of statement (c) and Theorem 3.1 of [Seneta \(1981\)](#) for a proof of statement (d). \square

Proof of Theorem 1

Consider two arbitrary states $\theta \neq \hat{\theta}$ and an arbitrary profile of beliefs $\mu \in \Delta^{\theta^n}$. Let $\bar{\Theta} = \{\theta, \hat{\theta}\}$. By definition of conditional probability,

$$\log \frac{f_{it}(\mu)(\theta)}{f_{it}(\mu)(\hat{\theta})} = \log \text{cond}_{\bar{\Theta}}(f_{it}(\mu))(\theta) - \log \text{cond}_{\bar{\Theta}}(f_{it}(\mu))(\hat{\theta}),$$

¹⁹The irreducibility of B_t follows from the strong connectivity of the social network and the monotonicity assumption.

where recall that as a consequence of imperfect recall we can restrict the domain of f_{it} to $\Delta\Theta^n$. On the other hand, IIA implies that

$$\log \frac{f_{it}(\mu)(\theta)}{f_{it}(\mu)(\hat{\theta})} = \log f_{it}(\text{cond}_{\Theta}(\mu))(\theta) - \log f_{it}(\text{cond}_{\Theta}(\mu))(\hat{\theta}).$$

Note that $\text{cond}_{\Theta}(\mu)$ depends on the belief profile μ only through the collection of likelihood ratios $\{\mu_j(\theta)/\mu_j(\hat{\theta})\}_{j=1}^n$. Consequently, for any given agent i , there exists a continuous function $g_{it} : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$\log \frac{f_{it}(\mu)(\theta)}{f_{it}(\mu)(\hat{\theta})} = g_{it} \left(\log \frac{\mu_1(\theta)}{\mu_1(\hat{\theta})}, \dots, \log \frac{\mu_n(\theta)}{\mu_n(\hat{\theta})} \right) \quad (20)$$

for all pairs of states $\theta \neq \hat{\theta}$ and all profiles of beliefs μ . Furthermore, LN guarantees that function g_{it} is independent of θ and $\hat{\theta}$.

Now, consider three distinct states θ , $\hat{\theta}$ and $\tilde{\theta}$. Given that (20) has to be satisfied for any arbitrary pair of states, we have

$$\begin{aligned} g_{it} \left(\log \frac{\mu_1(\theta)}{\mu_1(\hat{\theta})}, \dots, \log \frac{\mu_n(\theta)}{\mu_n(\hat{\theta})} \right) + g_{it} \left(\log \frac{\mu_1(\hat{\theta})}{\mu_1(\tilde{\theta})}, \dots, \log \frac{\mu_n(\hat{\theta})}{\mu_n(\tilde{\theta})} \right) &= \log \frac{f_{it}(\mu)(\theta)}{f_{it}(\mu)(\hat{\theta})} + \log \frac{f_{it}(\mu)(\hat{\theta})}{f_{it}(\mu)(\tilde{\theta})} \\ &= g_{it} \left(\log \frac{\mu_1(\theta)}{\mu_1(\tilde{\theta})}, \dots, \log \frac{\mu_n(\theta)}{\mu_n(\tilde{\theta})} \right). \end{aligned}$$

Since μ was arbitrary, the above equation implies that for any arbitrary $x, y \in \mathbb{R}^n$, it must be the case that $g_{it}(x) + g_{it}(y) = g_{it}(x + y)$. This equation is nothing but Cauchy's functional equation, with linear functions as its single family of continuous solutions. Therefore, there exist constants a_{ijt} such that $g_{it}(x) = \sum_{j=1}^n a_{ijt}x_j$. Thus, using (20) one more time implies that $\log(f_{it}(\mu)(\theta)/f_{it}(\mu)(\hat{\theta})) = \sum_{j=1}^n a_{ijt} \log(\mu_j(\theta)/\mu_j(\hat{\theta}))$ for all $\theta, \hat{\theta} \in \Theta$. Finally, monotonicity implies that $a_{ijt} > 0$ for all $j \in N_i$. \square

Proof of Theorem 2

Separability and imperfect recall imply that, for any given agent i and any time t , there exists a function $g_{it} : [0, 1]^n \rightarrow [0, 1]$ such that $f_{it}(\mu)(\theta) = g_{it}(\mu_1(\theta), \dots, \mu_n(\theta))$ for all belief profiles $\mu \in \Delta\Theta$ and all states $\theta \in \Theta$. Furthermore, LN guarantees that g_{it} is independent of θ . Fix an arbitrary belief profile $\mu \in \Delta\Theta^n$ and a pair of states $\theta \neq \hat{\theta}$. The above equation implies that

$$g_{it}(\mu(\theta)) + g_{it}(\mu(\hat{\theta})) = f_{it}(\mu)(\theta) + f_{it}(\mu)(\hat{\theta}) = 1 - \sum_{\tilde{\theta} \notin \{\theta, \hat{\theta}\}} f_{it}(\mu)(\tilde{\theta}) = 1 - \sum_{\tilde{\theta} \notin \{\theta, \hat{\theta}\}} g_{it}(\mu(\tilde{\theta})).$$

Note that changing $\mu(\theta)$ and $\mu(\hat{\theta})$ does not impact the right-hand side of the above equality as long as $\mu(\tilde{\theta})$ is kept unchanged for all $\tilde{\theta} \notin \{\theta, \hat{\theta}\}$. Consequently, $g_{it}(\mu(\theta)) + g_{it}(\mu(\hat{\theta})) = g_{it}(\mu(\theta) + \mu(\hat{\theta})) + g_{it}(0)$. Given that $\mu \in \Delta\Theta$ in the previous equation is arbitrary, we have $h_{it}(x) + h_{it}(y) = h_{it}(x + y)$ for any arbitrary $x, y \in \mathbb{R}^n$ such that $x, y \geq 0$ and $x + y \leq 1$ and with the function h_{it} defined as $h_{it}(z) = g_{it}(z) - g_{it}(0)$. Hence, by Cauchy's functional equation, there exist constants a_{ijt} such that $h_{it}(x) = \sum_{j=1}^n a_{ijt}x_j$. Hence, $g_{it}(x) = c_{it} + \sum_{j=1}^n a_{ijt}x_j$ for some constant c_{it} , which in turn implies that $f_{it}(\mu)(\theta) = c_{it} + \sum_{j \in N_i} a_{ijt}\mu_j(\theta)$. Finally, the fact that $f_{it}(\mu)(\theta)$ has to be non-negative for all belief profiles μ ensures that $c_{it} \geq 0$, while monotonicity guarantees that $a_{ijt} > 0$ for all $j \in N_i$ and all t . \square

Proof of Theorem 3

As a first step, we show that agent i 's belief dynamic is given by

$$\log \frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \log \frac{\ell_{it}^{\theta}(\omega_{it+1})}{\ell_{it}^{\hat{\theta}}(\omega_{it+1})} + \sum_{\tau=1}^t \sum_{j \in N_i} a_{it+1,j\tau} \log \frac{\mu_{j\tau}(\theta)}{\mu_{j\tau}(\hat{\theta})} \quad (21)$$

for a collection of constants $a_{it,j\tau}$ that are independent of agents' beliefs and signal structures. Note that establishing the relationship immediately implies parts (a) and (c) of the theorem.

For a given pair of (not necessarily neighboring) agents i and j and pair of time instances $t, \tau \geq 0$, define constants $a_{it,j\tau}$ recursively as

$$a_{it,j\tau} = \begin{cases} -1 & \text{if } i = j \text{ and } t = \tau, \\ -\sum_{r=\tau}^{t-1} \sum_{\substack{k: d(i,k) \leq t-r \\ d(k,j) \leq r-\tau}} a_{kr,j\tau} & \text{otherwise,} \end{cases} \quad (22)$$

where $d(l, k)$ denotes the length of the shortest directed path from vertex k to vertex l over the social network. Note that, by construction, $a_{it,j\tau} = 0$ for all $t < \tau$ and $a_{it,jt} = 0$ for all $i \neq j$, thus indicating that the above expression defines constants $a_{it,j\tau}$ recursively (and hence uniquely).²⁰ Also note that these constants are independent of agents' beliefs and signal structures and are solely dependent on the social network structure.

Our next lemma, which is a consequence of Menger's theorem (McCuaig, 1984), establishes that, under Assumption 1, $a_{it,j\tau}$ is non-zero only if agent i can observe agent j 's beliefs. The proof is provided in Online Appendix B.

Lemma A.3. *Suppose Assumption 1 is satisfied. If $j \notin N_i$, then $a_{it,j\tau} = 0$ for all $t, \tau \geq 0$.*

Lemma A.4. *Suppose Assumption 1 is satisfied and that agents' beliefs follow the dynamics*

$$\log \frac{\nu_{it+1}(\theta)}{\nu_{it+1}(\hat{\theta})} = \log \frac{\ell_{it+1}^{\theta}(\omega_{it+1})}{\ell_{it+1}^{\hat{\theta}}(\omega_{it+1})} + \sum_{\tau=1}^t \sum_{j \in N_i} a_{it+1,j\tau} \log \frac{\nu_{j\tau}(\theta)}{\nu_{j\tau}(\hat{\theta})}, \quad (23)$$

where $a_{it,j\tau}$ are defined as in (22). Then the resulting beliefs satisfy $\nu_{it}(\theta) = \mathbb{P}(\theta | \mathcal{G}_{it})$ for all $i \in N, t > 0$, and $\theta \in \Theta$, where $\mathcal{G}_{it} = \sigma(\{\omega_{j\tau} : d(i, j) \leq t - \tau\})$ is the σ -field generated by all the signals that could have potentially reached agent i up to time t .

Proof. By Lemma A.3, equation (23) can be rewritten as

$$\log \frac{\nu_{it}(\theta)}{\nu_{it}(\hat{\theta})} = \log \frac{\ell_{it}^{\theta}(\omega_{it})}{\ell_{it}^{\hat{\theta}}(\omega_{it})} + \sum_{\tau=1}^{t-1} \sum_{j=1}^n a_{it,j\tau} \log \frac{\nu_{j\tau}(\theta)}{\nu_{j\tau}(\hat{\theta})}.$$

Fix an arbitrary agent k and an arbitrary time instance $T \geq t$. The above equation implies that

$$\sum_{i=1}^n \sum_{t=1}^{T-d(k,i)} \log \frac{\nu_{it}(\theta)}{\nu_{it}(\hat{\theta})} = \sum_{i=1}^n \sum_{t=1}^{T-d(k,i)} \log \frac{\ell_{it}^{\theta}(\omega_{it})}{\ell_{it}^{\hat{\theta}}(\omega_{it})} + \sum_{j=1}^n \sum_{\tau=1}^{T-1} \sum_{t=\tau+1}^T \sum_{\substack{i: d(i,j) \leq t-\tau \\ d(k,i) \leq T-t}} a_{it,j\tau} \log \frac{\nu_{j\tau}(\theta)}{\nu_{j\tau}(\hat{\theta})}, \quad (24)$$

²⁰The constants defined in (22) are closely linked to the Möbius inverse function of summations over partially ordered sets. See Rota (1964) for more details on Möbius inversion and Eyster and Rabin (2014) for an application in the context of observational learning models in which agents take actions sequentially.

where we are using the fact that $a_{it,j\tau} = 0$ if $d(i,j) > t - \tau$. On the other hand, note that

$$\sum_{t=\tau+1}^T \sum_{\substack{i: d(i,j) \leq t-\tau \\ d(k,i) \leq T-t}} a_{it,j\tau} = \sum_{t=\tau}^{T-1} \sum_{\substack{i: d(i,j) \leq t-\tau \\ d(k,i) \leq T-t}} a_{it,j\tau} + a_{kT,j\tau} - a_{j\tau,j\tau} \mathbb{I}_{\{d(k,j) \leq T-\tau\}},$$

where \mathbb{I} denotes the indicator function. By (22), the first two terms on the right-hand side of the above expression cancel out, whereas $a_{j\tau,j\tau} = -1$. Consequently, equation (24) reduces to

$$\sum_{t=1}^T \sum_{i: d(k,i) \leq T-t} \log \frac{\nu_{it}(\theta)}{\nu_{it}(\hat{\theta})} = \sum_{i=1}^n \sum_{t=1}^{T-d(k,i)} \log \frac{\ell_{it}^{\theta}(\omega_{it})}{\ell_{it}^{\hat{\theta}}(\omega_{it})} + \sum_{\tau=1}^{T-1} \sum_{j: d(k,j) \leq T-\tau} \log \frac{\nu_{j\tau}(\theta)}{\nu_{j\tau}(\hat{\theta})},$$

and as a result, $\log(\nu_{kT}(\theta)/\nu_{kT}(\hat{\theta})) = \sum_{i=1}^n \sum_{t=1}^{T-d(k,i)} \log(\ell_{it}^{\theta}(\omega_{it})/\ell_{it}^{\hat{\theta}}(\omega_{it}))$. In other words, agent k 's log-likelihood ratio at time T is simply the sum of the log-likelihood ratios of all realized signals throughout the social network that could have reached her by time T . As such, agent k 's belief on state θ at time T coincides with $\mathbb{P}(\theta|\omega_{it} : d(k,i) \leq T-t)$, which completes the proof. \square

Proof of parts (a) and (c). Suppose agents' beliefs follow the dynamics in (23) with weights given by (22). It is sufficient to show that the resulting beliefs coincide with those of Bayesian agents under the common knowledge of Bayesian rationality.

Recall from Lemma A.4 that if agents' beliefs satisfy (23), the corresponding beliefs also satisfy $\nu_{it}(\theta) = \mathbb{P}(\theta|\mathcal{G}_{it})$, where $\mathcal{G}_{it} = \sigma(\{\omega_{j\tau} : d(i,j) \leq t - \tau\})$ is the σ -field generated by all the signals that could have potentially reached agent i up to time t . On the other hand, under the common knowledge of Bayesian rationality, the belief that a Bayesian agent i assigns to state θ at time t is, by definition, given by $\mu_{it}(\theta) = \mathbb{P}(\theta|\mathcal{F}_{it})$, where $\mathcal{F}_{it} = \sigma(\omega_{i1}, \dots, \omega_{it}, \{\mu_{j\tau} : j \in N_i, \tau \leq t-1\})$ is agent i 's information set at time t , consisting of all her private signals up to time t and her neighbors' beliefs up to time $t-1$. The proof is therefore complete once we show that $\mu_{it}(\theta) = \nu_{it}(\theta)$ for all i , all $\theta \in \Theta$, and all t . We establish this step by relying on strong induction. The statement trivially holds for $t=0$, as both beliefs coincide with agents' prior beliefs. As the induction hypothesis, assume that $\mu_{j\tau} = \nu_{j\tau}$ for all $\tau < t$ and all agents j and consider the beliefs at time t . Since $\nu_{it}(\theta) = \mathbb{P}(\theta|\mathcal{G}_{it})$, by definition of conditional expectations, we have $\mathbb{E}[(\nu_{it}(\theta) - \mathbb{I}_{\{\theta\}})^2|\mathcal{G}_{it}] \leq \mathbb{E}[(\mu_{it}(\theta) - \mathbb{I}_{\{\theta\}})^2|\mathcal{G}_{it}]$, where we are using the fact that $\mu_{it} \in \mathcal{F}_{it} \subseteq \mathcal{G}_{it}$. Therefore,

$$\mathbb{E}[(\nu_{it}(\theta) - \mathbb{I}_{\{\theta\}})^2|\mathcal{F}_{it}] \leq \mathbb{E}[(\mu_{it}(\theta) - \mathbb{I}_{\{\theta\}})^2|\mathcal{F}_{it}], \quad (25)$$

where once again we are using the fact that $\mathcal{F}_{it} \subseteq \mathcal{G}_{it}$. On the other hand, recall from (23) that the beliefs ν_{it} are recursively constructed as a function of the reports of i 's neighbors up to time $t-1$. Consequently, $\nu_{it} \in \mathcal{H}_{it}$, where $\mathcal{H}_{it} = \sigma(\omega_{i1}, \dots, \omega_{it}, \{\nu_{j\tau} : j \in N_i, \tau \leq t-1\})$. But by the induction hypothesis, $\mathcal{H}_{it} = \mathcal{F}_{it}$, thus guaranteeing that $\nu_{it} \in \mathcal{F}_{it}$. Consequently, the fact $\mu_{it}(\theta) = \mathbb{P}(\theta|\mathcal{F}_{it})$ implies that

$$\mathbb{E}[(\mu_{it}(\theta) - \mathbb{I}_{\{\theta\}})^2|\mathcal{F}_{it}] \leq \mathbb{E}[(\nu_{it}(\theta) - \mathbb{I}_{\{\theta\}})^2|\mathcal{F}_{it}]. \quad (26)$$

The juxtaposition of inequalities (25) and (26) thus guarantees that $\nu_{it}(\theta) = \mu_{it}(\theta)$. \square

Proof of part (b). We already established that agents' beliefs follow the dynamics in (21) with weights $a_{it,j\tau}$ given by (22). Since these weights are independent of the underlying states, it is immediate that the corresponding learning rules satisfy LN for all agents i at all time t .

To establish monotonicity, note that, for any i and j such that $j \in N_i$, equation (22) implies that

$$a_{it,jt-1} = - \sum_{\substack{k: d(i,k) \leq 1 \\ d(k,j) \leq 0}} a_{kt-1,jt-1} = -a_{jt-1,jt-1} = 1.$$

Consequently, the weight $a_{it,jt-1}$ is strictly positive at all times t .

Finally, to establish IIA, consider an arbitrary subset of states $\bar{\Theta} \subseteq \Theta$ and an arbitrary state θ . If $\theta \notin \bar{\Theta}$, it is immediate that both sides of (2) are equal to zero. Hence, it is sufficient to consider the case in which $\theta \in \bar{\Theta}$. From (6) and the definition of conditional probability, it is immediate that

$$\log \text{cond}_{\bar{\Theta}}(f_{it}(\mu_i^t))(\theta) = \sum_{\tau=1}^t \sum_{j \in N_i} a_{it+1,j\tau} \log \mu_{j\tau}(\theta) - \log \left(\sum_{\hat{\theta} \in \bar{\Theta}} \prod_{\tau=1}^t \prod_{j \in N_i} (\mu_{j\tau}(\hat{\theta}))^{a_{it+1,j\tau}} \right), \quad (27)$$

where note that weights $a_{it,j\tau}$ are constants that are independent of the underlying beliefs. On the other hand, equation (6) also implies that

$$\log f_{it}(\text{cond}_{\bar{\Theta}}(\mu_i^t)) = \sum_{\tau=1}^t \sum_{j \in N_i} a_{it+1,j\tau} \log \text{cond}_{\bar{\Theta}}(\mu_{j\tau})(\theta) - \log \left(\sum_{\hat{\theta} \in \bar{\Theta}} \prod_{\tau=1}^t \prod_{j \in N_i} (\text{cond}_{\bar{\Theta}}(\mu_{j\tau})(\hat{\theta}))^{a_{it+1,j\tau}} \right),$$

where we are relying on the fact that $\text{cond}_{\bar{\Theta}}(\mu_{j\tau})(\hat{\theta}) = 0$ for all $\hat{\theta} \notin \bar{\Theta}$. Replacing for $\text{cond}_{\bar{\Theta}}(\mu_{j\tau})(\hat{\theta})$ in the above expression by its definition, i.e., $\text{cond}_{\bar{\Theta}}(\mu_{j\tau})(\theta) = \mu_{j\tau}(\theta) / \sum_{\hat{\theta} \in \bar{\Theta}} \mu_{j\tau}(\hat{\theta})$, immediately implies that the right-hand sides of the above equation coincides with (27), thus establishing (2). \square

Proof of Theorem 4

Let θ denote the underlying state and $A_t = [a_{ijt}]$ denote the matrix of weights in learning rule (3) that agents assign to their neighbors' beliefs at time t , with the convention that $a_{ijt} = 0$ if $j \notin N_i$. Note that when agents' learning rules are unanimous, A_t is a stochastic matrix for all t . For any given state $\hat{\theta} \neq \theta$, equation (4) implies that $x_{t+1} = A_t x_t + y_{t+1}(\omega_{t+1})$, where $x_{it} = \log(\mu_{it}(\theta) / \mu_{it}(\hat{\theta}))$ and $y_{it}(\omega_{it}) = \log(\ell_i^\theta(\omega_{it}) / \ell_i^{\hat{\theta}}(\omega_{it}))$. Consequently,

$$x_t = y_t(\omega_t) + \sum_{\tau=1}^{t-1} A_{t-1} \dots A_{\tau+1} A_\tau y_\tau(\omega_\tau), \quad (28)$$

where we are using the fact that, under uniform prior beliefs, $x_{i0} = 0$ for all agents i . By Lemma A.1, there exists a sequence of uniformly lower-bounded probability vectors v_τ that jointly satisfy (16) for all $t \geq \tau$. Therefore, pre-multiplying both sides of (28) by v_t' and using (16) implies that $v_t' x_t = \sum_{\tau=1}^t v_\tau' y_\tau(\omega_\tau)$. As a result,

$$\liminf_{t \rightarrow \infty} \frac{1}{t} v_t' x_t = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t v_\tau' (y_\tau(\omega_\tau) - h(\theta, \hat{\theta})) + \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t v_\tau' h(\theta, \hat{\theta}),$$

where $h_i(\theta, \hat{\theta}) = \mathbb{E}^\theta[y_{it}(\omega_{it})]$. Since agents' private signals are independently distributed over time, the strong law of large numbers guarantees that the first term on the right-hand side above is equal to zero almost surely, which in turn implies that

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^n v_{it} x_{it} = \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^n v_{i\tau} h_i(\theta, \hat{\theta}) \quad (29)$$

almost surely. Note that Lemma A.1 also guarantees that $\liminf_{t \rightarrow \infty} v_{it} > 0$ for all i . Furthermore, the assumption that agents do not face an identification problem collectively guarantees that there exists an agent i such that $h_i(\theta, \hat{\theta}) > 0$. As a result, with probability one,

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^n v_{it} x_{it} > 0. \quad (30)$$

With the above inequality in hand, it is sufficient to establish that, for any pair of agents i and j ,

$$\lim_{t \rightarrow \infty} \frac{1}{t} (x_{it} - x_{jt}) = 0 \quad (31)$$

almost surely. In particular, (30) and (31), together with the fact that v_t is a probability vector, imply that $\liminf_{t \rightarrow \infty} x_{it}/t > 0$ almost surely for all agents i . Therefore, $\lim_{t \rightarrow \infty} x_{it} = \infty$ with probability one, which subsequently guarantees that $\mu_{it}(\hat{\theta}) \rightarrow 0$ almost surely for all $\hat{\theta} \neq \theta$. In other words, all agents learn the underlying state with probability one.

To establish (31), we follow an approach similar to Liu et al. (2014). Applying part (d) of Lemma A.2 to (28) implies that

$$\max_i x_{it} - \min_i x_{it} \leq \max_i y_{it}(\omega_{it}) - \min_i y_{it}(\omega_{it}) + \sum_{\tau=1}^{t-1} \pi(A_{t-1} \dots A_{\tau+1} A_\tau) (\max_i y_{i\tau}(\omega_{i\tau}) - \min_i y_{i\tau}(\omega_{i\tau})).$$

On the other hand, recall from the proof of Lemma A.1 that $A_t = \mathbb{E}^*[\Lambda_t B_t]$, where Λ_t is a sequence of independent Bernoulli random variables that take value 1 with probability λ_t and B_t is a stochastic matrix whose non-zero elements are uniformly lower bounded by a constant $\eta \in (0, 1)$ for all t . Hence,

$$\begin{aligned} \max_i x_{it} - \min_i x_{it} &\leq \sum_{\tau=1}^t \pi \left(\mathbb{E}^* \prod_{\substack{r: \Lambda_r=1 \\ \tau \leq r < t}} B_r \right) (\max_i y_{i\tau}(\omega_{i\tau}) - \min_i y_{i\tau}(\omega_{i\tau})) \\ &\leq \sum_{\tau=1}^t \mathbb{E}^* \left[\pi \left(\prod_{\substack{r: \Lambda_r=1 \\ \tau \leq r < t}} B_r \right) \right] (\max_i y_{i\tau}(\omega_{i\tau}) - \min_i y_{i\tau}(\omega_{i\tau})), \end{aligned}$$

where the expectation \mathbb{E}^* is over the collection of random variables Λ_t and the second inequality follows from the convexity of π , established in Lemma A.2. Since the set of signals S is finite, there exists a constant $c \geq 0$, independent of t , such that

$$\max_i x_{it} - \min_i x_{it} \leq c \mathbb{E}^* \left[\sum_{\tau=1}^t \pi \left(\prod_{\substack{r: \Lambda_r=1 \\ \tau \leq r < t}} B_r \right) \right].$$

Recall that all matrices in the sequence B_t are irreducible, with non-zero elements that are uniformly lower bounded by η for all t . Therefore, any product of length n of these matrices is element-wise strictly positive, with elements that are lower bounded by η^n . Dividing the matrix product $\prod_{r: \substack{\tau \leq r < t \\ \Lambda_r = 1}} B_r$ into groups of length n and using parts (b) and (c) of Lemma A.2 therefore implies that

$$\pi \left(\prod_{r: \substack{\Lambda_r = 1 \\ \tau \leq r < t}} B_r \right) \leq (1 - \eta^n)^{\lfloor (\Lambda_\tau + \dots + \Lambda_{t-1})/n \rfloor}, \quad (32)$$

where $\lfloor z \rfloor$ denotes the integer part of z . Consequently,

$$\max_i x_{it} - \min_i x_{it} \leq \frac{c}{\beta^n} \mathbb{E}^* \left[\sum_{\tau=1}^t \beta^{(\Lambda_\tau + \dots + \Lambda_{t-1})} \right], \quad (33)$$

where $\beta = (1 - \eta^n)^{1/n} < 1$. Since random variables Λ_t are independent, we have

$$\limsup_{t \rightarrow \infty} \frac{1}{t} (\max_i x_{it} - \min_i x_{it}) \leq \frac{c}{\beta^n} \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t (1 - (1 - \beta)\underline{\lambda}_t)^{t-\tau},$$

where $\underline{\lambda}_t = \min_{1 \leq r < t} \lambda_r$ and we are using the fact that $\mathbb{E}^*[\beta^{\Lambda_t}] = 1 - (1 - \beta)\lambda_t$. As a result,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} (\max_i x_{it} - \min_i x_{it}) \leq \frac{c}{\beta^n} \limsup_{t \rightarrow \infty} \frac{1 - (1 - (1 - \beta)\underline{\lambda}_t)^t}{(1 - \beta)t\underline{\lambda}_t} \leq \frac{c}{(1 - \beta)\beta^n} \limsup_{t \rightarrow \infty} \frac{1}{t\underline{\lambda}_t}.$$

The assumption that $\lim_{t \rightarrow \infty} t\underline{\lambda}_t = \infty$ now establishes (31). \square

Proof of Proposition 1

Let θ denote the underlying state and fix a state $\hat{\theta} \neq \theta$. Consider a pair of agents i and j such that the directed distance from j to i is equal to the social network's diameter, δ , and suppose j is the only agent in the social network with informative signals to distinguish between θ and $\hat{\theta}$, i.e., $\ell_k^\theta = \ell_k^{\hat{\theta}}$ for all $k \neq j$. We prove the result by establishing that $\limsup_{t \rightarrow \infty} x_{it} < \infty$ almost surely, where $x_{it} = \log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta}))$. Recall from (28) that $x_t = y_t(\omega_t) + \sum_{\tau=1}^{t-1} A_{t-1} \dots A_{\tau+1} A_\tau y_\tau(\omega_\tau)$, where $y_{jt}(\omega_{jt}) = \log(\ell_j^\theta(\omega_{jt})/\ell_j^{\hat{\theta}}(\omega_{jt}))$. On the other hand, recall from the proof of Lemma A.1 that $A_t = \mathbb{E}^*[(1 - \Lambda_t)I + \Lambda_t B_t]$, where B_t is a stochastic matrix whose non-zero elements are uniformly lower bounded by some positive constant η and Λ_t is a Bernoulli random variable that takes value 1 with probability λ_t . Consequently, $x_t = y_t(\omega_t) + \sum_{\tau=1}^{t-1} \mathbb{E}^*[Q_{t,\tau}] y_\tau(\omega_\tau)$, where $Q_{t,\tau} = \prod_{r: \substack{\Lambda_r = 1 \\ \tau \leq r < t}} B_r$. Since j is the only agent whose signals allow her to tell the two states apart, we have

$$x_{it} = \sum_{\tau=1}^{t-1} \mathbb{E}^*[(Q_{t,\tau})_{ij}] y_{j\tau}(\omega_{j\tau}) \leq \bar{y} \sum_{\tau=1}^{t-1} \mathbb{E}^*[(Q_{t,\tau})_{ij}] \quad (34)$$

where $\bar{y} = \max_{\omega \in S} |y_{j\tau}(\omega)| < \infty$. Note that $Q_{t,\tau}$ is a product of stochastic matrices, which means that all its elements are upper bounded by 1. Furthermore, the fact that the (directed) distance from agent j to agent i is equal to δ means that the (i, j) element of matrix product $Q_{t,\tau} = \prod_{r: \substack{\Lambda_r = 1 \\ \tau \leq r < t}} B_r$ is non-zero only if the product contains at least δ many matrices, which is equivalent to the condition that

$\Lambda_\tau + \dots + \Lambda_{t-1} \geq \delta$. Consequently,

$$\sum_{\tau=1}^{t-1} (Q_{t,\tau})_{ij} \leq \sum_{\tau=1}^{t-1} \mathbb{I}_{\{\Lambda_\tau + \dots + \Lambda_{t-1} \geq \delta\}} \leq \max \left\{ \tau : \sum_{r=\tau}^{\infty} \Lambda_r = \delta \right\}.$$

Plugging the above inequality into (34) therefore implies that

$$x_{it} \leq \bar{y} \mathbb{E}^* \left[\max \left\{ \tau : \sum_{r=\tau}^{\infty} \Lambda_r = \delta \right\} \right] = \bar{y} \sum_{\tau=1}^{\infty} \tau \lambda_\tau \mathbb{P}^* \left(\sum_{r=\tau+1}^{\infty} \Lambda_r = \delta - 1 \right). \quad (35)$$

Since Λ_t 's are independent Bernoulli random variables, Le Cam's Theorem (Steele, 1994, p. 48) implies

$$\mathbb{P}^* \left(\sum_{r=\tau+1}^{\infty} \Lambda_r = \delta - 1 \right) \leq \bar{\lambda}_\tau^{\delta-1} \frac{\exp(-\bar{\lambda}_\tau)}{(\delta-1)!} + 2 \sum_{r=\tau+1}^{\infty} \lambda_r^2,$$

where $\bar{\lambda}_\tau = \sum_{r=\tau+1}^{\infty} \lambda_r$. In addition, the fact that $\lambda_t = 1/t^\alpha$ implies that $\bar{\lambda}_\tau \leq \tau^{1-\alpha}/(\alpha-1)$. Hence,

$$\mathbb{P}^* \left(\sum_{r=\tau+1}^{\infty} \Lambda_r = \delta - 1 \right) \leq \frac{\tau^{(1-\alpha)(\delta-1)}}{(\alpha-1)^{\delta-1}(\delta-1)!} + \frac{2\tau^{1-2\alpha}}{2\alpha-1},$$

where we are using the fact that $\sum_{r=\tau+1}^{\infty} r^{-2\alpha} \leq \tau^{1-2\alpha}/(2\alpha-1)$ for all $\alpha > 1$. Therefore, (35) leads to

$$x_{it} \leq \frac{\bar{y}}{(\alpha-1)^{\delta-1}(\delta-1)!} \sum_{\tau=1}^{\infty} \tau^{\delta(1-\alpha)} + \frac{2\bar{y}}{2\alpha-1} \sum_{\tau=1}^{\infty} \tau^{2-3\alpha}.$$

The facts that $\delta(1-\alpha) < -1$ and $2-3\alpha < -1$ now guarantee that the right-hand side above is finite. \square

Proof of Theorem 5

Let θ denote the underlying state and $\zeta_{it} = f_{it}(\mu_i^t)$ be the interim belief of agent i after observing her neighbors' time- t reports but before observing her private signal ω_{it+1} . Since $\mu_{it+1} = \text{BU}(\zeta_{it}; \omega_{it+1})$, it is sufficient to show that $\zeta_{it}(\theta) \rightarrow 1$ for all i .

First, we show agent i 's interim belief on the true state is uniformly bounded away from zero, i.e.,

$$\liminf_{t \rightarrow \infty} \zeta_{it}(\theta) > 0 \quad (36)$$

with \mathbb{P}^θ -probability one. To this end, recall from Theorem 2 that $\zeta_{it}(\hat{\theta}) = \sum_{j \in N_i} a_{ijt} \mu_{jt}(\hat{\theta})$, where $\sum_{j \in N_i} a_{ijt} = 1$ due to the unanimity of i 's learning rule. Therefore, agent i 's interim belief satisfies

$$\zeta_{it+1}(\hat{\theta}) = \sum_{j=1}^n a_{ijt+1} \frac{\zeta_{jt}(\hat{\theta}) \ell_j^{\hat{\theta}}(\omega_{jt+1})}{m_{jt}(\omega_{jt+1})} \quad (37)$$

for all $\hat{\theta} \in \Theta$, where $m_{jt}(\omega_{jt+1}) = \sum_{\tilde{\theta} \in \Theta} \zeta_{jt}(\tilde{\theta}) \ell_j^{\tilde{\theta}}(\omega_{jt+1})$. On the other hand, by Lemma A.1, there exists a sequence of uniformly lower-bounded probability vectors v_τ that jointly satisfy (16) for all $t \geq \tau$. Thus, taking logarithms from both sides of (37), multiplying by v_{it+2} , and summing over all i lead to

$$\sum_{i=1}^n v_{it+2} \log(\zeta_{it+1}(\hat{\theta})) \geq \sum_{i=1}^n v_{it+1} \log(\zeta_{jt}(\hat{\theta})) + \sum_{i=1}^n v_{it+1} \log \left(\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} \right),$$

where we are using the concavity of logarithm and the fact that $\sum_{i=1}^n v_{it+2} a_{ijt+1} = v_{jt+1}$ (established in equation (16)). Evaluate the beliefs at the underlying state θ and take expectations from both sides of the above equation to obtain

$$\mathbb{E}_t^\theta \left[\sum_{i=1}^n v_{it+2} \log(\zeta_{it+1}(\theta)) \right] \geq \sum_{i=1}^n v_{it+1} \log(\zeta_{it}(\theta)),$$

where note that $\mathbb{E}_t^\theta \log(\ell_i^\theta(\omega_{it+1})/m_{it}(\omega_{it+1}))$ is the relative entropy of ℓ_i^θ with respect to m_{it} and as a result is always non-negative. The above inequality implies that $\sum_{i=1}^n v_{it+1} \log(\zeta_{it}(\theta))$ is a non-positive submartingale and hence converges \mathbb{P}^θ -almost surely. Since Lemma A.1 also guarantees that v_{it} is uniformly bounded away from zero for all i , it is immediate that $\zeta_{it}(\theta)$ remains bounded away from zero for all i with \mathbb{P}^θ -probability one.

With inequality (36) in hand, we next establish that, for any arbitrary signal $\omega \in S$,

$$\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_\tau^\theta \left[\sum_{\hat{\theta} \in \Theta} \zeta_{it}(\hat{\theta}) \frac{\ell_i^{\hat{\theta}}(\omega)}{\ell_i^\theta(\omega)} - 1 \right]^2 = 0 \quad (38)$$

with \mathbb{P}^θ -probability one. To this end, multiply both sides of (37) by v_{it+2} , sum over all i , and evaluate the beliefs at the underlying state θ to obtain $\sum_{i=1}^n v_{it+2} \zeta_{it+1}(\theta) = \sum_{i=1}^n v_{it+1} \zeta_{it}(\theta) \ell_i^\theta(\omega_{it+1})/m_{it}(\omega_{it+1})$, where once again we are relying on (16) to conclude that $\sum_{i=1}^n v_{it+2} a_{ijt+1} = v_{jt+1}$. As a result,

$$\mathbb{E}_t^\theta \left[\sum_{i=1}^n v_{it+2} \zeta_{it+1}(\theta) \right] - \sum_{i=1}^n v_{it+1} \zeta_{it}(\theta) = \sum_{i=1}^n v_{it+1} \zeta_{it}(\theta) \mathbb{E}_t^\theta \left[\frac{\ell_i^\theta(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right]. \quad (39)$$

By Jensen's inequality, $\mathbb{E}_t^\theta [\ell_i^\theta(\omega_{it+1})/m_{it}(\omega_{it+1})] \geq (\mathbb{E}_t^\theta [m_{it}(\omega_{it+1})/\ell_i^\theta(\omega_{it+1})])^{-1} = 1$, which implies that $\sum_{i=1}^n v_{it+1} \zeta_{it}(\theta)$ is a bounded submartingale and hence converges \mathbb{P}^θ -almost surely. This, coupled with the observation that the right-hand side of (39) is non-negative, thus guarantees that

$$\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_\tau^\theta \left[\sum_{i=1}^n v_{it+1} \zeta_{it}(\theta) \mathbb{E}_t^\theta \left[\frac{\ell_i^\theta(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] \right] = 0$$

with \mathbb{P}^θ -probability one. Inequality (36) and Lemma A.1 guarantee, respectively, that $\zeta_{it}(\theta)$ and v_{it+1} are uniformly bounded away from zero. Therefore,

$$\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_\tau^\theta \mathbb{E}_t^\theta \left[\frac{\ell_i^\theta(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] = 0 \quad (40)$$

for all agents i with \mathbb{P}^θ -probability one. Furthermore,

$$\mathbb{E}_t^\theta \left[\frac{\ell_i^\theta(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] = \sum_{\omega \in S} \frac{(\ell_i^\theta(\omega))^2}{m_{it}(\omega)} \left(\sum_{\hat{\theta} \in \Theta} \zeta_{it}(\hat{\theta}) \frac{\ell_i^{\hat{\theta}}(\omega)}{\ell_i^\theta(\omega)} - 1 \right)^2,$$

where we are using the fact that, by definition, $m_{it}(\omega_{it+1}) = \sum_{\hat{\theta} \in \Theta} \zeta_{it}(\hat{\theta}) \ell_i^{\hat{\theta}}(\omega_{it+1})$. Plugging the above into (40) and using the fact that ℓ_i^θ is has full support over S establishes (38).

Lemma A.5. For any given integer k and all collections of signals $(\omega_1, \dots, \omega_k) \in S^k$,

$$\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left[\sum_{\hat{\theta} \in \Theta} \zeta_{it}(\hat{\theta}) \prod_{r=1}^k \frac{\ell_i^{\hat{\theta}}(\omega_r)}{\ell_i^{\theta}(\omega_r)} - 1 \right]^2 = 0 \quad \mathbb{P}^{\theta}\text{-almost surely.} \quad (41)$$

Proof. Equation (38) establishes the lemma for $k = 1$. The proof for $k > 1$ follows an inductive argument on k and is provided in Online Appendix B. \square

Lemma A.6. Let $\hat{\Theta}_i = \{\hat{\theta} \in \Theta : \ell_i^{\hat{\theta}} \neq \ell_i^{\theta}\}$. Then, $\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} [\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta})]^2 = 0$ with \mathbb{P}^{θ} -probability one. Furthermore, $\lim_{t \rightarrow \infty} \zeta_{it}(\hat{\theta}) = 0$ with \mathbb{P}^{θ} -probability one for all $\hat{\theta} \in \hat{\Theta}_i$.

Proof. By Lemma 4 of Jadbabaie et al. (2012), there exist a constant $\beta \in (0, 1)$, a positive integer \tilde{k} , and a collection of signals $(\tilde{\omega}_1, \dots, \tilde{\omega}_{\tilde{k}})$ such that $\prod_{r=1}^{\tilde{k}} \ell_i^{\hat{\theta}}(\tilde{\omega}_r) / \ell_i^{\theta}(\tilde{\omega}_r) < 1 - \beta$ for all $\hat{\theta} \in \hat{\Theta}_i$. Therefore,

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left[\sum_{\hat{\theta} \in \Theta} \zeta_{it}(\hat{\theta}) \prod_{r=1}^{\tilde{k}} \frac{\ell_i^{\hat{\theta}}(\tilde{\omega}_r)}{\ell_i^{\theta}(\tilde{\omega}_r)} - 1 \right]^2 &= \lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left[\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta}) \left(\prod_{r=1}^{\tilde{k}} \frac{\ell_i^{\hat{\theta}}(\tilde{\omega}_r)}{\ell_i^{\theta}(\tilde{\omega}_r)} - 1 \right) \right]^2 \\ &\geq \beta^2 \lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left[\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta}) \right]^2. \end{aligned}$$

On the other hand, Lemma A.5 guarantees that the left-hand side of the above inequality is equal to zero with \mathbb{P}^{θ} -probability one. Since $\beta > 0$, it is immediate that $\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} [\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta})]^2 = 0$. The second claim then follows as $\lim_{\tau \rightarrow \infty} (\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta}))^2 \leq \lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} [\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta})]^2 = 0$. \square

Lemma A.7. $\sum_{i=1}^n v_{it+1} \zeta_{it}(\hat{\theta})$ converges \mathbb{P}^{θ} -almost surely for all $\hat{\theta} \in \Theta$.

Proof. Equation (39) guarantees that $\sum_{i=1}^n v_{it+1} \zeta_{it}(\theta)$ is a bounded submartingale and hence converges \mathbb{P}^{θ} -almost surely. It is therefore sufficient to establish the lemma for $\hat{\theta} \neq \theta$. Given any such state, equations (16) and (37) imply that

$$\mathbb{E}_t^{\theta} \left[\sum_{i=1}^n v_{it+2} \zeta_{it+1}(\hat{\theta}) \right] - \sum_{i=1}^n v_{it+1} \zeta_{it}(\hat{\theta}) = \sum_{i=1}^n v_{it+1} \zeta_{it}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right].$$

Denoting the left-hand side of the above equation by $\Delta_t(\hat{\theta})$, we have

$$\Delta_t(\hat{\theta}) = \sum_{i: \hat{\theta} \notin \hat{\Theta}_i} v_{it+1} \zeta_{it}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] + \sum_{i: \hat{\theta} \in \hat{\Theta}_i} v_{it+1} \zeta_{it}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right],$$

where recall that $\hat{\Theta}_i = \{\hat{\theta} \in \Theta : \ell_i^{\hat{\theta}} \neq \ell_i^{\theta}\}$. Consequently,

$$\sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} |\Delta_t(\hat{\theta})| \leq \sum_{i: \hat{\theta} \notin \hat{\Theta}_i} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] + \sum_{i: \hat{\theta} \in \hat{\Theta}_i} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left| \zeta_{it}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] \right|, \quad (42)$$

where we are using the fact that $v_{it+1} \in [0, 1]$ for all i and t and that $\mathbb{E}_t^{\theta} [\ell_i^{\hat{\theta}}(\omega_{it+1}) / m_{it}(\omega_{it+1})] \geq 1$. Equation (40) guarantees that the first term on the right-hand side above converges to zero \mathbb{P}^{θ} -almost surely as $\tau \rightarrow \infty$. We show the second term also has a limit of zero. Note that if $\hat{\theta} \in \hat{\Theta}_i$, then

$$\left| \zeta_{it}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] \right| \leq \zeta_{it}(\hat{\theta}) \sum_{\omega \in S} \ell_i^{\hat{\theta}}(\omega) \left| \frac{\ell_i^{\theta}(\omega)}{m_{it}(\omega)} - 1 \right| \leq \sum_{\omega \in S} \zeta_{it}(\hat{\theta}) \frac{\ell_i^{\hat{\theta}}(\omega)}{m_{it}(\omega)} \sum_{\tilde{\theta} \in \hat{\Theta}_i} \zeta_{it}(\tilde{\theta}) \left| \ell_i^{\theta}(\omega) - \ell_i^{\tilde{\theta}}(\omega) \right|.$$

As a result,

$$\sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left| \zeta_{it}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right] \right| \leq c \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} \left[\sum_{\hat{\theta} \in \hat{\Theta}_i} \zeta_{it}(\hat{\theta}) \right]^2 \quad (43)$$

for some positive constant c , where once again we are using the fact $\ell_i^{\hat{\theta}}$ has full support over the finite set S for all $\hat{\theta}$. By Lemma A.6, the right-hand side of the above inequality converges to zero \mathbb{P}^{θ} -almost surely as $\tau \rightarrow \infty$. Therefore, inequality (42) implies that $\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} |\Delta_t(\hat{\theta})| = 0$ with \mathbb{P}^{θ} -probability one. This observation implies that

$$\limsup_{\tau \rightarrow \infty} \sup_{t > \tau} \left| \mathbb{E}_{\tau}^{\theta} \left[\sum_{i=1}^n v_{it+1} \zeta_{it}(\hat{\theta}) \right] - \sum_{i=1}^n v_{i\tau+1} \zeta_{i\tau}(\hat{\theta}) \right| = \limsup_{\tau \rightarrow \infty} \sup_{t > \tau} \left| \mathbb{E}_{\tau}^{\theta} \sum_{r=\tau}^{t-1} \Delta_r(\hat{\theta}) \right| \leq \limsup_{\tau \rightarrow \infty} \sup_{t > \tau} \sum_{r=\tau}^{t-1} \mathbb{E}_{\tau}^{\theta} |\Delta_r(\hat{\theta})| = 0.$$

Hence, by Theorem 1 of Mucci (1976), the sequence $\sum_{i=1}^n v_{it+1} \zeta_{it}(\hat{\theta})$ converges \mathbb{P}^{θ} -almost surely. \square

Lemma A.8. $\lim_{\tau \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbb{E}_{\tau}^{\theta} [\zeta_{it}(\hat{\theta})] = 0$ for all $\hat{\theta} \neq \theta$ and all i with \mathbb{P}^{θ} -probability one.

Proof. For a given state $\hat{\theta} \neq \theta$, define $y_{it+1}(\hat{\theta}) = \sum_{j=1}^n a_{ijt+1} \zeta_{jt}(\hat{\theta}) \left(\frac{\ell_j^{\hat{\theta}}(\omega_{jt+1})}{m_{jt}(\omega_{jt+1})} - 1 \right)$. We first show that

$$\lim_{\tau \rightarrow \infty} \sum_{t=\tau}^{\infty} |\mathbb{E}_{\tau}^{\theta} [y_{it+1}]| = 0 \quad (44)$$

for all agents i with \mathbb{P}^{θ} -probability one. To this end, note that

$$\sum_{t=\tau}^{\infty} \mathbb{E}_{\tau}^{\theta} |\mathbb{E}_t^{\theta} [y_{it+1}]| \leq \sum_{t=\tau}^{\infty} \sum_{j: \hat{\theta} \notin \hat{\Theta}_j} \mathbb{E}_{\tau}^{\theta} \left[\frac{\ell_j^{\hat{\theta}}(\omega_{jt+1})}{m_{jt}(\omega_{jt+1})} - 1 \right] + \sum_{t=\tau}^{\infty} \sum_{j: \hat{\theta} \in \hat{\Theta}_j} \mathbb{E}_{\tau}^{\theta} \left| \zeta_{jt}(\hat{\theta}) \mathbb{E}_t^{\theta} \left[\frac{\ell_j^{\hat{\theta}}(\omega_{jt+1})}{m_{jt}(\omega_{jt+1})} - 1 \right] \right|.$$

Equation (40) guarantees that the first term on the right-hand side of the above expression converges to zero \mathbb{P}^{θ} -almost surely as $\tau \rightarrow \infty$. Moreover, the juxtaposition of (43) with Lemma A.6 implies that the second term also vanishes as $\tau \rightarrow \infty$. A simple application of Jensen's inequality leads to (44).

As our next step, we note that the interim beliefs of agent i satisfy equation (37), which can be written in matrix form as $\zeta_{t+1}(\hat{\theta}) = A_{t+1} \zeta_t(\hat{\theta}) + y_{t+1}(\hat{\theta})$. Iterating the belief dynamics equation forward and taking expectations from both sides implies that $\mathbb{E}_{\tau}^{\theta} [\zeta_t(\hat{\theta})] = A_t A_{t-1} \dots A_{\tau+1} \zeta_{\tau}(\hat{\theta}) + \sum_{r=\tau+1}^t A_t A_{t-1} \dots A_{r+1} \mathbb{E}_{\tau}^{\theta} [y_r(\hat{\theta})]$. Thus, by part (d) of Lemma A.2,

$$\begin{aligned} \max_i \mathbb{E}_{\tau}^{\theta} [\zeta_{it}(\hat{\theta})] - \min_i \mathbb{E}_{\tau}^{\theta} [\zeta_{it}(\hat{\theta})] &\leq \pi(A_t \dots A_{\tau+1}) + \sum_{r=\tau+1}^t \pi(A_t \dots A_{r+1}) (\max_i \mathbb{E}_{\tau}^{\theta} [y_{ir}(\hat{\theta})] - \min_i \mathbb{E}_{\tau}^{\theta} [y_{ir}(\hat{\theta})]) \\ &\leq \pi(A_t \dots A_{\tau+1}) + \sum_{r=\tau+1}^t (\max_i \mathbb{E}_{\tau}^{\theta} [y_{ir}(\hat{\theta})] - \min_i \mathbb{E}_{\tau}^{\theta} [y_{ir}(\hat{\theta})]) \end{aligned}$$

where in the first inequality we are using the fact that $\max_i \zeta_{i\tau}(\hat{\theta}) - \min_i \zeta_{i\tau}(\hat{\theta}) \leq 1$ and the second inequality is a consequence of part (a) of Lemma A.2. On the other hand, recall from the proof of Lemma A.1 that $A_t = \mathbb{E}^*[\Lambda_t B_t]$, where Λ_t is a sequence of independent Bernoulli random variables

that take value 1 with probability λ_t and B_t is a stochastic matrix whose non-zero elements are lower bounded by a constant $\eta \in (0, 1)$ that is independent of t . Consequently,

$$\begin{aligned} \max_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] - \min_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] &\leq \pi \left(\mathbb{E}^* \prod_{r: \Lambda_r=1, \tau < r \leq t} B_r \right) + \sum_{r=\tau+1}^t (\max_i \mathbb{E}_\tau^\theta[y_{ir}(\hat{\theta})] - \min_i \mathbb{E}_\tau^\theta[y_{ir}(\hat{\theta})]) \\ &\leq \mathbb{E}^* \left[\pi \left(\prod_{r: \Lambda_r=1, \tau < r \leq t} B_r \right) \right] + 2 \sum_{r=\tau+1}^t \max_i |\mathbb{E}_\tau^\theta[y_{ir}(\hat{\theta})]|, \end{aligned}$$

where the second inequality relies on the convexity of π , established in Lemma A.2. Following (32) and (33), the above inequality further reduces to

$$\max_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] - \min_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] \leq \frac{1}{\beta^n} \mathbb{E}^* \left[\beta^{(\Lambda_{\tau+1} + \dots + \Lambda_t)} \right] + 2 \sum_{r=\tau+1}^t \max_i |\mathbb{E}_\tau^\theta[y_{ir}(\hat{\theta})]|,$$

where $\beta = (1 - \eta^n)^{1/n} < 1$. Since $\sum_{t=0}^\infty \lambda_t = \infty$, the Borel-Cantelli lemma implies that $\Lambda_t = 1$ infinitely often with probability one. Hence, by the dominated convergence theorem, the limit of the first term on the right-hand side of the above inequality is equal to zero as $t \rightarrow \infty$. Consequently,

$$\lim_{t \rightarrow \infty} (\max_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] - \min_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})]) \leq 2 \sum_{r=\tau+1}^\infty \max_i |\mathbb{E}_\tau^\theta[y_{ir}(\hat{\theta})]|$$

for all τ . Recall that (44) guarantees that the right-hand side of the above inequality converges to zero with \mathbb{P}^θ -probability one as $\tau \rightarrow \infty$. Therefore, $\lim_{\tau \rightarrow \infty} \lim_{t \rightarrow \infty} (\max_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] - \min_i \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})]) = 0$. Finally, note that, by assumption, there exists an agent i such that $\ell_i^\theta \neq \ell_i^{\hat{\theta}}$ for whom, by Lemma A.6, $\lim_{t \rightarrow \infty} \zeta_{it}(\hat{\theta}) = 0$ with \mathbb{P}^θ -probability one. Hence, $\lim_{\tau \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] = 0$ for all i . \square

The juxtaposition of Lemmas A.7 and A.8 completes the proof of Theorem 5. In particular, whereas the former establishes that $\sum_{i=1}^n v_{it+1} \zeta_{it}(\hat{\theta})$ converges to some limit $\zeta^*(\hat{\theta})$ with \mathbb{P}^θ -probability one for all $\hat{\theta} \in \Theta$, the latter implies that $\lim_{\tau \rightarrow \infty} \lim_{t \rightarrow \infty} \sum_{i=1}^n v_{it+1} \mathbb{E}_\tau^\theta[\zeta_{it}(\hat{\theta})] = 0$ almost surely for all $\hat{\theta} \neq \theta$. Hence, it must be the case that $\zeta^*(\hat{\theta}) = 0$ with \mathbb{P}^θ -probability one for all $\hat{\theta} \neq \theta$, which in turn implies that $\lim_{t \rightarrow \infty} \zeta_{it}(\hat{\theta}) = 0$ for all $\hat{\theta} \neq \theta$ and all i . \square

Proof of Theorem 6

Proof of part (a). For any pair of states $\theta, \hat{\theta} \in \Theta$, the weak-separability assumption implies that

$$\frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \frac{\ell_i^\theta(\omega_{it+1}) \psi_i(\mu_t(\theta))}{\ell_i^{\hat{\theta}}(\omega_{it+1}) \psi_i(\mu_t(\hat{\theta}))}. \quad (45)$$

Multiplying both sides of the above equation by $\mu_{it+1}(\hat{\theta})$ and rewriting it in vector form leads to

$$\mu_{t+1}(\theta) = \text{diag} \left(\frac{\ell_i^\theta(\omega_{it+1})}{\ell_i^{\hat{\theta}}(\omega_{it+1})} \right) \text{diag} \left(\frac{\psi_i(\mu_t(\theta))}{\psi_i(\mu_t(\hat{\theta}))} \right) \mu_{t+1}(\hat{\theta}),$$

where $\text{diag}(z)$ is a square diagonal matrix whose diagonal elements are given by the elements of vector z . For a given agent j , applying the mapping ψ_j to both sides of the above inequality results in

$$\psi_j(\mu_{t+1}(\theta)) \leq \left(\max_i \frac{\ell_i^\theta(\omega_{it+1})}{\ell_i^{\hat{\theta}}(\omega_{it+1})} \right)^\rho \left(\max_i \frac{\psi_i(\mu_t(\theta))}{\psi_i(\mu_t(\hat{\theta}))} \right)^\rho \psi_j(\mu_{t+1}(\hat{\theta})),$$

where ρ is the degree of homogeneity of ψ_j and we are using the fact that ψ_j is non-decreasing in all its arguments (a consequence of monotonicity of j 's learning rule). Dividing both sides of the inequality by $\psi_j(\mu_{t+1}(\hat{\theta}))$ and taking logarithms from both sides implies that

$$x_{t+1} \leq \rho x_t + \rho y(\omega_t), \quad (46)$$

where $x_t = \max_j \log(\psi_j(\mu_t(\theta))/\psi_j(\mu_t(\hat{\theta})))$ and $y(\omega_t) = \max_i \log(\ell_i^\theta(\omega_{it})/\ell_i^{\hat{\theta}}(\omega_{it}))$. Therefore,

$$x_t \leq \rho^{t-n} x_n + \sum_{\tau=n}^{t-1} \rho^{t-\tau} y(\omega_\tau)$$

for all $t \geq n + 1$, where n denotes the number of agents in the social social network.

Fix an arbitrary time $T \geq n$ and an arbitrary pair of states $\theta \neq \hat{\theta}$ and suppose that $x_n < 0$ with positive probability — a claim that we prove below. For any $t > T$, the above inequality implies that $\rho^{-t} x_t \leq \rho^{-n} x_n + \sum_{\tau=n}^{T-1} \rho^{-\tau} y(\omega_\tau) + \bar{y} \sum_{\tau=T}^{t-1} 1/\rho^\tau$, where $\bar{y} = \max_i \max_{\omega_i} \log(\ell_i^\theta(\omega_i)/\ell_i^{\hat{\theta}}(\omega_i))$. As a result,

$$\limsup_{t \rightarrow \infty} \rho^{-t} x_t \leq \rho^{-n} x_n + \sum_{\tau=n}^{T-1} \rho^{-\tau} y(\omega_\tau) + \frac{\bar{y}/\rho^T}{1 - 1/\rho}. \quad (47)$$

On the other hand, note that there exists a signal profile $\tilde{\omega} = (\tilde{\omega}_1, \dots, \tilde{\omega}_n)$ such that $\ell_i^\theta(\tilde{\omega}_i) \leq \ell_i^{\hat{\theta}}(\tilde{\omega}_i)$ for all i , with the inequality holding strictly for at least one agent j .²¹ It is immediate that for such a signal profile, $y(\tilde{\omega}) \leq 0$. Furthermore, since $\rho > 1$ and $x_n < 0$, there exists a finite T large enough such that $\rho^{-n} x_n + \frac{\bar{y}/\rho^T}{1 - 1/\rho} < 0$. Consequently, there exists a large enough T such that the right-hand side of (47) is strictly negative on any path on which signal profile $\tilde{\omega}$ is realized in every period over the finite interval between $t = n$ and $t = T - 1$. This observation, together with the fact that $\rho > 1$, subsequently guarantees that $\lim_{t \rightarrow \infty} \log(\psi_i(\mu_t(\theta))/\psi_i(\mu_t(\hat{\theta}))) = -\infty$ with some strictly positive probability for all i , regardless of the underlying state. Subsequently, equation (45) guarantees that with positive probability $\lim_{t \rightarrow \infty} \log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta})) = -\infty$ for all i and all pairs $\hat{\theta} \neq \theta$. That is, all agents assign an asymptotic belief of zero to θ with a positive probability, even when θ is the underlying state.

The proof is complete once we show that $x_n < 0$ with strictly positive probability. We rely on a two-step argument to show that $x_n < 0$ on any path on which the signal profile $\tilde{\omega}$ is realized in every period between $t = 1$ and $t = n$.

First, we show inductively that, on any such path, $\mu_{it}(\theta) \leq \mu_{it}(\hat{\theta})$ for all i and all $t \leq n$ and that the inequality is strict for any agent i with a neighbor j such that $\mu_{jt-1}(\theta) < \mu_{jt-1}(\hat{\theta})$. Recall that, by the uniform prior assumption, $\mu_{i0}(\theta) \leq \mu_{i0}(\hat{\theta})$ for all i , an inequality that serves as the induction's base. As the induction hypothesis, suppose $\mu_{it-1}(\theta) \leq \mu_{it-1}(\hat{\theta})$ for all i . By monotonicity, $\psi_i(\mu_{t-1}(\theta)) \leq \psi_i(\mu_{t-1}(\hat{\theta}))$ for all i and $\psi_i(\mu_{t-1}(\theta)) < \psi_i(\mu_{t-1}(\hat{\theta}))$ if i has a neighbor j such that $\mu_{jt-1}(\theta) < \mu_{jt-1}(\hat{\theta})$. Furthermore, recall that $\ell_i^\theta(\tilde{\omega}_i) \leq \ell_i^{\hat{\theta}}(\tilde{\omega}_i)$ for all i . The above claim then follows from (45).

²¹The existence of such an agent j is a consequence of the assumption that the underlying state is globally identifiable.

With the above claim in hand, we next show that $\mu_{in}(\theta) < \mu_{in}(\hat{\theta})$ for all i on any path on which the signal profile $\tilde{\omega}$ is realized in every period between $t = 1$ and $t = n$. Once again, recall that the signal profile $\tilde{\omega}$ is such that $\ell_i^\theta(\tilde{\omega}_i) \leq \ell_i^{\hat{\theta}}(\tilde{\omega}_i)$ for all i and that $\ell_j^\theta(\tilde{\omega}_j) < \ell_j^{\hat{\theta}}(\tilde{\omega}_j)$ for some agent j . We show inductively that $\mu_{it}(\theta) < \mu_{it}(\hat{\theta})$ for all i that are at distance $t - 1$ or less from agent j on the social network, where $t \leq n$. Given that $\ell_j^\theta(\tilde{\omega}_j) < \ell_j^{\hat{\theta}}(\tilde{\omega}_j)$, equation (45) implies that $\mu_{j1}(\theta) < \mu_{j1}(\hat{\theta})$, an inequality that will serve as the induction's base. As the induction hypothesis, suppose that $\mu_{it}(\theta) < \mu_{it}(\hat{\theta})$ for all agents i who are at distance $t - 1$ or less from agent j on the social network. The claim in the previous paragraph then implies that $\mu_{it+1}(\theta) < \mu_{it+1}(\hat{\theta})$ for all agents i who are at distance t or less from agent j , thus completing the inductive argument.

Since the diameter of the social network is at most n , all nodes are at most at distance $n - 1$ from agent j . Therefore, $\mu_{in}(\theta) < \mu_{in}(\hat{\theta})$ for all i . Consequently, $x_n = \max_j \log(\psi_j(\mu_n(\theta))/\psi_j(\mu_n(\hat{\theta}))) < 0$ on all paths on which the signal profile $\tilde{\omega}$ is realized in every period between $t = 1$ and $t = n$. Noting the fact that such paths have a strictly positive probability completes the proof. \square

Proof of part (b). Consider an arbitrary pair of states $\theta \neq \hat{\theta}$ and recall from the proof of part (a) that, under weak separability, agents' belief dynamics satisfy (45) and (46). Equation (46) implies that $x_{t+1} \leq \rho x_t + \rho \bar{y}$, where $x_t = \max_j \log(\psi_j(\mu_t(\theta))/\psi_j(\mu_t(\hat{\theta})))$ and $\bar{y} = \max_i \max_{\omega_i} \log(\ell_i^\theta(\omega_i)/\ell_i^{\hat{\theta}}(\omega_i))$. Therefore $x_t \leq \bar{y} \sum_{\tau=1}^t \rho^\tau$. The fact that $\rho < 1$ implies that $\sum_{\tau=1}^t \rho^\tau$ is bounded above, even as $t \rightarrow \infty$. Therefore, $\limsup_{t \rightarrow \infty} \log(\psi_i(\mu_t(\theta))/\psi_i(\mu_t(\hat{\theta}))) < \infty$ for all i . Subsequently, equation (45) guarantees that $\limsup_{t \rightarrow \infty} \log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta})) < \infty$ for all i and all pair of states $\hat{\theta} \neq \theta$. In other words, all agents remain asymptotically uncertain, regardless of the realization of the underlying state. \square

Proof of Theorem 7

We prove parts (b) and (c) of Theorem 7 by constructing a social network, a signal structure, and social learning rules with logarithmic curvatures outside the $[-1, 1]$ interval for which agents fail to learn the state asymptotically. The details are provided in Online Appendix B.

In what follows, we prove part (a) of the theorem. Throughout, we assume that $\psi_i : [0, 1]^n \rightarrow \mathbb{R}$ satisfying (10) is such that $\psi_i(\mathbf{1}) = 1$ for all i , where $\mathbf{1} = (1, \dots, 1)$ denotes the vector of all ones. Note that since the social learning rule of agent i can be rewritten as (11), this is a simple normalization and is without loss of generality. We also extend the domain of ψ_i to \mathbb{R}_+^n by defining $\psi_i(x)$ for any $x \in \mathbb{R}_+^n$ as $\psi_i(x) = \psi_i(\epsilon x)/\epsilon$, where $\epsilon > 0$ is sufficiently small to ensure that $\epsilon x \in [0, 1]^n$.²²

We have the following lemmas, the proofs of which are provided in Online Appendix B.

Lemma A.9. *Suppose $g : (0, \infty)^n \rightarrow (0, \infty)$ is smooth, homogeneous of degree 1, weakly increasing, and satisfies $g(\mathbf{1}) = 1$. If the logarithmic curvature of g is smaller than 1, then g is concave. Furthermore, $g(x) \leq \sum_{j=1}^n g^{(j)}(\mathbf{1})x_j$, where $g^{(j)}(x) = \partial g(x)/\partial x_j$.*

Lemma A.10. *The function $\phi_i : (0, \infty)^n \rightarrow (0, \infty)$ defined as $\phi_i(x) = [\psi_i(1/x_1, \dots, 1/x_n)]^{-1}$ is concave, homogeneous of degree 1, strictly increasing in x_j for $j \in N_i$, with a logarithmic curvature that is less than or equal to 1 throughout its domain.*

²²Note that, since ψ_i is homogeneous of degree 1, $\psi_i(\epsilon x)/\epsilon$ is independent of the choice ϵ .

Lemma A.11. *The mapping $\Phi(x) = \lim_{t \rightarrow \infty} \overbrace{\phi \circ \phi \circ \dots \circ \phi}^{t \text{ times}}(x)$ is well-defined, element-wise strictly positive, smooth, homogeneous of degree 1, and concave. Moreover, $\Phi_i^{(j)}(x) > 0$ for all $x > 0$ and all i, j .*

Proof. Theorem 3.2 of [Nussbaum \(1988, p. 93\)](#) guarantees that $\Phi(x)$ is well-defined, smooth, homogeneous of degree 1, and element-wise strictly positive for all $x > 0$. The concavity of Φ is an immediate consequence of the fact that it is the composition of the non-decreasing and concave function ϕ with itself, where both properties were established in [Lemma A.10](#).

We now turn to proving the last statement. By Theorem 3.2 of [Nussbaum \(1988\)](#), there exists a function $q : (0, \infty)^n \rightarrow (0, \infty)$ such that $\Phi(x) = q(x)\mathbf{1}$. It is therefore sufficient to show that $q^{(j)}(x) > 0$ for all $x > 0$. Note that $q(\phi(x)) = q(x)$. Therefore, by chain rule, $q^{(j)}(x) = \sum_{i=1}^n q^{(i)}(\phi(x))\phi_i^{(j)}(x)$. Let $u(x)$ denote the vector whose i -th element is given by $q^{(i)}(x)$ and let $U(x)$ denote the matrix with the (i, j) element equal to $\phi_i^{(j)}(x)$. Hence,

$$u(x)' = u(\phi(x))'U(x) \quad (48)$$

Iterating the above expression thus implies that $u(x)' = u(\phi^t(x))'U(\phi^{t-1}(x)) \dots U(\phi(x))U(x)$, where ϕ^t denotes the composition of ϕ with itself t times. Let x be an element-wise positive vector. The continuity of U and the fact that $\lim_{t \rightarrow \infty} \phi^t(x) = \Phi(x) = q(x)\mathbf{1}$ imply that $\lim_{t \rightarrow \infty} U(\phi^t(x)) = U(q(x)\mathbf{1}) = U(\mathbf{1})$. The last equality is a consequence of the fact that U is homogeneous of degree 0. On the other hand, the monotonicity assumption and the fact that the social network is strongly connected guarantee that $U(\mathbf{1})$ is an irreducible non-negative matrix, whereas Euler's theorem and the assumption that ψ is homogeneous of degree 1 imply that $U(\mathbf{1})$ is row stochastic. Taken together, the above observations imply that $U(\phi^{t-1}(x))$ converges to an irreducible stochastic matrix as $t \rightarrow \infty$. Therefore, by [Chatterjee and Seneta \(1977, p. 93\)](#), $U(\phi^{t-1}(x)) \dots U(\phi(x))U(x)$ converges to an element-wise strictly positive, rank 1 matrix $U^*(x)$. Consequently, $u(x)' = u(\Phi(x))'U^*(x)$, where $U_{ij}^*(x) > 0$. The proof is therefore complete once we show that $u(\Phi(x))$ is element-wise strictly positive. Note that $u(\Phi(x)) = u(q(x)\mathbf{1}) = u(\mathbf{1})$. On the other hand, equation (48) implies that $u(\mathbf{1})' = u(\phi(\mathbf{1}))'U(\mathbf{1}) = u(\mathbf{1})'U(\mathbf{1})$. That is, $u(\Phi(x)) = u(\mathbf{1})$ is the left Perron vector of the irreducible stochastic matrix $U(\mathbf{1})$ and so is element-wise positive. \square

Lemma A.12. *Let $\Psi(x) = [\Phi(1/x_1, \dots, 1/x_n)]^{-1}$. Then, $\lim_{t \rightarrow \infty} \Psi(\mu_t(\theta)) = \Psi^*$ with \mathbb{P}^θ -probability one.*

Proof. Let $\nu_{it}(\theta) = 1/\mu_{it}(\theta)$. The belief update rule (1) of agent i implies that

$$\nu_{it+1}(\theta) = \frac{\sum_{\hat{\theta}} \ell_i^\theta(\omega_{it+1}) f_i(\mu_t)(\hat{\theta})}{\ell_i^\theta(\omega_{it+1}) f_i(\mu_t)(\theta)},$$

and therefore, $\mathbb{E}_t^\theta[\nu_{it+1}(\theta)] = 1/f_i(\mu_t)(\theta)$. Hence, by (11), $\mathbb{E}_t^\theta[\nu_{it+1}(\theta)] = \sum_{\hat{\theta} \in \Theta} \psi_i(\mu_t(\hat{\theta}))/\psi_i(\mu_t(\theta))$. On the other hand, [Lemma A.9](#) guarantees that $\psi_i(\mu_t(\hat{\theta})) \leq \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1})\mu_{jt}(\hat{\theta})$. As a result, $\sum_{\hat{\theta} \in \Theta} \psi_i(\mu_t(\hat{\theta})) \leq \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) = 1$, where the equality is a consequence of the fact that ψ_i is homogeneous of degree 1. Therefore,

$$\mathbb{E}_t^\theta[\nu_{it+1}(\theta)] \leq \frac{1}{\psi_i(\mu_t(\theta))} = \phi_i(\nu_t(\theta)). \quad (49)$$

Writing the above inequality in vector form, we obtain, $\mathbb{E}_t^\theta[\nu_{t+1}(\theta)] \leq \phi(\nu_t(\theta))$. Furthermore, the fact that Φ is increasing (established in Lemma A.11) guarantees that $\Phi(\mathbb{E}_t^\theta[\nu_{t+1}(\theta)]) \leq \Phi(\phi(\nu_t(\theta))) = \Phi(\nu_t(\theta))$, where the equality holds by definition of Φ . Finally, concavity of Φ , also established in Lemma A.11, alongside Jensen's inequality, guarantees that $\mathbb{E}_t^\theta \Phi(\nu_{t+1}(\theta)) \leq \Phi(\nu_t(\theta))$. In other words, each element of the vector $\Phi(\nu_t(\theta))$ is a supermartingale. Since Φ is lower bounded, it converges almost surely, which in turn implies that $\Psi(\mu_t(\theta))$ converges to some $\Psi^* \in (0, 1]^n$ almost surely. \square

Lemma A.13. *If θ is the underlying state, then $\mu_{it}(\theta)$ remains bounded away from zero almost surely for all agents i . Furthermore,*

$$\psi_i(\mu_t(\hat{\theta})) - \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1})\mu_{jt}(\hat{\theta}) \rightarrow 0 \quad (50)$$

for all $\hat{\theta} \in \Theta$ as $t \rightarrow \infty$ with \mathbb{P}^θ -probability one.

Proof. To prove the first statement, recall from (49) that $\mathbb{E}_t^\theta[\nu_{it+1}(\theta)] \leq \phi_i(\nu_t(\theta))$. Also recall from Lemma A.10 that ϕ_i is homogeneous of degree 1 and weakly increasing, with a logarithmic curvature that is less than or equal to 1 throughout its domain. Therefore, Lemma A.9 implies that $\mathbb{E}_t^\theta[\nu_{it+1}(\theta)] \leq \sum_{j=1}^n \phi_i^{(j)}(\mathbf{1})\nu_{jt}(\theta)$. On the other hand, recall from the proof of Lemma A.11 that $u_i(\mathbf{1}) = q^{(i)}(\mathbf{1}) > 0$ for all agents i . Therefore, multiplying both sides of the inequality by $q^{(i)}(\mathbf{1})$, summing over all agents i , and using (48) implies that

$$\mathbb{E}_t^\theta \left[\sum_{i=1}^n q^{(i)}(\mathbf{1})\nu_{it+1}(\theta) \right] \leq \sum_{j=1}^n \sum_{i=1}^n q^{(i)}(\mathbf{1})\phi_i^{(j)}(\mathbf{1})\nu_{jt}(\theta) = \sum_{j=1}^n q^{(j)}(\mathbf{1})\nu_{jt}(\theta).$$

Thus, by the martingale convergence theorem, the expression $\sum_{j=1}^n q^{(j)}(\mathbf{1})\nu_{jt}(\theta)$ converges almost surely as $t \rightarrow \infty$. The fact that $q^{(j)}(\mathbf{1})$ is strictly positive for all j in turn implies that $\nu_{jt}(\theta)$ remains bounded \mathbb{P}^θ -almost surely for all j . As a result, $\mu_{jt}(\theta)$ is bounded away from zero for all j as $t \rightarrow \infty$.

To establish (50), recall that i 's belief dynamics are given by $\mu_{it+1} = \text{BU}(f_i(\mu_t); \omega_{it+1})$. Therefore,

$$\mu_{it+1}(\hat{\theta}) = \frac{f_i(\mu_t)(\hat{\theta})\ell_i^{\hat{\theta}}(\omega_{it+1})}{\sum_{\tilde{\theta} \in \Theta} f_i(\mu_t)(\tilde{\theta})\ell_i^{\tilde{\theta}}(\omega_{it+1})}$$

for all $\hat{\theta} \in \Theta$. Since f_i belongs to the class of weakly-separable learning rules, there exists a homogeneous function ψ_i such that $f_i(\mu_t)(\hat{\theta})/f_i(\mu_t)(\tilde{\theta}) = \psi_i(\mu_t(\hat{\theta}))/\psi_i(\mu_t(\tilde{\theta}))$ for all $\hat{\theta}, \tilde{\theta} \in \Theta$. Hence,

$$\mu_{it+1}(\hat{\theta}) = \psi_i(\mu_t(\hat{\theta})) + \left(\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}(\omega_{it+1})} - 1 \right) \psi_i(\mu_t(\hat{\theta})), \quad (51)$$

where $m_{it}(\omega_{it+1}) = \sum_{\tilde{\theta} \in \Theta} \ell_i^{\tilde{\theta}}(\omega_{it+1})\psi_i(\mu_t(\tilde{\theta}))$. Applying function Ψ_i defined in Lemma A.12 to both sides of (51), evaluating the result at the underlying state θ , using the mean value theorem, and noting that $\Psi(\psi(x)) = \Psi(x)$, we have

$$\Psi_i(\mu_{t+1}(\theta)) = \Psi_i(\mu_t(\theta)) + \sum_{j=1}^n \Psi_i^{(j)}(z_{t+1})\psi_j(\mu_t(\theta)) \left(\frac{\ell_j^\theta(\omega_{jt+1})}{m_{jt}(\omega_{jt+1})} - 1 \right),$$

where z_{t+1} is a point on the line segment connecting $\psi(\mu_t(\theta))$ to $\mu_{t+1}(\theta)$. Lemma A.12 guarantees that $\Psi_i(\mu_{t+1}(\theta)) - \Psi_i(\mu_t(\theta)) \rightarrow 0$ for all i as $t \rightarrow \infty$ almost surely. Therefore,

$$\sum_{j=1}^n \Psi_i^{(j)}(z_{t+1}) \psi_j(\mu_t(\theta)) \left(\frac{\ell_j^\theta(\omega_{jt+1})}{m_{jt}(\omega_{jt+1})} - 1 \right) \rightarrow 0 \quad \mathbb{P}^\theta\text{-a.s.} \quad (52)$$

On the other hand, the fact that $\mu_t(\theta)$ remains bounded away from zero almost surely as $t \rightarrow \infty$ guarantees that $\psi_j(\mu_t(\theta))$ and hence z_{t+1} — which belongs to the line segment connecting $\psi(\mu_t(\theta))$ to $\mu_{t+1}(\theta)$ — remain bounded away from zero almost surely as $t \rightarrow \infty$. Thus, by Lemma A.11, $\Psi_i^{(j)}(z_{t+1})$ is also bounded away from zero for all j . The juxtaposition of these observations with (52) therefore implies that $\ell_j^\theta(\omega_{jt+1})/m_{jt}(\omega_{jt+1}) \rightarrow 1$ almost surely for all j . Hence, by the dominated convergence theorem for conditional expectation, we have, $\mathbb{E}_t^\theta |\ell_j^\theta(\omega_{jt+1})/m_{jt}(\omega_{jt+1}) - 1| \rightarrow 0$ as $t \rightarrow \infty$. Expressing the conditional expectation as a sum and using the assumption that $\ell_i^\theta(\omega) > 0$ for all ω implies that

$$\ell_j^\theta(\omega) - \sum_{\tilde{\theta} \in \Theta} \ell_j^{\tilde{\theta}}(\omega) \psi_j(\mu_t(\tilde{\theta})) \rightarrow 0 \quad \mathbb{P}^\theta\text{-a.s.} \quad (53)$$

for all $\omega \in S$. Summing both sides of the above equation over ω , we obtain

$$\sum_{\omega \in S} \sum_{\tilde{\theta} \in \Theta} \ell_j^{\tilde{\theta}}(\omega) \psi_j(\mu_t(\tilde{\theta})) - 1 \rightarrow 0. \quad (54)$$

On the other hand, part (b) of Lemma A.9 guarantees that

$$\psi_i(\mu_t(\hat{\theta})) \leq \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) \mu_{jt}(\hat{\theta}). \quad (55)$$

Multiplying both sides of the above inequality by $\ell_i^{\hat{\theta}}(\omega)$ and summing over $\hat{\theta} \in \Theta$ and $\omega \in S$ leads to

$$\sum_{\omega \in S} \sum_{\hat{\theta} \in \Theta} \ell_i^{\hat{\theta}}(\omega) \psi_i(\mu_t(\hat{\theta})) \leq \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) \sum_{\hat{\theta} \in \Theta} \mu_{jt}(\hat{\theta}) \sum_{\omega \in S} \ell_i^{\hat{\theta}}(\omega) = 1,$$

where the last equality is a consequence of the fact that $\sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) = 1$. The juxtaposition of the above inequality with (54) implies that (55) holds as an equality as $t \rightarrow \infty$, thus establishing (50). \square

Lemma A.14. *Let θ denote the underlying state. Then,*

$$\lim_{t \rightarrow \infty} \sum_{\tilde{\theta} \in \Theta} \psi_i(\mu_t(\tilde{\theta})) \prod_{r=1}^k \ell_i^{\tilde{\theta}}(\omega_r) / \ell_i^\theta(\omega_r) = 1$$

with \mathbb{P}^θ -probability one for all integers k and all collections of signals $(\omega_1, \dots, \omega_k) \in S^k$.

Proof. Equation (53) establishes the lemma for $k = 1$. We prove the statement for $k > 1$ using an inductive argument on k , the details of which are provided in Online Appendix B. \square

Lemma A.15. *Let θ denote the underlying state of the world. If $\hat{\theta} \in \Theta$ is such that $\ell_i^{\hat{\theta}} \neq \ell_i^\theta$, then*

$$\lim_{t \rightarrow \infty} \psi_i(\mu_t(\hat{\theta})) = 0 \quad \mathbb{P}^\theta\text{-almost surely.} \quad (56)$$

Proof. Let $\hat{\Theta}_i$ denote the set of states $\hat{\theta} \neq \theta$ for which $\ell_i^{\hat{\theta}} \neq \ell_i^\theta$. Recall from Lemma A.14 that $\sum_{\hat{\theta} \in \hat{\Theta}_i} \psi_i(\mu_t(\hat{\theta})) \prod_{r=1}^k \ell_i^{\hat{\theta}}(\omega_r) / \ell_i^\theta(\omega_r) + \sum_{\hat{\theta} \notin \hat{\Theta}_i} \psi_i(\mu_t(\hat{\theta})) \rightarrow 1$ with \mathbb{P}^θ -probability one for all integers k and all collections of signals $(\omega_1, \dots, \omega_k) \in S^k$. On the other hand, summing both sides of (50) over $\hat{\theta} \in \Theta$ implies that $\sum_{\hat{\theta} \in \Theta} \psi_i(\mu_t(\hat{\theta})) - 1 \rightarrow 0$ almost surely. Hence,

$$\sum_{\hat{\theta} \in \hat{\Theta}_i} \psi_i(\mu_t(\hat{\theta})) \left(1 - \prod_{r=1}^k \frac{\ell_i^{\hat{\theta}}(\omega_r)}{\ell_i^\theta(\omega_r)} \right) \rightarrow 0 \quad (57)$$

\mathbb{P}^θ -almost surely for all integers k and all collections of signals $(\omega_1, \dots, \omega_k) \in S^k$.

Lemma 4 of [Jadbabaie et al. \(2012\)](#) guarantees that there exist a constant $\beta > 0$, a positive integer \tilde{k} , and a collection of signals $(\tilde{\omega}_1, \dots, \tilde{\omega}_{\tilde{k}})$ such that $1 - \prod_{r=1}^{\tilde{k}} \ell_i^{\hat{\theta}}(\tilde{\omega}_r) / \ell_i^\theta(\tilde{\omega}_r) > \beta$ for all $\hat{\theta} \in \hat{\Theta}_i$. This observation, alongside the fact that (57) has to hold for all integers k and any arbitrary sequence of signals $(\omega_1, \dots, \omega_k)$, therefore implies that $\psi_i(\mu_t(\hat{\theta})) \rightarrow 0$ with \mathbb{P}^θ -probability one for all $\hat{\theta} \in \hat{\Theta}_i$. \square

Proof of Theorem 7(a). Let θ denote the underlying state and consider an arbitrary $\hat{\theta} \neq \theta$. Since agents' information structures are collectively informative, there exists an agent i such that $\ell_i^{\hat{\theta}} \neq \ell_i^\theta$. Therefore, Lemma A.15 guarantees that the agent-state pair $(i, \hat{\theta})$ satisfies (56). Thus, by equation (50), $\lim_{t \rightarrow \infty} \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) \mu_{jt}(\hat{\theta}) = 0$. Since i 's learning rule satisfies monotonicity, it must be the case that $\psi_i^{(j)}(\mathbf{1}) > 0$ for all $j \in N_i$, and as a result $\lim_{t \rightarrow \infty} \mu_{jt}(\hat{\theta}) = 0$ with \mathbb{P}^θ -probability one. Subsequently, (51) implies that $\lim_{t \rightarrow \infty} \psi_j(\mu_t(\hat{\theta})) = 0$ with \mathbb{P} -probability one for all $j \in N_i$. In other words, the fact that the agent-state pair $(i, \hat{\theta})$ satisfies (56) is sufficient to guarantee that the agent-state pair $(j, \hat{\theta})$ also satisfies (56) for all $j \in N_i$. Repeating the same argument and using the fact that the social network is strongly connected guarantees that (56) is satisfied for all agents in the social network, which in turn — using (50) and monotonicity — implies that $\mu_j(\hat{\theta}) \rightarrow 0$ with \mathbb{P}^θ -probability one as $t \rightarrow \infty$ for all j . Since $\hat{\theta} \neq \theta$ was arbitrary, this means that all agents learn the underlying state almost surely. \square

Proof of Proposition 2

Recall from Lemma A.1 that when $\lim_{t \rightarrow \infty} t\lambda_t = \infty$, there exists a unique sequence of probability vectors v_t uniformly lower bounded away from zero such that $v'_{t+1} A_t A_{t-1} \dots A_\tau = v'_\tau$ for all $t \geq \tau$. It is immediate that the above equation reduces to (14) when $\tau = t$. Furthermore, the fact that the sequence v_t is uniformly bounded away from zero implies that $\liminf_{t \rightarrow \infty} v_{it} > 0$ for all agents i . \square

Proof of Theorem 8

Let θ denote the underlying state and fix an arbitrary $\hat{\theta} \neq \theta$. Also let $x_{it} = \log(\mu_{it}(\theta) / \mu_{it}(\hat{\theta}))$. Equations (29) and (31) imply that for any agent j , $\liminf_{t \rightarrow \infty} \frac{1}{t} x_{jt} = \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^n v_{i\tau} h_i(\theta, \hat{\theta})$ almost surely, where we are using the fact that $\sum_{j=1}^n v_{jt} = 1$. Moreover, recall from Theorem 4 that all agents learn the underlying state almost surely, that is, $\lim_{t \rightarrow \infty} \mu_{jt}(\theta) = 1$ with probability one. As a result, $\liminf_{t \rightarrow \infty} \frac{1}{t} \log \mu_{jt}(\hat{\theta}) = -\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^n v_{i\tau} h_i(\theta, \hat{\theta})$ with probability one. Finally, the fact that the sum of the beliefs $e_{jt}^\theta = \sum_{\hat{\theta} \neq \theta} \mu_{jt}(\hat{\theta})$ that agent j assigns to states other than θ satisfies $\max_{\hat{\theta} \neq \theta} \mu_{jt}(\hat{\theta}) \leq e_{jt}^\theta \leq (|\Theta| - 1) \max_{\hat{\theta} \neq \theta} \mu_{jt}(\hat{\theta})$ implies that $\gamma_i^\theta = \liminf_{t \rightarrow \infty} \frac{1}{t} |\log e_{jt}^\theta| = \min_{\hat{\theta} \neq \theta} \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^n v_{i\tau} h_i(\theta, \hat{\theta})$ almost surely. \square

Proof of Proposition 3

We first state a lemma, the proof of which is provided in the Online Appendix.

Lemma A.16. *Let $b_{t,\tau} = 1 - \prod_{r=\tau}^{t-1} (1 - r^{-\alpha})$. If $\alpha \in (1, 2)$, then $\lim_{t \rightarrow \infty} t^{\alpha-2} b_{t,\tau} = 0$ and $\lim_{t \rightarrow \infty} t^{\alpha-2} \sum_{\tau=1}^{t-1} b_{t,\tau} = 1/(2 - \alpha)$.*

Next, we show that agent i learns the underlying state θ . Consider an arbitrary $\hat{\theta} \neq \theta$ and let $x_{jt} = \log(\mu_{jt}(\theta)/\mu_{jt}(\hat{\theta}))$ and $y_{jt}(\omega_{jt}) = \log(\ell_j^\theta(\omega_{jt})/\ell_j^{\hat{\theta}}(\omega_{jt}))$. By (28), $x_t = \sum_{\tau=1}^{t-1} A_{t-1} \dots A_{\tau+1} A_\tau y_\tau(\omega_\tau) + y_t(\omega_t)$, where $A_t = (1 - t^{-\alpha})I + (t^{-\alpha}/n)\mathbf{1}\mathbf{1}'$. Therefore, $x_{it} = \sum_{\tau=1}^t (1 - b_{t,\tau})y_{i\tau}(\omega_{i\tau}) + \sum_{\tau=1}^{t-1} b_{t,\tau}\tilde{y}_\tau(\omega_\tau)$, where $\tilde{y}_\tau(\omega_\tau) = \frac{1}{n} \sum_{j=1}^n y_{j\tau}(\omega_{j\tau})$, $b_{t,\tau} = 1 - \prod_{r=\tau}^{t-1} (1 - r^{-\alpha})$ for $t > \tau$, and $b_{t,t} = 0$. As a result, by Lemma A.16,

$$\lim_{t \rightarrow \infty} t^{\alpha-2} x_{it} = \lim_{t \rightarrow \infty} t^{\alpha-2} \sum_{\tau=1}^t (1 - b_{t,\tau})y_{i\tau}(\omega_{i\tau}) + \frac{1}{2 - \alpha} \mathbb{E}^\theta[\tilde{y}_\tau(\omega_\tau)] + \lim_{t \rightarrow \infty} t^{\alpha-2} \sum_{\tau=1}^{t-1} b_{t,\tau}(\tilde{y}_\tau(\omega_\tau) - \mathbb{E}^\theta[\tilde{y}_\tau(\omega_\tau)]).$$

We next show that the last term on the right-hand side above is equal to zero almost surely. Recall from Lemma A.16 that $\lim_{t \rightarrow \infty} t^{\alpha-2} b_{t,\tau} = 0$ and $\lim_{t \rightarrow \infty} t^{\alpha-2} \sum_{\tau=1}^{t-1} b_{t,\tau} = \frac{1}{2-\alpha}$. Furthermore, note that $\max_{\tau \leq t} t^{\alpha-2} b_{t,\tau} = t^{\alpha-2}$. Thus, Theorem 2 of Pruitt (1966) implies that the weighted sum $t^{\alpha-2} \sum_{\tau=1}^{t-1} b_{t,\tau}(\tilde{y}_\tau(\omega_\tau) - \mathbb{E}^\theta[\tilde{y}_\tau(\omega_\tau)])$ of centered i.i.d. random variables converges to zero almost surely as $t \rightarrow \infty$. Similarly, it is immediate that $t^{\alpha-2} \sum_{\tau=1}^{t-1} b_{t,\tau}(y_{i\tau}(\omega_{i\tau}) - \mathbb{E}^\theta[y_{i\tau}(\omega_{i\tau})])$ converges to zero almost surely. Hence, $\lim_{t \rightarrow \infty} t^{\alpha-2} x_{it} = \lim_{t \rightarrow \infty} t^{\alpha-2} \sum_{\tau=1}^t y_{i\tau}(\omega_{i\tau}) + \frac{1}{2-\alpha} \mathbb{E}^\theta[\tilde{y}_\tau(\omega_\tau)] - \frac{1}{2-\alpha} \mathbb{E}^\theta[y_{i\tau}(\omega_{i\tau})]$. By the strong law of large numbers, $\lim_{t \rightarrow \infty} t^{-1} \sum_{\tau=1}^t y_{i\tau}(\omega_{i\tau}) = \mathbb{E}^\theta[y_{it}(\omega_{it})] = h_i(\theta, \hat{\theta})$. Therefore,

$$\lim_{t \rightarrow \infty} t^{\alpha-2} \log \frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} = h_i(\theta, \hat{\theta}) \lim_{t \rightarrow \infty} t^{\alpha-1} + \frac{1}{n(2 - \alpha)} \sum_{j \neq i} h_j(\theta, \hat{\theta}) \quad (58)$$

almost surely. The assumption that agents' signal structures are collectively informative guarantees that the right-hand side of the above equation is strictly positive, thus guaranteeing that $\lim_{t \rightarrow \infty} x_{it} = \infty$ almost surely, which means that $\lim_{t \rightarrow \infty} \mu_{it}(\hat{\theta}) = 0$. Since $\hat{\theta} \neq \theta$ was an arbitrary state, this establishes that agent i learns the underlying state, i.e., $\lim_{t \rightarrow \infty} \mu_{it}(\theta) = 1$ with probability one.

To determine the rate of learning, note that $\max_{\hat{\theta} \neq \theta} \mu_{it}(\hat{\theta}) \leq e_{it}^\theta \leq (|\Theta| - 1) \max_{\hat{\theta} \neq \theta} \mu_{it}(\hat{\theta})$, which implies that $\lim_{t \rightarrow \infty} t^{\alpha-2} |\log e_{it}^\theta| = \min_{\hat{\theta} \neq \theta} \lim_{t \rightarrow \infty} t^{\alpha-2} |\log \mu_{it}(\hat{\theta})|$. This equality, in conjunction with (58) and the fact that $\lim_{t \rightarrow \infty} \mu_{it}(\theta) = 1$, leads to

$$\lim_{t \rightarrow \infty} t^{\alpha-2} |\log e_{it}^\theta| = \min_{\hat{\theta} \neq \theta} \left\{ h_i(\theta, \hat{\theta}) \lim_{t \rightarrow \infty} t^{\alpha-1} + \frac{1}{n(2 - \alpha)} \sum_{j \neq i} h_j(\theta, \hat{\theta}) \right\}.$$

The assumption that agent i faces an identification problem in isolation guarantees that there exists a state $\hat{\theta} \neq \theta$ such that $h_i(\theta, \hat{\theta}) = 0$, whereas the assumption that agents do not face an identification problem collectively guarantees that there exists an agent $j \neq i$ such that $h_j(\theta, \hat{\theta}) > 0$. Taken together, these observations imply that the right-hand side above is finite and strictly positive almost surely. \square

References

- Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), “Bayesian learning in social networks.” *The Review of Economic Studies*, 78, 1201–1236.
- Banerjee, Abhijit (1992), “A simple model of herd behavior.” *The Quarterly Journal of Economics*, 107, 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), “A theory of fads, fashion, custom, and cultural change as information cascades.” *Journal of Political Economy*, 100, 992–1026.
- Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri (2016), “Testing models of social learning on networks: Evidence from a lab experiment in the field.” NBER Working paper No. 21468.
- Chatterjee, S. and Eugene Seneta (1977), “Towards consensus: Some convergence theorems on repeated averaging.” *Journal of Applied Probability*, 89–97.
- DeGroot, Morris H. (1974), “Reaching a consensus.” *Journal of American Statistical Association*, 69, 118–121.
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffery Zwiebel (2003), “Persuasion bias, social influence, and unidimensional opinions.” *The Quarterly Journal of Economics*, 118, 909–968.
- Eyster, Erik and Matthew Rabin (2010), “Naïve herding in rich-information settings.” *American Economic Journal: Microeconomics*, 2, 221–243.
- Eyster, Erik and Matthew Rabin (2014), “Extensive imitation is irrational and harmful.” *The Quarterly Journal of Economics*, 129, 1861–1898.
- Golub, Benjamin and Matthew O. Jackson (2010), “Naïve learning in social networks and the wisdom of crowds.” *American Economic Journal: Microeconomics*, 2, 112–149.
- Golub, Benjamin and Matthew O. Jackson (2012), “How homophily affects the speed of learning and best-response dynamics.” *The Quarterly Journal of Economics*, 127, 1287–1338.
- Golub, Benjamin and Evan Sadler (2016), “Learning in social networks.” In *The Oxford Handbook on the Economics of Networks* (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.), chapter 19, 504–542, Oxford University Press.
- Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz (2017), “Groupthink and the failure of information aggregation in large groups.” Working paper, available at <https://ssrn.com/abstract=2541707>.
- Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012), “Non-Bayesian social learning.” *Games and Economic Behavior*, 76, 210–225.

- Jadbabaie, Ali, Elchanan Mossel, and M. Amin Rahimian (2017), “Bayesian group decisions: Algorithms and complexity.” Working paper, available at <https://ssrn.com/abstract=3055954>.
- Kolmogorov, Andrey (1936), “Zur Theorie der Markoffschen Ketten.” *Mathematische Annalen*, 112, 155–160.
- Lehrer, Keith and Carl Wagner (1981), *Rational Consensus in Science and Society, A Philosophical and Mathematical Study*. D. Reidel Publishing Company, Dordrecht, Holland.
- Levy, Gilat and Ronny Razin (2016), “Information diffusion in networks with the Bayesian Peer Influence heuristic.” Working paper.
- Li, Wei and Xu Tan (2017), “Learning in local networks.” Working paper, available at <https://www.aeaweb.org/conference/2017/preliminary/paper/B36rRFsi>.
- Liu, Qipeng, Aili Fang, Lin Wang, and Xiofan Wang (2014), “Social learning with time-varying weights.” *Journal of System Science and Complexity*, 27, 581–593.
- McCuaig, William (1984), “A simple proof of Menger’s theorem.” *Journal of Graph Theory*, 8, 427–429.
- Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2016), “Foundations of non-Bayesian social learning.” Columbia Business School Working Paper No. 15-95.
- Mossel, Elchanan, Noah Olsman, and Omer Tamuz (2016), “Efficient Bayesian learning in social networks with Gaussian estimators.” In *Annual Allerton Conference on Communication, Control, and Computing*, 425–432, Allerton House, Illinois.
- Mucci, Anthony G. (1976), “Another martingale convergence theorem.” *Pacific Journal of Mathematics*, 64, 539–541.
- Nussbaum, Roger D. (1988), *Hilbert’s Projective Metric and Iterated Nonlinear Maps*, volume 75. American Mathematical Society.
- Pruitt, William E. (1966), “Summability of independent random variables.” *Journal of Mathematics and Mechanics*, 15, 769–776.
- Rahimian, Mohammad Amin, Pooya Molavi, and Ali Jadbabaie (2014), “(Non-)Bayesian learning without recall.” In *Proceedings of the 53rd IEEE Conference on Decision and Control*, 5730–5735.
- Rota, Gian-Carlo (1964), “On the foundations of combinatorial theory I: Theory of Möbius functions.” *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 2, 340–368.
- Seneta, Eugene (1981), *Non-negative Matrices and Markov Chains*, second edition. Springer, New York.
- Steele, J. Michael (1994), “Le Cam’s inequality and Poisson approximation.” *The American Mathematical Monthly*, 101, 48–54.