

Foundations of non-Bayesian Social Learning^{*}

Pooya Molavi[†]

Alireza Tahbaz-Salehi[‡]

Ali Jadbabaie[§]

August 2016

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 establish that as long as imperfect recall represents the only point of departure from Bayesian rationality, agents’ social learning rules take a log-linear form. Our result thus establishes that all other non-Bayesian models of social learning (such as the canonical model of DeGroot) deviate from Bayesian rationality in ways above and beyond the assumption of imperfect recall. We then obtain general long-run learning results that are not tied to the specific functional form of agents’ learning rules, thus identifying the fundamental forces that lead to learning, non-learning, and mislearning in social networks. Our characterization results establish that long-run beliefs are closely linked to (i) the curvature of agents’ social learning rule and (ii) whether agents’ initial tendencies are amplified or moderated as a result of their social interactions.

Keywords: social networks, non-Bayesian learning, bounded rationality.

JEL Classification: D83, D85.

^{*}We thank Daron Acemoglu, Kostas Bimpikis, Emily Breza, Arun Chandrasekhar, Matt Elliott, Ben Golub, Sanjeev Goyal, Ilan Lobel, Mohamed Mostagir, Pietro Ortoleva, Andy Postlewaite, Amin Rahimian, Matthew Rabin, Kamiar Rahnama-Rad, Evan Sadler, Alvaro Sandroni, Juuso Toikka, Andrea Vedolin, Ricky Vohra, Alex Wolitzky, and seminar participants at Cambridge, Cleveland Fed, Columbia GSB, IESE, MIT, Penn, Pompeu Fabra, 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 Air Force Office of Scientific Research (MURI Award No. FA9550-10-1-0567) and the Army Research Office (MURI Award No. W911NF-12-1-0509).

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

[‡]Columbia Business School, Columbia University, alirezat@gsb.columbia.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 economic mechanisms.

In this paper, we address these challenges by taking an alternative 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 enables us to not only provide a systematic way of capturing deviations from Bayesian inference, but also to reveal fundamental forces that are central to information aggregation but may be obscured by the non-generic 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. Our results illustrate that asymptotic learning relies on knife-edge conditions that are satisfied in DeGroot's model but are bound to be violated in all but few of its generalizations.

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 the manifestation 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.

As our starting point, we focus on a case with a close relationship to Bayesian learning by imposing three restrictions on agents' social learning rules — other than imperfect recall — that are satisfied by Bayesian agents under fairly general conditions. First, we assume that agents' social learning rules are label neutral (LN), in the sense 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.

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 ratio of her and her neighbors' beliefs in the previous period. This learning rule, which we refer to as *log-linear learning*, serves as the benchmark non-Bayesian learning rule for the rest of our results.

Our representation theorem establishes that so long as imperfect recall is the only point of departure from Bayesian rationality, the learning rule must take the above-mentioned log-linear form. Furthermore, this result reveals that all other non-Bayesian models of social learning (such as DeGroot's model) deviate from Bayesian rationality in ways above and beyond the assumption of imperfect recall. 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, LN, 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. This result thus illustrates that DeGroot's model is the result of a double deviation from Bayesian rationality.

Given their different functional forms and distinct foundations, it is not surprising that agents who follow the log-linear and DeGroot's learning 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 identical implications for agents' beliefs in the long-run. Motivated by this observation, 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 social networks.

We achieve this by replacing IIA and separability with a weaker notion of independence and

obtaining a general class of learning rules that encompasses 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 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 (including DeGroot and log-linear models). We then establish that irrespective of the functional form used by the agents to incorporate their neighbors' information, long-run beliefs only depend on two specific characteristics of the learning rule: (i) the degree of homogeneity of ψ and (ii) its logarithmic curvature (defined as the curvature in the log-log scale), which measures the underlying learning rule's departure from the benchmark of log-linear learning.

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 homogenous function of degree greater than 1. This is due to the fact that with a degree of homogeneity greater than 1, agents overweigh evidence they encounter early on at the expense of the more recent pieces of information, thus resulting in what we refer to as *group polarization*: every round of social interaction reinforces individuals' initial tendencies and hence leads to asymptotic mislearning. We also obtain a diametrically opposite result by establishing that agents remain uncertain forever about the underlying state whenever ψ is a homogenous function of degree less than 1, as they downplay the already accumulated information in favor of their more recent observations, an effect which we refer to as *group depolarization*. These results thus imply that for asymptotic learning to be successful, it is necessary for the degree of homogeneity of ψ to be exactly equal to 1, guaranteeing that each agent assigns weights of the same order of magnitude to any independent piece of information. They also illustrate that asymptotic learning relies on knife-edge conditions that are satisfied in the linear learning model of DeGroot but are bound to be violated in all but few of its generalizations.

Despite its key role in ensuring that agents assign roughly equal weights to any independent piece of information, homogeneity of degree 1 is not sufficient for the asymptotic aggregation of information over the social network. As our next result, we show that information aggregation also requires a restriction on the curvature of agents' social learning rules. Formally, this is captured by requiring that the logarithmic curvature of ψ to fall 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 processes her neighbors' information (relative to the benchmark of log-linear learning) is so large that either she mislearns the underlying state or remains uncertain forever, even if ψ is homogenous of degree 1.

Taken together, our results provide a tight characterization of conditions that lead to asymptotic learning, non-learning, and mislearning. Crucially, these conditions are not tied to the specific functional form of agents' learning rules and hence, uncover the key forces that underpin information aggregation in social networks.

One key implication of our results is that whether dispersed information is successfully aggregated in the long-run is orthogonal to (i) the detailed structure of the social network and (ii)

how information is dispersed among different agents in the society. We thus end the paper by showing that these features are the key variables in determining the *rate* of information aggregation. More specifically, we show that the speed at which information is aggregated throughout the society has a simple analytical characterization in terms of the relative entropy of agents' signal structures and various notions of network centrality.

Related Literature 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).¹ 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 this approach in one significant way: instead of assuming functional forms for how agents incorporate their neighbors' opinions into their beliefs, we study a broader class of learning rules and determine the behavioral assumptions that underpin various non-Bayesian models of social learning. This approach enables us to (i) establish that as long as imperfect recall represents the only point of departure from Bayesian rationality, agents' social learning rules take a log-linear form; (ii) characterize the behavioral assumptions that underpin the linear learning mode of DeGroot; and most importantly, (iii) obtain novel and general long-run learning results that are not tied to the specific functional form of agents' learning rules.

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 (2016) depart from the Bayesian framework by assuming that each agent updates her beliefs as if her local neighborhood is the entire society.

The contrast between some of the predictions of Bayesian and non-Bayesian learning models has led to a growing empirical and experimental literature that aims to test the details of how agents aggregate information. For instance, Chandrasekhar, Larreguy, and Xandri (2016) conduct an experiment in rural India to test whether a variant of DeGroot model can outperform the Bayesian framework in describing learning in social networks. Relatedly, Grimm and Mengel (2015) use a lab experiment to study belief formation in social networks and observe that the DeGroot model

¹Some of the other contributions in this literature include Acemoglu, Ozdaglar, and ParandehGheibi (2010), Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012), and Banerjee, Breza, Chandrasekhar, and Möbius (2016).

²Other works in this literature include Bala and Goyal (1998), Smith and Sørensen (2000), Gale and Kariv (2003), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Mueller-Frank (2013), Acemoglu, Bimpikis, and Ozdaglar (2014), Lobel and Sadler (2015, 2016), Mossel, Sly, and Tamuz (2015), and Mueller-Frank and Pai (2016). See Golub and Sadler (2016) for a thorough survey of the social learning literature, covering both Bayesian and non-Bayesian paradigms.

does a better job in explaining individual choices compared to Bayesian learning, though it cannot explain changes in behavior induced by participants' information about the network structure.³ Despite its theoretical take, the approach taken in our paper contributes to this debate by illustrating the behavioral assumptions that underpin DeGroot and other non-Bayesian learning models. It suggests that testing the behavioral assumptions that underpin a given learning model can serve as an alternative (and complementary) approach to testing the model's predictions for agents' entire path of actions.

Our results on the rate of learning is related to the recent work of [Harel, Mossel, Strack, and Tamuz \(2015\)](#), who show that increased interactions between Bayesian agents can lower the speed of learning. [Mueller-Frank \(2013\)](#) shows that when Bayesian agents have access to a single piece of information, the speed of learning is determined by the diameter of the underlying social network. In contrast to his setting, agents in our model receive a stream of informative signals over time. As such, speed of information aggregation in our model is tightly linked to the relative entropies of individuals' signal structures. Finally, our results also align with those of [Jackson \(2008\)](#) and [Golub and Jackson \(2010, 2012\)](#) who show that agents' asymptotic beliefs in the DeGroot model is tightly linked to their eigenvector centralities, a statistic that captures the extent of each agent's influence on others. Generalizing these results, we show that with a constant flow of new information, the rate of learning depends not only on agents' eigenvector centralities, but also on how information is distributed throughout the social network as well as a second notion of centrality that captures how each agent is influenced by others.

Outline of the Paper 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. Section 4 is dedicated to the DeGroot learning model. We then focus on a more general class of social learning rules in Section 5 and characterize the conditions that lead to learning, non-learning, and mislearning in the long-run. The rate of learning is characterized in Section 6. Section 7 concludes. All proofs and some additional mathematical details are provided in the Appendix.

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 $t = 0$ according to a probability distribution with full support over some finite set Θ .

Even though the realized state remains unobservable to the individuals, they 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 $\omega_{it} \in S$ which is drawn according to

³Also see the recent works of [Banerjee, Chandrasekhar, Duflo, and Jackson \(2013\)](#), [Möbius, Phan, and Szeidl \(2015\)](#), [Mueller-Frank and Neri \(2015\)](#), [Eyster, Rabin, and Weizsäcker \(2015\)](#), and [Alatas, Banerjee, Chandrasekhar, Hanna, and Olken \(2016\)](#). See [Breza \(2016\)](#) and [Choi, Gallo, and Kariv \(2016\)](#) for surveys of field and lab experiments on social networks.

distribution $\ell_i^\theta \in \Delta S$. We assume that the signal space S is finite and that ℓ_i^θ has full support over S for all i and all $\theta \in \Theta$. The realized signals are independent across individuals and over time. Each agent may face an identification problem in the sense that she may not be able to distinguish between two states. However, agents' observations are *collectively* informative: for any distinct pair of states $\theta, \hat{\theta} \in \Theta$, there exists an agent i such that $\ell_i^\theta \neq \ell_i^{\hat{\theta}}$.

In addition to her private signals, each agent also observes the beliefs of a subset of other agents, to whom we refer 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 (j, i) is present from vertex j to vertex i if agent j can observe the beliefs of agent i . 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, in the sense that there exists a directed path from each vertex to any other. This assumption ensures that the information available to any given agent can potentially flow to other individuals in the social network.

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 of the world. 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 beliefs. 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 $f_{it} : \Delta\Theta^{|N_i|(t+1)} \rightarrow \Delta\Theta$, 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 beliefs.⁵

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

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 the actions whenever the action space is "rich" enough, so that an individual's actions fully reveal her beliefs. See [Eyster and Rabin \(2010\)](#) and [Ali \(2016\)](#) for a discussion of observational learning models in information-rich settings.

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.

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 $\tau \leq t - 1$ and does not depend on time index t .

The restriction imposed by imperfect recall represents a departure from Bayesian rationality. For instance, in the absence of a common prior, a Bayesian agent i can make inferences about j ’s latest private signal only by comparing j ’s current belief to her beliefs in the previous periods. 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 structure of the social network — require a high level of sophistication on the part of the agents. In contrast, under IR, agent i simply treats her neighbors’ most recent opinions as sufficient statistics for all the information available to them, while ignoring the rest of the history.⁶

We remark that imperfect recall is distinct from the notions of “persuasion bias” and “redundancy neglect” studied by DeMarzo, Vayanos, and Zwiebel (2003) and Eyster and Rabin (2010, 2014), according to which agents treat the information provided by their neighbors at each period as entirely novel and fail to account for the fact that some of that information may have already been incorporated into their own (or other agents’) beliefs in prior periods. In contrast to these notions, imperfect recall represents a behavioral bias whereby agents simply rely on the beliefs currently held by their neighbors and discard all the information buried in the rest of their observation histories.

⁶The 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 non-Bayesian learning models in the literature. Most notably, the DeGroot model and its different variations (e.g., Golub and Jackson (2010, 2012), Acemoglu et al. (2010), and Chandrasekhar et al. (2016)) rely on imperfect recall by assuming that agents only use the last period beliefs of their neighbors.

Consequently, depending on the entire path of beliefs, agents that suffer from imperfect recall may either under- or over-react to the information provided to them by their neighbors (compared to what is prescribed by Bayesian learning).

In the remainder of the paper, we supplement the assumption of imperfect recall by a variety of other restrictions on agents' behavior. By putting more structure on the social learning rules, these restrictions enable us to uncover the forces that shape agents' long-run beliefs in the presence of imperfect recall.

3 Log-Linear Learning

We start our analysis by considering the special case of the general framework introduced in the previous section that exhibits a minimal deviation from Bayesian updating. In particular, we impose three restrictions — other than imperfect recall — that are satisfied by Bayesian agents under fairly general conditions and 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 ratio of her and her neighbors' beliefs in the previous period. Besides its close relationship to Bayesian learning, the log-linear learning rule will serve as the natural benchmark for our results on agents' long-run beliefs in Section 5.

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

Consequently, any asymmetry in how an individual updates her opinion about different states is only due to asymmetries in her or her neighbors' subjective priors about those states — as opposed to how different states are labeled. It is not hard to see that, in the presence of common knowledge of Bayesian rationality, relabeling the underlying states has no bearing on how Bayesian agents process information.

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

Monotonicity. $f_{it}(\mu_i^t)(\theta)$ is strictly increasing in $\mu_{jt}(\theta)$ for all $j \in N_i$ and all $\theta \in \Theta$ whenever all beliefs are interior.

The rationale behind this assumption is as follows: keeping the history of observations μ_i^{t-1} fixed, agent i interprets an increase in $\mu_{jt}(\theta)$ as evidence that either (i) agent j has observed a private signal

in favor of θ at period t ; or that (ii) 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.

In our environment, monotonicity is consistent with Bayesian updating: *ceteris paribus*, the posterior belief of a Bayesian agent assigned to a given state is increasing in her neighbors' beliefs about that same state in the previous time period. This is in contrast to the canonical environment in the observational learning literature in which agents observe a single private signal and take actions sequentially. In particular, as argued by [Eyster and Rabin \(2014\)](#), Bayesian updating in such environments may entail a significant amount of "anti-imitative" behavior: agents may revise their beliefs *downwards* in response to an increase in beliefs by some of their predecessors. This disparity in behavior is due to the difference in how new information is revealed in the two environments. In an environment where individuals observe a single signal and take actions according to a pre-specified order, the information in j 's action is fully incorporated into the actions of all agents who observe her behavior. Consequently, in order to avoid double-counting the same piece of information, a Bayesian agent i may have to anti-imitate her predecessors whose signals are already incorporated into the actions of i 's other predecessors. In contrast, each agent in our environment observes an informative private signal in every period. Therefore, to properly account for the information in her neighbors' most recent private signals, Bayesian rationality requires each agent to monotonically increase her posterior belief in her neighbors' beliefs in the previous time period.

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

In other words, the conditional belief of agent i after aggregating her neighbors' opinions is identical to the belief obtained by aggregating her neighbors' conditional beliefs using the same social learning rule. Thus, i 's posterior belief conditional on $\bar{\Theta}$ exclusively depends on the history of her and her neighbors' beliefs on states in $\bar{\Theta}$ and is independent of beliefs assigned by any individual to $\theta \notin \bar{\Theta}$ in any of the previous time periods.⁷

Independence of irrelevant alternatives requires that, as far as agent i is concerned, her neighbors' beliefs about states in $\bar{\Theta}$ are sufficient statistics for their collective information regarding all $\theta \in \bar{\Theta}$. Put differently, once given access to her neighbors' beliefs about $\bar{\Theta}$, agent i does not change her opinions about any $\theta \in \bar{\Theta}$ if she learns her neighbors' beliefs about some $\hat{\theta} \notin \bar{\Theta}$. One can show that, for a wide-range of networks and information structures, IIA is consistent with Bayesian

⁷Note that IIA is trivially satisfied when $|\Theta| = 2$, and hence does not impose any restrictions on agents' social learning rules. Throughout the rest of the paper, we assume that Θ consists of at least three elements.

updating.⁸ This is due to the fact that Bayes' rule guarantees — almost tautologically — that the belief of a Bayesian agent on $\bar{\Theta}$ is a sufficient statistic for the likelihood of any $\theta \in \bar{\Theta}$ given all the information available to that agent.

3.1 Representation Theorem

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

Theorem 1. *If agents' social learning rules satisfy IR, LN, monotonicity, and IIA, there exist constants $a_{ij} > 0$ such that*

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

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. More importantly, given that Bayesian updating satisfies LN, monotonicity, and IIA, Theorem 1 also establishes that as long as imperfect recall is the only point of departure from Bayesian rationality, agents' social learning rules take the log-linear form of equation (3). As a consequence, this result implies that other non-Bayesian models of social learning in which agents interact with one another repeatedly (such as DeGroot's model) deviate from Bayesian rationality in ways above and beyond the assumption of imperfect recall.

It is instructive to elaborate on the role of each assumption in determining the functional form of the social learning rule in (3). First, note that imperfect recall requires i 's posterior beliefs 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 should only depend on her and her neighbors' likelihood ratios for those two states. Given that such independence should hold 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_{ij} do not depend on the pair of states θ and $\hat{\theta}$ under consideration. Finally, the non-negativity of these constants is an immediate implication of the monotonicity assumption.

We can now use the representation in Theorem 1 to characterize the dynamics of agents' beliefs over the social network.

Corollary 1. *If agents' social learning rules satisfy IR, LN, monotonicity, and IIA, then*

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

⁸The exception is the case in which beliefs on irrelevant states help agents gain additional information by allowing them to disentangle different observations that lead to similar beliefs.

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 ratio of her and her neighbors' beliefs in the previous time period, with constant a_{ij} representing the weight that i assigns to the belief of agent j in her neighborhood.

One immediate consequence of the log-linear nature of the learning rule in (4) is that whenever any of i 's neighbors rules out state θ as impossible, agent i would follow suit by assigning a belief of zero to θ in the next period. This property is an implication of IIA, according to which i 's posterior conditional beliefs have to be consistent with her neighbors' conditional beliefs in the previous period.

Note that besides positivity, the representation in (4) does not impose any restrictions on constants a_{ij} .⁹ This observation highlights that the assumption of imperfect recall is orthogonal to whether agents under- or over-weight their neighbors' opinions (relative to what is prescribed by Bayesian learning), as the belief update rule in (4) is consistent with IR for all a_{ij} . We also emphasize that the assumption that agents incorporate their private signals into their beliefs in a Bayesian fashion does not play a crucial rule in our characterization. More specifically, altering the way agents process the information content of their private signals only impacts how 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 so called "Bayesian Peer Influence" heuristic of Levy and Razin (2016), who show that individuals who treat their marginal information sources as (conditionally) independent must follow learning rule (3) with corresponding weights $a_{ij} = 1$ (also see Levy and Razin (2015)). Our representation Theorem 1 thus provides a distinct foundation for log-linear updating by demonstrating the key role played by the imperfect recall and IIA assumptions.

3.2 Information Aggregation

The fact that agents' signals are not publicly observable means that information about the underlying state of the world is dispersed throughout society. At the same time, the social network provides a potential channel over which each individual's private information can disseminate to others. In this subsection, we use our representation theorem to characterize the conditions under which the form of bounded rationality captured by imperfect recall can act as an impediment to efficient aggregation of information over the social network.

We say agents' social learning rules are *group polarizing* if there exist a pair of states $\theta \neq \hat{\theta}$ and a belief profile $\mu \in \Delta\Theta^n$ satisfying $\mu_i(\theta) \geq \mu_i(\hat{\theta})$ for all i such that

$$\frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} \geq \frac{\mu_i(\theta)}{\mu_i(\hat{\theta})} \quad (5)$$

⁹For example, these constants need not add up to one or any other constant.

for all agents i .¹⁰ In other words, in the presence of group polarization, an initial tendency of individual agents toward a given direction is (weakly) enhanced following social interactions (Isenberg, 1986; Sunstein, 2000). Likewise, we say social learning rules are *group depolarizing* if exchange of information with neighbors leads to the adoption of less extreme viewpoints relative to agents' prior opinions; that is, if there exist a pair of states $\theta \neq \hat{\theta}$ and a belief profile μ satisfying $\mu_i(\theta) \geq \mu_i(\hat{\theta})$ for all i such that (5) holds with the inequality sign reversed. As a natural extension of these concepts, we say social learning rules are *strictly* group polarizing or depolarizing if the corresponding inequalities hold strictly for all i .

We emphasize that the concepts of group polarization and depolarization only require the corresponding inequalities to hold for a particular pair of states and a single profile of beliefs (as opposed to all possible beliefs and states). In fact, it may indeed be the case that the collection of social learning rules exhibit (weak) group polarization and depolarization simultaneously, a property which we refer to as *non-polarization*. We have the following result:

Theorem 2. *Suppose agents' social learning rules satisfy IR, LN, monotonicity, and IIA.*

- (a) *If learning rules are strictly group polarizing, agents mislearn the state with positive probability.*
- (b) *If learning rules are strictly group depolarizing, agents remain uncertain forever.*
- (c) *If learning rules are non-polarizing, all agents learn the underlying state almost surely.*

Statement (a) establishes that as long as social interactions intensify agents' prior biases in at least one direction, individuals may assign probability one to a false state as $t \rightarrow \infty$. This is despite the fact that they have access to enough information to (collectively) uncover the underlying state. The intuition for this result is that in the presence of strict group polarization, certain constellation of opinions become self-reinforcing. Thus, if early signals are sufficiently misleading, agents may end up hearing echoes of their own voices, and as a result mislearn the state. Part (b), on the other hand, shows that in the presence of strict group depolarization, agents downplay the already accumulated information in favor of their more recent observations, and as a result remain uncertain about the underlying state forever.¹¹ Finally, part (c) of Theorem 2 shows that the information dispersed throughout the social network is efficiently aggregated as long as agents' social learning rules are non-polarizing. Thus, any such learning rule asymptotically coincides with Bayesian learning 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 redundancies in different information sources, and may be unaware of the intricate details of the social network.

¹⁰Note that by IR, agent i 's posterior belief only depends on her neighbors' beliefs in the previous time period. Hence, we can drop time index t and focus on the learning rules $f_i : \Delta\Theta^n \rightarrow \Delta\Theta$ that map a belief profile to a single belief.

¹¹In this sense, strict group depolarization is the collective counterpart to the "This Time is Different" bias (Collin-Dufresne, Johannes, and Lochstoer, 2015), whereby agents (as a group) fail to put enough weight on their own and their neighbors' past experiences, whereas strict group polarization is akin to the concept of "social confirmation bias" (Rabin and Schrag, 1999).

To see the intuition underlying the above result, note that as long as IIA is satisfied, a Bayesian agent assigns in any period an equal weight to any independent piece of information that has reached her. This property, however, is clearly violated by the learning rule in (4), implying that agents' beliefs depart from the prescriptions of Bayesian learning at any given finite time t . Nonetheless, when social learning rules are non-polarizing, the weight that agent i assigns to any independent piece of information remains finite and bounded away from zero as $t \rightarrow \infty$. Therefore, despite imperfect recall, each piece of information is eventually accounted for in the long-run (even if with unequal weights). This guarantees that all agents eventually learn the underlying state. In contrast, group polarization and depolarization act as impediments to learning exactly because a subset of signals are assigned weights that, respectively, grow unboundedly and converge to zero as $t \rightarrow \infty$.

A sufficient condition for non-polarization is for agents' learning rules to be *unanimous*, in the sense that $f_i(\mu, \dots, \mu) = \mu$ for all beliefs $\mu \in \Delta\Theta$ and all i . Under unanimous learning rules, each agent adopts the shared beliefs of her neighbors whenever they all agree with one another. We thus have the following corollary to Theorem 2(c):

Corollary 2. *Suppose agents' social learning rules satisfy IR, LN, monotonicity, and IIA. If learning rules are unanimous, then all agents learn the underlying state almost surely.*

We conclude this section by remarking that the key role played by group polarization and depolarization in shaping asymptotic outcomes is not specific to the log-linear learning rule in (4). Rather, as we show in Section 5, the predictions of Theorem 2 and Corollary 2 generalize to a large class of learning rules.

4 DeGroot Learning

A key implication of Theorem 1 is that any social learning rule that satisfies imperfect recall but is distinct from (3) has to violate either LN, IIA, or monotonicity, thus exhibiting a second departure from Bayesian rationality above and beyond imperfect recall.

One such model is the learning model of DeGroot (1974), which serves as the canonical model of non-Bayesian 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 non-negative and independent of the underlying state θ , label neutrality and monotonicity are also 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).

The juxtaposition of these observations with the fact that IIA is satisfied by Bayesian updating reveals the second dimension along which DeGroot learning deviates from Bayesian rationality. To further clarify the nature of this deviation, 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 all j , all $\tau \leq t$, and all $\hat{\theta} \neq \theta$.

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 3. *Suppose agents' learning rules satisfy IR, LN, monotonicity, and separability. Then, there exists a set of constants $a_{ij} > 0$ and $a_{i0} \geq 0$ such that*

$$f_{it}(\mu_i^t)(\theta) = a_{i0} + \sum_{j \in N_i} a_{ij} \mu_j(\theta) \quad (6)$$

for all $\theta \in \Theta$.

Therefore, replacing IIA with separability results in a learning rule according to which each agent's beliefs depend linearly on her neighbors' opinions in the previous time period, in line with DeGroot's model.¹² Thus, in addition to providing the key assumptions that underpin DeGroot learning, Theorem 3 also formalizes the point we made earlier: that DeGroot learning is the result of a double deviation from Bayesian rationality.

Not surprisingly, agents who follow the linear learning rule in (6) process their neighbors' information differently from those who follow the log-linear learning rule of (3). As an example, recall from the discussion following Corollary 1 that, under IIA, agent i rules out state θ whenever any of her neighbors do so — a property that also holds under Bayesian learning. In contrast, when learning rules are separable, agent i ends up with a positive posterior belief on θ even if some (but not all) of her neighbors rule out that state.

Our next theorem characterizes the conditions under which DeGroot learning leads to the long-run aggregation of information throughout the social network:

Theorem 4. *Suppose agents' social learning rules satisfy IR, LN, monotonicity, and separability. Then, all agents learn the underlying state almost surely if and only if their learning rules are unanimous.*

Contrasting this result with Corollary 2 highlights that if social learning rules are unanimous, the DeGroot and log-linear learning rules result in asymptotic learning, despite the fact that they have different behavioral foundations and may lead to different sets of beliefs at any given finite time.¹³ In the next section, we show that the convergence of these two learning rules to the same limit is no coincidence and is a much more general phenomenon.

We end this section with a remark on terminology. Throughout the paper, we use the term DeGroot learning to refer to a model in which agents set their beliefs as a weighted average of

¹²Lehrer and Wagner (1981) provide a characterization similar to ours, albeit under a different set of restrictions.

¹³The canonical DeGroot model frequently used in the literature assumes that agents set their beliefs as a convex combination of their neighbors' beliefs, and thus imposes unanimity by design. In our setting, this translates to assuming that $\sum_{j \in N_i} a_{ij} = 1$ and $a_{i0} = 0$ for all agents i .

their neighbors' *beliefs*. This is consistent with many of the papers in the literature, including the original work of DeGroot (1974) as well as some of the subsequent works, such as Acemoglu et al. (2010), Jadbabaie et al. (2012), and Banerjee et al. (2016). At the same time, some other papers, such as DeMarzo et al. (2003), use the term DeGroot learning to refer to a model where agents linearly combine their neighbors' *point estimates*. Despite their seemingly similar natures, the two models are not identical: averaging two probability distributions does not result in a distribution whose mean is equal to the average of the two original means. In fact, our representation theorems highlight that the two models impose fundamentally different assumptions on how agents process information. On the one hand, Theorem 3 shows that agents average their neighbors' beliefs if their learning rules are separable. On the other hand, our characterization result in Theorem 1 establishes that, under IIA, agents follow a log-linear updating rule, which reduces to taking weighted averages of their neighbors' point estimates when all signals and beliefs are normally distributed.¹⁴ This observation also illustrates that more generally (e.g., when Θ is discrete as in our model), a learning rule that is based on averaging of point estimates violates both IIA and separability. Irrespective of terminology, our characterization results clarify the distinction in the behavioral assumptions that underpin each model.

5 Long-Run Learning, Mislearning, and non-Learning

Theorems 1 and 3 reveal that the functional form of agents' social learning rules is closely tied to the underlying notion of independence, that is, whether and how agent j 's beliefs about state θ impact i 's opinions about $\hat{\theta} \neq \theta$ in the next period. At the same time, Corollary 2 and Theorem 4 establish that, despite their distinct foundations, unanimity serves as a sufficient condition for asymptotic learning for both the DeGroot and log-linear learning models.

In this section, we develop a taxonomy of results on asymptotic learning, non-learning, and mislearning that is not tied to the specific functional form of agents' learning rules, thus enabling us to uncover the more fundamental forces that underpin information aggregation in social networks. 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.

For a learning rule that satisfies IR and LN, we define the following concept:

Weak Separability (WS). There exists a smooth function $\psi_i : [0, 1]^n \rightarrow \mathbb{R}_+$ such that

$$\frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} = \frac{\psi_i(\mu(\theta))}{\psi_i(\mu(\hat{\theta}))} \quad (7)$$

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

¹⁴See Appendix A, where we also show that any learning rule according to which agents update their beliefs by linearly combining their neighbors' point estimates using time-invariant weights violates the joint restriction imposed by imperfect recall and IIA, even if signals and beliefs are normally distributed.

In other words, in determining the relative likelihoods of two given states, agent i solely relies on her neighbors' relative beliefs, transformed by some mapping ψ_i , about those states. Thus, any function ψ_i results in a distinct social learning rule of the form

$$f_i(\mu)(\theta) = \frac{\psi_i(\mu(\theta))}{\sum_{\hat{\theta} \in \Theta} \psi_i(\mu(\hat{\theta}))} \quad (8)$$

for all $\theta \in \Theta$. The key observation is that WS imposes a weaker requirement on agents' social learning rules than both IIA and separability. In fact, our representation Theorems 1 and 3 reveal that the log-linear and DeGroot learning rules belong to the general 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.

Given its general nature, the class of weakly separable learning rules does not lend itself to a representation theorem similar to the ones we obtained in Sections 3 and 4. Nonetheless, our next set of results provide a tight characterization of how different features of agents' social learning rules determine their long-run beliefs.

Theorem 5. *Suppose agents' social learning rules satisfy IR, LN, monotonicity, and WS.*

(a) *Agents mislearn the state with positive probability if ψ_i 's are homogenous of degree $\rho > 1$.*

(b) *Agents remain uncertain forever with probability one if ψ_i 's are homogenous of degree $\rho < 1$.*

This result generalizes Theorem 2 to the class of weakly separable learning rules, highlighting the key role played by the degree of homogeneity of ψ_i 's for agents' long-run beliefs. In particular, statement (a) of the theorem illustrates that homogeneity of a degree larger than 1 results in potential mislearning. This is due to the fact that, for such learning rules, agents' initial beliefs are reinforced following their social interactions. Put differently, agents end up assigning progressively growing weights to their earlier observations at the expense of the signals they receive later on. In contrast, with $\rho < 1$, agents downplay their current beliefs in favor of their most recent observations and as a result never accumulate enough information to reach certainty. In fact, in a close parallel to the predictions of Theorem 2, it can be immediately verified that when agents follow the log-linear learning rule (4), the conditions that $\rho > 1$ and $\rho < 1$ reduce, respectively, to the notions of group polarization and depolarization defined in Subsection 3.2. Taken together, the two parts of Theorem 5 illustrate that 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 already accumulated in their neighborhoods. The theorem also clarifies that asymptotic learning relies on knife-edge conditions that are satisfied in DeGroot's model but are bound to be violated in all but few of its generalizations.

In what follows, we study the dynamics of beliefs when agents employ weakly separable learning rules for which $\rho = 1$. But first, we define the *logarithmic curvature* of agent i 's social learning rule (in the direction of j and k) as

$$\delta_i^{(kj)}(x) = - \left(\frac{\partial^2 \log \psi_i}{\partial \log x_k \partial \log x_j} \right) / \left(\frac{\partial \log \psi_i}{\partial \log x_k} \frac{\partial \log \psi_i}{\partial \log x_j} \right), \quad (9)$$

where $k \neq j$ are agents in i 's neighborhood. The key observation is that δ_i parameterizes the extent to which i 's learning rule departs from the benchmark of log-linear learning. In particular, it is easy to verify that $\delta_i^{(kj)}(x) = 0$ for all $j, k \in N_i$ and all $x \in [0, 1]^n$ whenever agent i follows the log-linear learning rule in (3). As δ_i deviates from 0 in either direction, the functional form used by agent i to combine 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 homogenous of degree 1 and takes a CES functional form given by

$$\psi_i(x) = \left[\sum_{j \in N_i} a_{ij} x_j^\xi \right]^{1/\xi}, \quad (10)$$

where a_{ij} 's are some positive constants and $\xi \in \mathbb{R}$. Note that the 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 $\delta_i^{(kj)}(x) = \xi$ for all pairs of agents $j, k \in N_i$ and all $x \in [0, 1]^n$. Therefore, any learning rule in this class for which $\xi > 0$ (such as the DeGroot model) 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 6. *Suppose agents' social learning rules satisfy IR, LN, monotonicity, and WS, where the corresponding ψ_i 's are homogenous of degree $\rho = 1$.*

- (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 are learning rules with logarithmic curvature less than -1 under which agents mislearn the state with strictly positive probability.*
- (c) *There are learning rules with logarithmic curvature larger than 1 under which agents learn the underlying state almost never.*

Part (a) of this theorem illustrates that for agents to learn the underlying state, their learning rules have to satisfy two key conditions. First, the requirement that $\rho = 1$ (which plays a role similar to unanimity in Corollary 2 and Theorem 4) 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 signals are neither discarded nor amplified as $t \rightarrow \infty$. Second, the restriction on the learning rules' logarithmic curvature 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 6 establish that the restriction on the learning rules' logarithmic curvature cannot be dispensed with, thus illustrating that homogeneity of degree one is not, in and of itself, sufficient for the long-run aggregation of information. More specifically, they illustrate that agents may either not learn or mislearn the state if the logarithmic curvature falls outside of the $[-1, 1]$ interval. Crucially, 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 would be successful.

The juxtaposition of Theorems 5 and 6 thus provides a complete picture of the forces that underpin learning, non-learning, and mislearning for the wide class of weakly separable learning rules. In particular, it clearly indicates that the long-run aggregation of information under the log-linear and DeGroot learning rules (established in Corollary 2 and Theorem 4, 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 (7)); (ii) the corresponding ψ_i 's are homogeneous of degree one; and (iii) the logarithmic curvature of agents' learning rules are confined to the interval $[-1, 1]$. Crucially, our results also establish that if either the homogeneity or curvature conditions are violated, agents may remain uncertain forever about the underlying state or assign probability one to a false state as $t \rightarrow \infty$.

6 Rate of Information Aggregation

Our results in Sections 3–5 characterize the conditions under which all agents will eventually uncover the underlying state of the world. These results, however, are silent on the precision of agents' beliefs in the short-run. 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 (4).

We emphasize that, by Theorem 2, agents do not learn the underlying state if their social learning rules are strictly polarizing or depolarizing. Therefore, throughout this section, we restrict our attention to an environment with non-polarizing learning rules to ensure that the rate of learning is a meaningful concept.

6.1 Information and Centrality

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

$$h_i(\theta, \hat{\theta}) = \mathbb{E}^\theta \left[\log \frac{\ell_i^\theta(\omega)}{\ell_i^{\hat{\theta}}(\omega)} \right]$$

denote the *relative entropy* of θ with respect to $\hat{\theta}$ in i 's signal structure.¹⁵ This metric captures the expected information (per signal) in agent i 's private observations in favor of the hypothesis that the underlying state is θ against the alternative hypothesis $\hat{\theta}$, when the underlying state is indeed θ . When $h_i(\theta, \hat{\theta})$ 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 fact, the number of observations required to reach a given pre-specified confidence is determined by the magnitude of $h_i(\theta, \hat{\theta})$: a larger $h_i(\theta, \hat{\theta})$ means that the agent can rule out $\hat{\theta}$ with fewer observations. On the other hand, if $h_i(\theta, \hat{\theta}) = 0$, agent i would not be able to distinguish between the states based on her private signals alone, no matter how many observations she makes.

Even though relative entropy captures the speed at which new information is revealed to the agents, the fact that this information has to be eventually relayed over the social network means that the collection $\{h_i(\theta, \hat{\theta})\}_{i \in N, (\theta, \hat{\theta}) \in \Theta^2}$ is not a sufficient statistic for the rate of learning. Rather, this rate also depends on the speed at which information travels from one agent to another, which in turn is determined by the structure of the social network.

To account for the differential roles of various agents in the social network, we define the *out-centrality* of agent i as

$$v_i = \sum_{j=1}^n v_j a_{ji}, \quad (11)$$

where a_{ji} is the weight in (4) used by agent j to incorporate the belief of agent i . The out-centrality of an agent, which coincides with the well-known notion of eigenvector centrality, is thus a measure of the agent's importance as a source of influence: an individual is more out-central if other more out-central agents put a large weight on her opinion. Similarly, we define the *in-centrality* of agent i as

$$w_i = \sum_{j=1}^n w_j a_{ij}. \quad (12)$$

Parallel to our earlier notion, an agent's in-centrality captures the extent of her reliance (directly or indirectly) on the information provided by other agents. Finally, note that equations (11) and (12) have strictly positive solutions as long as the underlying social network is strongly connected and the spectral radius of matrix A is equal to 1, a condition that is satisfied whenever agents' social learning rules are non-polarizing (see Lemma B.1 in the Appendix). This condition also guarantees that in- and out-centralities are uniquely determined up to a scaling. We normalize these values by setting $\sum_{i=1}^n v_i w_i = 1$.¹⁶

6.2 Learning Rate

Let $e_{it}^\theta = \sum_{\hat{\theta} \neq \theta} \mu_{it}(\hat{\theta})$ denote the belief that agent i assigns to states other than θ at time t when the underlying state is indeed θ . As already discussed, Theorem 2(c) guarantees that $e_{it}^\theta \rightarrow 0$ almost

¹⁵For more on relative entropy and related concepts in information theory, see Cover and Thomas (1991).

¹⁶This normalization assumption is made to simplify the analytical expressions, with no bearing on our results.

surely as $t \rightarrow \infty$ whenever learning rules are non-polarizing. We define agent i 's *rate of learning* as

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

This quantity 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. Note that if the above limit is finite and positive, agent i learns the underlying state exponentially fast.

Theorem 7. *Suppose agents' learning rules satisfy IR, LN, monotonicity, and IIA. Furthermore, suppose the learning rules are non-polarizing. Then, the rate of learning of agent i is*

$$\lambda_i^\theta = \min_{\hat{\theta} \neq \theta} w_i \sum_{j=1}^n v_j h_j(\theta, \hat{\theta}) \quad (13)$$

almost surely, where v and w denote the out-centrality and in-centrality, respectively.

As a first implication, the above result guarantees that all agents' rates of learning are non-zero and finite, thus implying that they learn the underlying state exponentially fast. The significance of Theorem 7, however, lies in establishing that the rate of learning depends not only on the total amount of information available throughout the network, but also on how this information is distributed among different agents, summarized via their in- and out-centralities.

Expression (13) for the rate of learning has an intuitive interpretation. Recall that relative entropy $h_j(\theta, \hat{\theta})$ is the expected rate at which agent j accumulates evidence in favor of θ against $\hat{\theta}$ when the realized state is indeed θ . Thus, it is not surprising that, *ceteris paribus*, an increase in the informativeness of agents' signals cannot lead to a slower rate of learning. In addition to the signal structures, the rate of learning also depends on the structure of the social network. In particular, the relative entropy between distributions ℓ_j^θ and $\ell_j^{\hat{\theta}}$ is weighted by agent j 's out-centrality, which measures the effective (direct and indirect) attention she receives from other agents in the social network. This characterization implies that with dispersed information, social learning exhibits a “network bottleneck effect”: the long-run dynamics of the beliefs is less sensitive to changes in the information of peripheral agents who receive little attention from others.

The characterization in (13) also highlights that agents may learn the underlying states at potentially different rates. In particular, agent i 's learning rate consists of a common term, $\sum_{j=1}^n v_j h_j(\theta, \hat{\theta})$, which is then weighted by her in-centrality w_i . This term captures the intuitive idea that agents who pay more (direct or indirect) attention to the information available to other agents learn at a faster rate. In the special case that learning rules are unanimous (that is, $\sum_{j=1}^n a_{ij} = 1$ for all i), all agents have identical in-centralities, and as a result they learn the underlying state at the same exact rate, with out-centralities serving as sufficient statistics for the structure of the social network:

Corollary 3. *If agents' social learning rules are unanimous, then $\lambda_i^\theta = \min_{\hat{\theta} \neq \theta} \sum_{j=1}^n v_j h_j(\theta, \hat{\theta})$.*

As a final remark, note that learning is complete only if agents can rule out all incorrect states. More specifically, conditional on the realization of θ , the speed of learning depends on the rate at

which agents rule out state $\hat{\theta} \neq \theta$ that is closest to θ in terms of relative entropy. Thus, as (13) suggests, the rate of learning is determined by minimizing the weighted sum of relative entropies over all other possible alternatives $\hat{\theta} \neq \theta$. This characterization points towards the presence of an “identification bottleneck effect,” according to which the rate of learning is determined by the pair of states $(\theta, \hat{\theta})$ that are hardest to distinguish from one another by the collection of all individuals in the social network.

7 Conclusions

The complexity involved in Bayesian inference over social networks has led to a growing literature on non-Bayesian models of social learning. This literature, which for the most part builds on the canonical model of DeGroot, imposes specific functional form assumptions on how agents incorporate other people’s opinions into their beliefs. Such non-Bayesian heuristics, however, are open to the challenge that even though Bayesian learning is a well-defined concept, deviations from Bayesian rationality are bound to be ad hoc.

In this article, we take an alternative approach by studying the link between different behavioral assumptions and various learning rules. In particular, we impose a set of restrictions on how individuals incorporate their neighbors’ beliefs and obtain representation theorems that identify the corresponding learning rules up to a set of constants. As a first result, we establish that as long as imperfect recall represents the only point of departure from Bayesian rationality, agents follow learning rules that are linear in their neighbors’ log-likelihood ratios. This approach also enables us to compare the behavioral assumptions that underpin the log-linear learning model with those of the canonical model of DeGroot. We then provide a taxonomy of conditions for the long-run success or failure of information aggregation that are not tied to the specific functional form of agents’ learning rules, thus identifying the forces that underpin learning, non-learning, and mislearning in social networks. In particular, our results establish that the information dispersed throughout the society is properly aggregated if (i) weights that any given agent assigns to each independent piece of information are of the same order of magnitude; and (ii) agents’ social learning rules do not exhibit too strong of a departure from the log-linear learning benchmark. Finally, we establish that the speed of information aggregation is the result of the interplay between the dispersion of information among individuals and the underlying structure of the social network.

A Log-Linear Learning under Normality

In this appendix, we show that if all signals and prior beliefs are normally distributed, then the log-linear learning rule (3) is equivalent to a learning rule in which agents update their beliefs by linearly combining their neighbors' point estimates about the underlying state. We also show that for this learning rule to be consistent with the joint restrictions imposed by imperfect recall and IIA, the corresponding weights cannot remain constant over time.

Suppose that $\Theta = \mathbb{R}$ and that all agents' belief are normally distributed at time t . In particular, suppose that $\mu_{it} \sim \mathcal{N}(m_{it}, 1/\kappa_{it})$. It is easy to see that for any given pair of states θ and $\hat{\theta}$,

$$\log \frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} = \frac{\kappa_{it}}{2} (\hat{\theta} - \theta)(\hat{\theta} + \theta - 2m_{it}). \quad (14)$$

On the other hand, recall from Theorem 1 that under IR, LN, monotonicity, and IIA, the learning rule of agent i at any given time satisfies

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

for a set of constants $a_{ij} > 0$. Combining the above with (14) thus implies that

$$\log \frac{f_{it}(\mu_i^t)(\theta)}{f_{it}(\mu_i^t)(\hat{\theta})} = \frac{1}{2\beta_{it}} (\hat{\theta} - \theta)(\hat{\theta} + \theta - 2\beta_{it}\gamma_{it}),$$

where $\beta_{it}^{-1} = \sum_{j=1}^n a_{ij}\kappa_{jt}$ and $\gamma_i = \sum_{j=1}^n a_{ij}\kappa_{jt}m_{jt}$. Therefore, the log-linear learning rule in (4) perseveres normality as long as all beliefs and private signals are normally distributed. Furthermore, the mean and precision of agent i 's beliefs after aggregating her neighbors' opinions (but before observing her private signal) are given by

$$m_{it+1} = \sum_{j=1}^n \left(\frac{a_{ij}\kappa_{jt}}{\sum_{r=1}^n a_{ir}\kappa_{rt}} \right) m_{jt} \quad (15)$$

$$\kappa_{it+1} = \sum_{j=1}^n a_{ij}\kappa_{jt}. \quad (16)$$

Consequently, under the assumption of normally distributed beliefs and signals the point estimate of agent i is a convex combination of the point estimates of all her neighbors. Crucially, however, note that imperfect recall and IIA require these weights to evolve with time. In particular, even though the weights a_{ij} are independent of time index t , the weights agents use in (15) to update the means of their beliefs depend on t . This observation thus illustrates that even in the presence of normally distributed signals and beliefs, a learning rule in which agents use fixed weights to incorporate their neighbors' point estimates violates the joint restriction imposed by imperfect recall and IIA.

B Proofs

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,

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

On the other hand, IIA implies that

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

Note that $\text{cond}_{\bar{\Theta}}(\mu)$ 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 $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$\log \frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} = g_i \left(\log \frac{\mu_1(\theta)}{\mu_1(\hat{\theta})}, \dots, \log \frac{\mu_n(\theta)}{\mu_n(\hat{\theta})} \right) \quad (17)$$

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

Now, consider three distinct states $\theta, \hat{\theta}$ and $\tilde{\theta}$. Given that (17) has to be satisfied for any arbitrary pair of states, we have

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

Since μ was arbitrary, the above equation implies that for any arbitrary $x, y \in \mathbb{R}^n$, it must be the case that

$$g_i(x) + g_i(y) = g_i(x + y).$$

The above equation is nothing but Cauchy's functional equation, with linear functions as its single family of continuous solutions, which means there exist constants a_{ij} such that $g_i(x) = \sum_{j=1}^n a_{ij}x_j$. Thus, using (17) one more time implies that

$$\log \frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} = \sum_{j=1}^n a_{ij} \log \frac{\mu_j(\theta)}{\mu_j(\hat{\theta})}$$

for all $\theta, \hat{\theta} \in \Theta$. Finally, monotonicity guarantees that $a_{ij} > 0$ for all $j \in N_i$, completing the proof. $\square$

Proof of Theorem 2

We start by stating and proving a lemma, relating the notions of weak and strict group polarization to the spectral radius of matrix $A = [a_{ij}]$.

Lemma B.1. *If agents' social learning rules satisfy group polarization, then $\rho(A) \geq 1$. Furthermore, if agent's social learning rules satisfy strict group polarization, then $\rho(A) > 1$.*

Proof. Suppose agents' social learning rules satisfy group polarization. By assumption, there exists a profile of beliefs μ and a pair of states $\theta \neq \hat{\theta}$ such that

$$\frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} \geq \frac{\mu_i(\theta)}{\mu_i(\hat{\theta})} \geq 1 \quad (18)$$

for all i . Furthermore, by Theorem 1,

$$\log \frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} = (Ay)_i,$$

where $y_i = \log(\mu_i(\theta)/\mu_i(\hat{\theta}))$. Combining the above with (18) implies that $(Ay)_i/y_i \geq 1$ for all i . On the other hand, Theorem 2.10 of [Berman and Plemmons \(1979, p. 31\)](#) guarantees that $\rho(A) \geq \min_i (Ay)_i/y_i$ for all non-negative vectors y , thus establishing that $\rho(A) \geq 1$. The proof for the case of strict group polarization is analogous. $\square$

We now proceed to the proof of Theorem 2. Let $v \in \mathbb{R}^n$ denote the left eigenvector of matrix A corresponding to its largest eigenvalue, that is, $v'A = \rho v'$, where ρ is the spectral radius of A . Since A is non-negative and irreducible, the Perron-Frobenius theorem guarantees that $v_i > 0$ for all i .

By Corollary 1, the belief update rule of agent i is given by (4) for any $\hat{\theta} \neq \theta$. Multiplying both sides of (4) by v_i and summing over all i leads to

$$\sum_{i=1}^n v_i \log \frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \sum_{i=1}^n v_i \log \frac{\ell_i^\theta(\omega_{it+1})}{\ell_i^{\hat{\theta}}(\omega_{it+1})} + \sum_{j=1}^n \sum_{i=1}^n v_i a_{ij} \log \frac{\mu_{jt}(\theta)}{\mu_{jt}(\hat{\theta})}.$$

The fact that v is the left eigenvector of A guarantees that

$$\sum_{i=1}^n v_i \log \frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \sum_{i=1}^n v_i \log \frac{\ell_i^\theta(\omega_{it+1})}{\ell_i^{\hat{\theta}}(\omega_{it+1})} + \rho \sum_{j=1}^n v_j \log \frac{\mu_{jt}(\theta)}{\mu_{jt}(\hat{\theta})},$$

and as a result,

$$x_t = \rho^t x_0 + \sum_{\tau=1}^t \rho^{t-\tau} r(\omega_\tau), \quad (19)$$

where $x_t = \sum_{i=1}^n v_i \log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta}))$ and $r(\omega) = \sum_{i=1}^n v_i \log(\ell_i^\theta(\omega_i)/\ell_i^{\hat{\theta}}(\omega_i))$.

Proof of part (a) Choose $\theta, \hat{\theta} \in \Theta$ such that $x_0 = \sum_{i=1}^n v_i \log(\mu_{i0}(\theta)/\mu_{i0}(\hat{\theta}))$ is non-positive. Given that agents' social learning rules satisfy strict group polarization, Lemma B.1 guarantees that $\rho > 1$. Therefore, equation (19) implies that

$$\begin{aligned} \rho^{-t} x_t &\leq x_0 + \sum_{\tau=1}^{T-1} \rho^{-\tau} r(\omega_\tau) + \sum_{\tau=T}^t \rho^{-\tau} r_{\max} \\ &\leq x_0 + \sum_{\tau=1}^{T-1} \rho^{-\tau} r(\omega_\tau) + \frac{\rho^{-T}}{1 - \rho^{-1}} r_{\max} \end{aligned} \quad (20)$$

for an arbitrary T , where $r_{\max} = \max_{\omega} r(\omega) > 0$. Since x_0 is non-positive, there exists a large enough T and a sequence of signal profiles $(\omega_1, \dots, \omega_T)$ such that the right-hand side of (20) is strictly negative, thus guaranteeing that $\limsup_{t \rightarrow \infty} \rho^{-t} x_t < 0$. Therefore,

$$\lim_{t \rightarrow \infty} \sum_{i=1}^n v_i \log \frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} = -\infty$$

with some strictly positive probability, regardless of the state of the world. Therefore, there exists at least one agent j who assigns an asymptotic belief of zero to θ . But since the learning rule satisfies IIA and the network is strongly connected, if one agent rules out state θ on some sample path, all other agents would eventually do so as well. Now, since there is an *ex ante* positive probability that the true state is θ , this means that there exists a positive probability that all agents mislearn the underlying state. $\square$

Proof of part(b) Suppose there exists a sample path over which agent i becomes certain that the underlying state of the world is not $\hat{\theta}$; that is, $\lim_{t \rightarrow \infty} \mu_{it}(\hat{\theta}) = 0$. This, alongside the strong connectivity of the social network, guarantees that all other agents also assign an asymptotic belief of zero to $\hat{\theta}$.

Let $\theta \neq \hat{\theta}$ denote a state over which all agents assign a positive probability infinitely often on that sample path. From (19) we have,

$$x_t \leq \rho^t x_0 + \sum_{\tau=1}^t \rho^{t-\tau} r_{\max},$$

where $r_{\max} = \max_{\omega} r(\omega)$ and ρ is the spectral radius of A . On the other hand, given that agents' social learning rules satisfy strict group depolarization, a result analogous to Lemma B.1 guarantees that $\rho < 1$. Consequently,

$$\limsup_{t \rightarrow \infty} \sum_{i=1}^n v_i \log \frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} \leq \frac{r_{\max}}{1 - \rho}. \quad (21)$$

This inequality, however, is inconsistent with the proposition that all agents assign an asymptotic belief of zero to $\hat{\theta}$, which requires that the left-hand side of (21) to diverge to $+\infty$. Consequently, all agents remain uncertain about the underlying state of the world on all sample paths. $\square$

Proof of part (c) Let θ denote the underlying state of the world. From Lemma B.1, it is immediate that $\rho = 1$ whenever agents' social learning rules are non-polarizing. Consequently, equation (19) implies that

$$\lim_{t \rightarrow \infty} \frac{1}{t} x_t = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t r(\omega_\tau).$$

Since agents' private signals are i.i.d. over time, the law of large numbers guarantees that

$$\lim_{t \rightarrow \infty} \frac{1}{t} x_t = \mathbb{E}^\theta[r(\omega)]$$

almost surely, and as a result,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^n v_i \log \frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} = \sum_{i=1}^n v_i \mathbb{E}^\theta \left[\log \frac{\ell_i^\theta(\omega)}{\ell_i^{\hat{\theta}}(\omega)} \right]$$

with probability one. Jensen's inequality and the fact that $v_i > 0$ for all i guarantee that the right-hand side of the above equation is strictly positive and finite. Consequently, there exists at least one agent j such that

$$\lim_{t \rightarrow \infty} \log \left(\frac{\mu_{jt}(\theta)}{\mu_{jt}(\hat{\theta})} \right) = \infty \quad (22)$$

with probability one. Now (4) guarantees that any agent i who has agent j in their neighborhood needs to also satisfy $\log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta})) \rightarrow \infty$. Given that the social network is strongly connected, an inductive argument then guarantees that (22) is satisfied for all agents with probability one. This in turn implies that

$$\lim_{t \rightarrow \infty} \mu_{jt}(\hat{\theta}) = 0$$

almost surely for all j and all $\hat{\theta} \neq \theta$. Hence, all agents almost surely learn the underlying state of the world. $\square$

Proof of Theorem 3

The separability property implies that for any given i , there exists a function $g_i : [0, 1]^n \rightarrow [0, 1]$ such that

$$f_i(\mu)(\theta) = g_i(\mu_1(\theta), \dots, \mu_n(\theta)) \quad (23)$$

for belief profiles $\mu \in \Delta\Theta$ and all states $\theta \in \Theta$, with LN guaranteeing that g_i is independent of θ .

Fix a pair of states $\theta \neq \hat{\theta}$ and an arbitrary belief profile $\mu \in \Delta\Theta^n$. Equation (23) guarantees that

$$g_i(\mu(\theta)) + g_i(\mu(\hat{\theta})) = f_i(\mu)(\theta) + f_i(\mu)(\hat{\theta}) = 1 - \sum_{\tilde{\theta} \notin \{\theta, \hat{\theta}\}} f_i(\mu)(\tilde{\theta}),$$

and as a result

$$g_i(\mu(\theta)) + g_i(\mu(\hat{\theta})) = 1 - \sum_{\tilde{\theta} \notin \{\theta, \hat{\theta}\}} g_i(\mu(\tilde{\theta})). \quad (24)$$

Note that changing $\mu(\theta)$ and $\mu(\hat{\theta})$ does not impact the right-hand side of (24) as long as $\mu(\tilde{\theta})$ is kept unchanged for all $\tilde{\theta} \notin \{\theta, \hat{\theta}\}$. Consequently,

$$g_i(\mu(\theta)) + g_i(\mu(\hat{\theta})) = g_i(\mu(\theta) + \mu(\hat{\theta})) + g_i(0).$$

Given that $\mu \in \Delta\Theta$ was arbitrary, the above equation implies that

$$h_i(x) + h_i(y) = h_i(x + y) \quad (25)$$

for any arbitrary $x, y \in \mathbb{R}^n$ such that $x, y \geq 0$ and $x + y \leq 1$, where $h_i(z) = g_i(z) - g_i(0)$. But (25) is nothing but Cauchy's functional equation, whose only set of continuous solutions is the family of linear functions. As a result, there exist constants a_{ij} such that

$$h_i(x) = \sum_{j=1}^n a_{ij} x_j.$$

Hence, $g_i(x) = a_{i0} + \sum_{j=1}^n a_{ij} x_j$ for some constant a_{i0} , which in turn guarantees that

$$f_i(\mu)(\theta) = a_{i0} + \sum_{j \in N_i} a_{ij} \mu_j(\theta).$$

Finally, the fact that $f_i(\mu)(\theta)$ has to be non-negative ensures that $a_{i0} \geq 0$ while monotonicity guarantees that $a_{ij} > 0$ for all $j \in N_i$. $\square$

Proof of Theorem 4

Suppose agents' social learning rules are unanimous. This implies that the constant a_{i0} in (6) is equal to zero for all i . Furthermore, it must be the case that $\sum_{j \in N_i} a_{ij} = 1$ for all i . Therefore, Proposition 3 of Jadbabaie et al. (2012) guarantees that all agents learn the underlying state almost surely.

To prove the converse, suppose that agent i 's learning rule is not unanimous. Since $f_{it}(\mu_i^t)$ is a belief vector within the simplex $\Delta\Theta$, it must be the case that $a_{i0} + \sum_{j \in N_i} a_{ij} = 1$. Therefore, the assumption that agent i 's learning rule is not unanimous implies that $a_{i0} > 0$. As a result, $\liminf_{t \rightarrow \infty} \mu_{it}(\theta) > 0$ for all $\theta \in \Theta$, regardless of the underlying state of the world, implying that agent i never reaches certainty. $\square$

Proof of Theorem 5

Proof of part (a) In order to simplify the exposition we assume that agents' social learning rules and signal structures are identical: $\psi_i = \psi_j = \psi$ and $\ell_i^\theta = \ell_j^\theta = \ell^\theta$ for all i and j .¹⁷ The weak separability assumption implies that agent i 's likelihood ratios satisfy

$$\frac{\mu_{i,t+1}(\theta)}{\mu_{i,t+1}(\hat{\theta})} = \frac{\ell_i^\theta(\omega_{i,t+1}) \psi_i(\mu_t(\theta))}{\ell_i^{\hat{\theta}}(\omega_{i,t+1}) \psi_i(\mu_t(\hat{\theta}))} = \frac{\ell^\theta(\omega_{i,t+1}) \psi(\mu_t(\theta))}{\ell^{\hat{\theta}}(\omega_{i,t+1}) \psi(\mu_t(\hat{\theta}))}, \quad (26)$$

¹⁷A proof for the general case can be made available upon request.

for all $\theta, \hat{\theta} \in \Theta$. Choose a pair of states $\theta, \hat{\theta} \in \Theta$ such that $x_0 = \log(\psi(\mu_0(\theta))/\psi(\mu_0(\hat{\theta})))$ is non-positive and consider a path of observations along which agents observe identical signals in the first T periods. Equation (26) implies that, along this path,

$$\frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\hat{\theta})} = \frac{\ell^\theta(\omega_t) \psi(\mu_{t-1}(\theta))}{\ell^{\hat{\theta}}(\omega_t) \psi(\mu_{t-1}(\hat{\theta}))}$$

for all $t \leq T$, where ω_t denotes the agents' common signal at period t . Writing the above equation in vector form, we obtain

$$\mu_t(\theta) = \frac{\ell^\theta(\omega_t) \psi(\mu_{t-1}(\theta))}{\ell^{\hat{\theta}}(\omega_t) \psi(\mu_{t-1}(\hat{\theta}))} \mu_t(\hat{\theta}),$$

where note that $\ell^\theta(\omega_t) \psi(\mu_{t-1}(\theta))$ and $\ell^{\hat{\theta}}(\omega_t) \psi(\mu_{t-1}(\hat{\theta}))$ are scalars, whereas $\mu_t(\theta)$ and $\mu_t(\hat{\theta})$ are vectors. Since ψ is homogeneous of degree $\rho > 1$,

$$\psi(\mu_t(\theta)) = \left(\frac{\ell^\theta(\omega_t) \psi(\mu_{t-1}(\theta))}{\ell^{\hat{\theta}}(\omega_t) \psi(\mu_{t-1}(\hat{\theta}))} \right)^\rho \psi(\mu_t(\hat{\theta})).$$

Consequently, $x_t = \rho r(\omega_t) + \rho x_{t-1}$ for all $t \leq T$, where $x_t = \log(\psi(\mu_t(\theta))/\psi(\mu_t(\hat{\theta})))$ and $r(\omega_i) = \log(\ell^\theta(\omega_i)/\ell^{\hat{\theta}}(\omega_i))$, and as a result,

$$x_T = \rho^T x_0 + \rho \sum_{\tau=1}^T \rho^{T-\tau} r(\omega_\tau). \quad (27)$$

Next define $r_{\max} = \max_{\omega} r(\omega)$. For $t > T$,

$$\frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\hat{\theta})} = \frac{\ell^\theta(\omega_{i,t}) \psi(\mu_{t-1}(\theta))}{\ell^{\hat{\theta}}(\omega_{i,t}) \psi(\mu_{t-1}(\hat{\theta}))} \leq e^{r_{\max}} \frac{\psi(\mu_{t-1}(\theta))}{\psi(\mu_{t-1}(\hat{\theta}))},$$

which can be written in vector form as

$$\mu_t(\theta) \leq e^{r_{\max}} \frac{\psi(\mu_{t-1}(\theta))}{\psi(\mu_{t-1}(\hat{\theta}))} \mu_{t-1}(\theta).$$

Furthermore, since ψ is increasing and homogeneous of degree ρ ,

$$\psi(\mu_t(\theta)) \leq e^{\rho r_{\max}} \left(\frac{\psi(\mu_{t-1}(\theta))}{\psi(\mu_{t-1}(\hat{\theta}))} \right)^\rho \psi(\mu_{t-1}(\theta)),$$

which in turn implies that, for all $t > T$,

$$x_t \leq \rho r_{\max} + \rho x_{t-1}. \quad (28)$$

Combining equations (27) and (28) implies that

$$\rho^{-t} x_t \leq x_0 + \rho \sum_{\tau=1}^T \rho^{-\tau} r(\omega_\tau) + \rho \sum_{\tau=T+1}^t \rho^{-\tau} r_{\max} < x_0 + \rho \sum_{\tau=1}^T \rho^{-\tau} r(\omega_\tau) + \frac{\rho^{-T}}{1 - \rho^{-1}} r_{\max}.$$

Since x_0 is non-positive, there exists a large enough T and a sequence of signals $(\omega_1, \dots, \omega_T)$ such that the right-hand side of (20) is strictly negative, thus guaranteeing that $\limsup_{t \rightarrow \infty} \rho^{-t} x_t < 0$. Therefore,

$$\lim_{t \rightarrow \infty} \log \frac{\psi(\mu_t(\theta))}{\psi(\mu_t(\hat{\theta}))} = -\infty$$

with some strictly positive probability, regardless of the underlying state of the world. Consequently, there exists at least one agent j who assigns an asymptotic belief of zero to θ , even if θ is indeed the underlying state. $\square$

Proof of part (b) Let θ denote the underlying state of the world. As in the proof of Theorem 5(a), define $x_t = \log(\psi(\mu_t(\theta))/\psi(\mu_t(\hat{\theta})))$ and $r_{\max} = \max_{\omega} r(\omega)$. An argument similar to the previous part implies that $x_t \leq \rho r_{\max} + \rho x_{t-1}$, where ρ denotes the degree of homogeneity of ψ . Consequently,

$$x_t \leq \rho^t x_0 + \sum_{t=1}^t \rho^{t-\tau} r_{\max}$$

for all t . Since $\rho < 1$, it is immediate that x_t is always bounded above, even as $t \rightarrow \infty$. In other words,

$$\limsup_{t \rightarrow \infty} \log \frac{\psi(\mu_t(\theta))}{\psi(\mu_t(\hat{\theta}))} < \infty$$

Therefore, there exists at least one agent j who does not rule out state $\hat{\theta} \neq \theta$ asymptotically. $\square$

Proof of Theorem 6(a)

We present the proof by stating and proving a sequence of lemmas. Throughout, we assume that functions $\psi_i : [0, 1]^n \rightarrow \mathbb{R}$ satisfying (7) are such that $\psi_i(1, \dots, 1) = 1$ for all i , where recall that, by assumption, ψ_i is a homogenous function of degree of 1. Also note that since the social learning rule of agent i can be rewritten as (8), this is simply a normalization and hence is without loss of generality. To simplify the proofs, we also assume that $\psi_i^{(j)}(x)$ is uniformly bounded away from zero for all x and all $j \in N_i$, where $\psi_i^{(j)}(x)$ denotes the partial derivate of $\psi_i(x)$ with respect to x_j .

Lemma B.2. *Suppose ψ_i is homogenous of degree 1 and that $\delta_i^{(kj)}(x) \leq 1$ for all $k \neq j$, where δ_i is the logarithmic curvature of i 's learning rule. Then, ψ_i is increasing and jointly concave in all arguments.*

Proof. The first statement is an immediate consequence of the fact that agent i 's social learning rule, f_i , is monotonically increasing.

To prove the second statement, it is sufficient to show that $\sum_{k=1}^n \sum_{j=1}^n y_k y_j \psi_i^{(kj)}(x) \leq 0$ for all $x \in [0, 1]^n$ and all $y \in \mathbb{R}^n$. Since ψ_i is homogenous of degree 1, $\psi_i^{(k)}$ is homogenous of degree zero for all k , which in turn implies that $\sum_{j=1}^n x_j \psi_i^{(kj)}(x) = 0$ for all $x \in [0, 1]^n$. As a result,

$$\psi_i^{(kk)}(x) = -\frac{1}{x_k} \sum_{j \neq k} x_j \psi_i^{(kj)}(x).$$

Therefore, for any given vector $y \in \mathbb{R}^n$,

$$\begin{aligned}
\sum_{k=1}^n \sum_{j=1}^n y_k y_j \psi_i^{(kj)}(x) &= \sum_{k=1}^n \sum_{j \neq k}^n y_k y_j \psi_i^{(kj)}(x) + \sum_{k=1}^n y_k^2 \psi_i^{(kk)}(x) \\
&= \sum_{k=1}^n \sum_{j \neq k}^n y_k y_j \psi_i^{(kj)}(x) - \sum_{k=1}^n \sum_{j \neq k}^n \frac{x_j}{x_k} y_k^2 \psi_i^{(kj)}(x) \\
&= -\frac{1}{2} \sum_{k=1}^n \sum_{j \neq k}^n \frac{1}{x_k x_j} (y_k x_j - y_j x_k)^2 \psi_i^{(kj)}(x).
\end{aligned}$$

As a result,

$$\sum_{k=1}^n \sum_{j=1}^n y_k y_j \psi_i^{(kj)}(x) = -\frac{1}{2} \sum_{k=1}^n \sum_{j \neq k}^n \frac{1}{x_k x_j} (y_k x_j - y_j x_k)^2 (1 - \delta_i^{(kj)}(x)) \psi_i^{(k)}(x) \psi_i^{(j)}(x) / \psi_i(x),$$

where $\delta_i^{(kj)}(x)$ is defined in (9). Now the fact that ψ_i is increasing, alongside the assumption that $\delta_i^{(kj)}(x) \leq 1$, establishes that the right-hand side of the above equation is non-positive. $\square$

Lemma B.3. *If $\sum_{k=1}^m z^k = \mathbf{1}$, then $\sum_{k=1}^m \psi_i(z^k) \leq 1$.*

Proof. The concavity of ψ_i , established in Lemma B.2, implies that for any $z^k \in [0, 1]^n$,

$$\psi_i(z^k) \leq \psi_i(\mathbf{1}) + \sum_{j=1}^n (z_j^k - 1) \psi_i^{(j)}(\mathbf{1}).$$

On the other hand, the fact that ψ_i is homogenous of degree 1 guarantees that $\sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) = \psi_i(\mathbf{1})$, and as a result, $\psi_i(z^k) \leq \sum_{j=1}^n z_j^k \psi_i^{(j)}(\mathbf{1})$. Summing both sides over k and using the assumption that $\sum_{j=1}^m z_j^k = 1$ leads to

$$\sum_{k=1}^m \psi_i(z^k) \leq \sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}).$$

Using the observation that $\sum_{j=1}^n \psi_i^{(j)}(\mathbf{1}) = \psi_i(\mathbf{1}) = 1$ completes the proof. $\square$

Lemma B.4. *The function $\phi_i : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ defined as*

$$\phi_i(x) = \left[\psi_i \left(\frac{1}{x_1}, \dots, \frac{1}{x_n} \right) \right]^{-1} \quad (29)$$

is increasing and jointly concave in all arguments.

Proof. The first statement is an immediate consequence of the fact that ψ_i is monotonically increasing. To prove the second statement, note that $\phi_i(x)$ is itself a homogeneous function of degree 1. Furthermore, the logarithmic curvature of ϕ_i is equal to the negative of the logarithmic curvature of ψ_i , that is,

$$\left(\frac{\partial^2 \log \phi_i}{\partial \log x_k \partial \log x_j} \right) \bigg/ \left(\frac{\partial \log \phi_i}{\partial \log x_k} \frac{\partial \log \phi_i}{\partial \log x_j} \right) = - \left(\frac{\partial^2 \log \psi_i}{\partial \log x_k \partial \log x_j} \right) \bigg/ \left(\frac{\partial \log \psi_i}{\partial \log x_k} \frac{\partial \log \psi_i}{\partial \log x_j} \right).$$

Therefore, the fact that $\delta_i^{(kj)}(x) \geq -1$ guarantees that the logarithmic curvature of ϕ_i is less than or equal to 1. The second part of the lemma therefore immediately follows from applying Lemma B.2 to function ϕ_i . $\square$

Lemma B.5. *The mapping*

$$\Phi(x) = \lim_{t \rightarrow \infty} \underbrace{\phi \circ \phi \circ \dots \circ \phi(x)}_{t \text{ times}}, \quad (30)$$

is well-defined, where $\phi : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ is the mapping obtained from concatenating functions $\phi_i(x)$ defined in (29).

Proof. Since ψ_i is homogenous of degree 1, it must be the case that

$$\min_{j \in N_i} \{x_j\} \leq \psi_i(x) \leq \max_{j \in N_i} \{x_j\} \quad (31)$$

for all $x \in [0, 1]^{N_i}$. This is a consequence of the assumption that $\psi_i(\mu(\theta))$ is strictly increasing in $\mu_j(\theta)$ for all $j \in N_i$.

Next, define the sequence of functions $\Psi^{m+1}(x) = \psi(\Psi^m(x))$ where $\psi : [0, 1]^n \rightarrow [0, 1]^n$ is the mapping obtained by concatenating functions ψ_i for all i and the convention that $\Psi^0(x) = x$. From (31) it is immediate that

$$\max_i \Psi_i^{m+1}(x) \leq \max_i \Psi_i^m(x).$$

Therefore, $\max_i \Psi_i^m(x)$ converges to some limit $\bar{x}$ from above as $m \rightarrow \infty$, where we are using the fact that ψ maps a compact set to itself. A similar argument implies that $\min_i \Psi_i^m(x)$ converges from below to some limit $\underline{x}$ as $m \rightarrow \infty$.

We next show that $\underline{x} = \bar{x}$. Let $V : [0, 1]^n \rightarrow [0, 1]$ be $V(x) = \max_i x_i - \min_i x_i$, which is clearly a continuous function. On the other hand, by (31),

$$V(\psi(x)) - V(x) \leq 0,$$

which means that V is a Lyapunov function for ψ .¹⁸ Hence, by LaSalle's Invariance Principle (LaSalle, 2012, p. 9), $\Psi^m(x)$ converges to some set X , where X is the largest set with the property that $\psi(x) \in X$ and $V(\psi(x)) - V(x) = 0$ whenever $x \in X$. The assumptions that the network is strongly connected and ψ_i is a strictly increasing function imply that $X = \{x : x_i = x_j, \forall i, j\}$. As a result, $\max_i \Psi_i^m(x) - \min_i \Psi_i^m(x)$ must converge to zero as $m \rightarrow \infty$, which guarantees that $\lim_{m \rightarrow \infty} \Psi_i^m(x)$ exists for all x . Consequently,

$$\Phi(x) = \lim_{t \rightarrow \infty} \underbrace{\phi \circ \phi \circ \dots \circ \phi(x)}_{t \text{ times}} = \left[\lim_{m \rightarrow \infty} \Psi^m \left(\frac{1}{x_1}, \dots, \frac{1}{x_n} \right) \right]^{-1}$$

also exists. $\square$

Lemma B.6. *Let Φ be defined as in (30).*

(a) Φ is non-decreasing, continuous, and jointly concave in all arguments over its domain.

(b) Φ is homogeneous of degree 1.

¹⁸For a definition, see, LaSalle (2012, p. 8).

Proof of part (a) Recall from Lemma B.4 that ϕ_i is non-decreasing in all arguments. The fact that Φ is obtained by composing function ϕ with itself in turn guarantees that Φ is also non-decreasing.

To show that Φ is concave, note that by Lemma B.4, ϕ_i is jointly concave for all i . This, coupled with the fact that ϕ_i are also non-decreasing guarantees that Φ is also concave. Concavity of Φ then guarantees that Φ is also continuous. $\square$

Proof of part (b) Define the sequence of functions $\phi^{(m)} : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ recursively as $\phi^{(m)}(x) = \phi(\phi^{(m-1)}(x))$, with $\phi^{(0)}(x) = x$. Given that ϕ_i is homogenous of degree 1 for all i , a simple inductive argument implies that $\phi^{(m)}(\alpha x) = \alpha \phi^{(m)}(x)$ for all $m \geq 1$. Taking limits from both sides of this equation as $m \rightarrow \infty$ proves the result. $\square$

For the next lemma, define the vector $\nu_t(\theta) \in \mathbb{R}_+^n$ as follows:

$$\nu_{it}(\theta) = \frac{1}{\mu_{it}(\theta)}. \quad (32)$$

Lemma B.7. $\Phi(\nu_t(\theta)) \rightarrow \Phi^*$ as $t \rightarrow \infty$ with $\mathbb{P}^\theta$ -probability one, where Φ^* is finite almost surely.

Proof. Recall from (1) that the belief of agent i satisfies $\mu_{it+1} = \text{BU}(f_{it}(\mu_i^t); \omega_{it+1})$, which means that

$$\mu_{i,t+1}(\theta) = \frac{\ell_i^\theta(\omega_{it+1}) f_i(\mu_t)(\theta)}{\sum_{\hat{\theta}} \ell_i^{\hat{\theta}}(\omega_{it+1}) f_i(\mu_t)(\hat{\theta})}. \quad (33)$$

Therefore,

$$\mathbb{E}_t^\theta[\nu_{i,t+1}(\theta)] = \frac{1}{f_i(\mu_t)(\theta)} \sum_{\omega_{it+1}} \sum_{\hat{\theta}} \ell_i^{\hat{\theta}}(\omega_{it+1}) f_i(\mu_t)(\hat{\theta}) = \frac{1}{f_i(\mu_t)(\theta)}.$$

Consequently, equation (8) for any weak separable learning rule f_i implies that

$$\mathbb{E}_t^\theta[\nu_{i,t+1}(\theta)] = \frac{\sum_{\hat{\theta} \in \Theta} \psi_i(\mu_t(\hat{\theta}))}{\psi_i(\mu_t(\theta))}.$$

On the other hand, since $\sum_{\hat{\theta} \in \Theta} \mu_t(\hat{\theta}) = 1$, Lemma B.3 implies that $\sum_{\hat{\theta} \in \Theta} \psi_i(\mu_t(\hat{\theta})) \leq 1$. Therefore,

$$\mathbb{E}_t^\theta[\nu_{i,t+1}(\theta)] \leq \frac{1}{\psi_i(\mu_t(\theta))} = \phi_i(\nu_t(\theta)),$$

where the equality follows from the definition of ϕ_i . 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 non-decreasing guarantees that

$$\Phi(\mathbb{E}_t^\theta[\nu_{t+1}(\theta)]) \leq \Phi(\phi(\nu_t(\theta))) = \Phi(\nu_t(\theta)),$$

where the equality is a consequence of the definition of Φ . Finally, concavity of Φ established in part (a) of Lemma B.6 alongside Jensen's inequality guarantees that

$$\mathbb{E}_t^\theta \Phi(\nu_{t+1}(\theta)) \leq \Phi(\nu_t(\theta)),$$

thus establishing that $\Phi(\nu_t(\theta))$ is an n -dimensional supermartingale. Since Φ is lower bounded, it thus converges to some vector Φ^* almost surely. $\square$

Lemma B.8. *If the underlying state of the world is θ , then $\mu_{it}(\theta)$ converges to zero almost never.*

Proof. Since $\psi_i^{(j)}$ is positive and uniformly bounded away from zero for all $j \in N_i$, if $\mu_{it}(\theta)$ converges to zero on some path, then $\mu_{jt}(\theta)$ also has to converge to zero on that path for all $j \in N_i$. Given that the social network is strongly connected, an inductive argument guarantees that $\mu_{it}(\theta) \rightarrow 0$ for all agents i . As a consequence, it is immediate that $\Phi(\nu_t(\theta)) \rightarrow \infty$ as $t \rightarrow \infty$ on any such path. However, Lemma B.7 guarantees that such paths have measure zero under the true probability distribution, as the limit Φ^* is finite almost surely. $\square$

Lemma B.9. *Suppose θ denotes the underlying state of the world. For almost all paths, there exists a vector $u \in \mathbb{R}_{++}^n$ such that $\lim_{t \rightarrow \infty} \mu_{it}(\theta) / \|\mu_t(\theta)\| = u_i$ for all i .*

Proof. Recall from (33) that

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

Replacing for f_i from (8) implies that

$$\mu_{i,t+1}(\theta) = \frac{\ell_i^\theta(\omega_{it+1}) \psi_i(\mu_t(\theta))}{\sum_{\hat{\theta}} \ell_i^{\hat{\theta}}(\omega_{it+1}) \psi_i(\mu_t(\hat{\theta}))}.$$

Furthermore, the fact that ψ_i is homogenous of degree 1 implies that $\psi_i(x) = \sum_{j=1}^n x_j \psi_i^{(j)}(x)$, and as a result, the vector of beliefs assigned to state θ by all agents satisfies

$$\mu_{t+1}(\theta) = C(\omega_{t+1}, \mu_t(\theta)) \mu_t(\theta), \quad (34)$$

where matrix $C(\omega, x) \in \mathbb{R}^{n \times n}$ is given by

$$[C(\omega, x)]_{ij} = \frac{\ell_i^\theta(\omega_i) \psi_i^{(j)}(x)}{\sum_{\hat{\theta} \in \Theta} \ell_i^{\hat{\theta}}(\omega_i) \psi_i(x)}.$$

Let M_t and m_t denote the largest and smallest non-zero elements of $C(\omega_{t+1}, \mu_t(\theta))$. Since Lemma B.8 guarantees that the elements of $\mu_t(\theta)$ converge to zero almost never, it is immediate that the ratio M_t/m_t converges to zero almost never. Therefore, by Corollary 5.1 of Hartfiel (2002), for almost all paths, the sequence of matrices $\{P_0, P_1, \dots\}$ is ergodic, where

$$P_t(\omega^t) = \prod_{\tau=0}^t C(\omega_{\tau+1}, \mu_\tau(\theta)).$$

Consequently, Theorem 5.1 of Hartfiel (2002) guarantees that

$$\lim_{t \rightarrow \infty} \tau_B(P_t(\omega^t)) = 0$$

with $\mathbb{P}^\theta$ -probability one, where τ_B is the Birkhoff contraction coefficient defined as

$$\tau_B(P) = \sup_{x, y \in \mathbb{R}_+^n} \frac{d(Px, Py)}{d(x, y)}$$

with $d(x, y) = \log \frac{\max_i (x_i/y_i)}{\min_i (x_i/y_i)}$ denoting the Hilbert projective metric. Hence, Lemma 5.1 of Hartfiel (2002) implies that the columns of $P_t(\omega^t)$ tend to column proportionality. The juxtaposition of this observation with (34) then implies the result immediately. $\square$

Proof of Theorem 6(a) We are now ready to prove the theorem. Let θ denote the underlying state of the world and recall from Lemma B.7 that $\Phi(\nu_t(\theta)) \rightarrow \Phi^*$ as $t \rightarrow \infty$ on almost all paths. On the other hand, Lemma B.9 guarantees that for any given path, there exists a positive vector u such that

$$\mu_{it}(\theta) - u_i \|\mu_t(\theta)\| \rightarrow 0$$

for all i . As a result, the continuity of Φ guarantees that with $\mathbb{P}^\theta$ -probability one

$$\Phi \left(\frac{1}{\|\mu_t(\theta)\|} v \right) \rightarrow \Phi^*,$$

where $v_i = 1/u_i$ for all i . Thus, by part (b) of Lemma B.6,

$$\frac{1}{\|\mu_t(\theta)\|} \Phi(v) \rightarrow \Phi^* \quad \mathbb{P}^\theta\text{-a.s.},$$

implying that $\|\mu_t(\theta)\|$ converges to some finite limit almost surely. Thus, Lemma B.9 guarantees that $\mu_{it}(\theta) \rightarrow \mu_i^*(\theta)$ as $t \rightarrow \infty$ with $\mathbb{P}^\theta$ -probability one, where $\mu_i^*(\theta) > 0$. Now, the fact that the limit of $\mu_{it}(\theta)$ exists and is strictly positive implies that $\mu_i^*(\theta) = 1$ for all i , establishing that all agents learn the underlying state of the world almost surely. $\square$

Proofs of Theorems 6(b) and 6(c)

To prove parts (b) and (c) of the theorem, it is sufficient to construct examples of learning rules with logarithmic curvatures outside the $[-1, 1]$ interval for which learning fails asymptotically.

Consider a network consisting of two agents. Suppose that there are only two states $\Theta = \{\theta, \hat{\theta}\}$ and that both agents follow the constant elasticity of substitution (CES) learning rule introduced in (10) with a common parameter ξ . More specifically,

$$\psi_i(\mu_1(\theta), \mu_2(\theta)) = \left(\alpha \mu_1(\theta)^\xi + (1 - \alpha) \mu_2(\theta)^\xi \right)^{1/\xi}$$

for $i \in \{1, 2\}$, where $\alpha \in (0, 1)$ is a constant. Recall that ψ_i is homogeneous of degree $\rho = 1$, with a logarithmic curvature equal to ξ throughout its domain. Suppose that agent 1's signals are uninformative, whereas agent 2 receives an informative signal of the following form: if the state is θ , with probability $p > 1/2$ she observes signal θ and with probability $1 - p$ she observes signal $\hat{\theta}$. On the other hand, if the state is $\hat{\theta}$, with probability p she observes $\hat{\theta}$ and with probability $1 - p$ she observes θ . Therefore, whenever the underlying state is θ , agents' beliefs evolve as

$$\begin{aligned} \frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} &= \frac{\psi(\mu_{1t}(\theta), \mu_{2t}(\theta))}{\psi(\mu_{1t}(\hat{\theta}), \mu_{2t}(\hat{\theta}))} \\ \frac{\mu_{2t+1}(\theta)}{\mu_{2t+1}(\hat{\theta})} &= \frac{\ell_2^\theta(\omega_{t+1}) \psi(\mu_{1t}(\theta), \mu_{2t}(\theta))}{\ell_2^{\hat{\theta}}(\omega_{t+1}) \psi(\mu_{1t}(\hat{\theta}), \mu_{2t}(\hat{\theta}))}, \end{aligned}$$

where ω_{t+1} is the signal observed by agent 2 at time $t + 1$. Therefore,

$$\frac{\mu_{2,t}(\theta)}{\mu_{2,t}(\hat{\theta})} = \frac{\ell_2^\theta(\omega_t) \mu_{1t}(\theta)}{\ell_2^{\hat{\theta}}(\omega_t) \mu_{1t}(\hat{\theta})}$$

As a result,

$$\frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} = \frac{\psi\left(\mu_{1t}(\theta), \frac{\ell_2^\theta(\omega_t)}{\mu_{1t}(\theta)\ell_2^\theta(\omega_t) + \mu_{1t}(\hat{\theta})\ell_2^{\hat{\theta}}(\omega_t)} \mu_{1t}(\theta)\right)}{\psi\left(\mu_{1t}(\hat{\theta}), \frac{\ell_2^{\hat{\theta}}(\omega_t)}{\mu_{1t}(\theta)\ell_2^\theta(\omega_t) + \mu_{1t}(\hat{\theta})\ell_2^{\hat{\theta}}(\omega_t)} \mu_{1t}(\hat{\theta})\right)}.$$

Since ψ is homogenous of degree 1, it is immediate to see that

$$\frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} = \frac{\psi(1 + \mu_{1t}(\theta)(\lambda(\omega_t) - 1), \lambda(\omega_t))}{\psi(1 + \mu_{1t}(\theta)(\lambda(\omega_t) - 1), 1)} \frac{\mu_{1t}(\theta)}{\mu_{1t}(\hat{\theta})},$$

where $\lambda(\omega_t) = \ell_2^\theta(\omega_t)/\ell_2^{\hat{\theta}}(\omega_t)$. Therefore, the log-likelihood ratio of agent 1's beliefs evolves according to

$$\log \frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} = D(\mu_{1t}(\theta), \lambda(\omega_t)) + \log \frac{\mu_{1t}(\theta)}{\mu_{1t}(\hat{\theta})}, \quad (35)$$

where

$$D(x, \lambda) = \log \frac{\psi(1 + x(\lambda - 1), \lambda)}{\psi(1 + x(\lambda - 1), 1)}. \quad (36)$$

We now state and prove a few lemmas.

Lemma B.10. *Let $D(x, \lambda)$ be defined as in (36).*

(a) *If $\xi < -1$ and $\alpha > -1/\xi$, then $\mathbb{E}^\theta D(x, \lambda(\omega_t)) < 0$ in a neighborhood of $x = 0$ and $p = 1/2$.*

(b) *If $\xi > 1$ and $\alpha > 1/\xi$, then $\mathbb{E}^\theta D(x, \lambda(\omega_t)) < 0$ in a neighborhood of $x = 1$ and $p = 1/2$.*

Proof. We only state the proof of part (a), as the proof of part (b) is identical. Suppose that $\xi < -1$ and $\alpha > -1/\xi$. It is easy to verify that when $p = 1/2$,

$$\mathbb{E}^\theta D(0, \lambda(\omega_t)) = \frac{d}{dp} \mathbb{E}^\theta D(0, \lambda(\omega_t)) = 0,$$

and

$$\frac{d^2}{dp^2} \mathbb{E}^\theta D(0, \lambda(\omega_t)) = 16(1 - \alpha)(\alpha\xi + 1),$$

which is negative since $\alpha > -1/\xi$. The lemma then follows from the continuity of $\mathbb{E}^\theta D(x, \lambda(\omega_t))$. $\square$

Lemma B.11. *Let $D(x, \lambda)$ be defined as in (36).*

(a) *If $\xi < 0$, then $D(x, \lambda)$ is increasing in x .*

(b) *If $\xi > 0$, then $D(x, \lambda)$ is decreasing in x .*

Proof. Once again, we only state the proof of the first statement, as the proof of the second statement is identical. Suppose that $\xi < 0$. It is immediate to verify that

$$\frac{d}{dx} D(x, \lambda) = (\lambda - 1) \left(\frac{\alpha(1 + x(\lambda - 1))^\xi}{\alpha(1 + x(\lambda - 1))^\xi + (1 - \alpha)\lambda^\xi} - \frac{\alpha(1 + x(\lambda - 1))^\xi}{\alpha(1 + x(\lambda - 1))^\xi + (1 - \alpha)} \right)$$

Note that when $\lambda > 1$, then $\lambda^\xi < 1$, which implies that the term in parentheses is positive. On the other hand, when $\lambda < 1$, then $\lambda^\xi > 1$, which implies that the term in parentheses is negative. Thus, either way, the expression on the right-hand side above is positive. $\square$

Proof of Theorem 6(b) Let θ denote the underlying state of the world and suppose that $\xi < -1$. By Lemma B.10, there exists a triple $(\bar{\alpha}, \underline{x}, \underline{p})$ such that if $\alpha \in (\bar{\alpha}, 1)$ and $p \in (1/2, \underline{p})$, then $\mathbb{E}^\theta D(x, \lambda(\omega_t)) < 0$ for all $x < \underline{x}$. On the other hand, equation (35) implies that

$$\log \frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} = \sum_{\tau=1}^t D(\mu_{1\tau}(\theta), \omega_\tau) + \log \frac{\mu_{11}(\theta)}{\mu_{11}(\hat{\theta})}.$$

Let $x^* < \underline{x}$. Note that for any belief $\mu_{1t}(\theta) \leq x^*$, Lemma B.11 implies that $D(\mu_{1t}(\theta), \omega_t) \leq D(x^*, \omega_t)$. Therefore, as long as the belief on θ remains below x^* , agent 1's log-likelihood ratio is upper bounded by a random walk with jump size given by $D(x^*, \omega_t)$. Since the expected drift of this random walk is negative, it converges to $-\infty$ without ever hitting x^* with some positive probability. Therefore, the log-likelihood ratio $\log(\mu_{1t+1}(\theta)/\mu_{1t+1}(\hat{\theta}))$ also converges to $-\infty$ with positive probability, which means that agents mislearn the underlying state with positive probability. $\square$

Proof of Theorem 6(c) Let θ denote the underlying state of the world and suppose that $\xi > 1$. By Lemma B.10, there exists a triple $(\bar{\alpha}, \bar{x}, \underline{p})$ such that if $\alpha \in (\bar{\alpha}, 1)$ and $p \in (1/2, \underline{p})$, then $\mathbb{E}^\theta D(x, \lambda(\omega_t)) < 0$ for all $x > \bar{x}$. On the other hand, recall from equation (35) that

$$\log \frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} = \sum_{\tau=1}^t D(\mu_{1\tau}(\theta), \omega_\tau) + \log \frac{\mu_{11}(\theta)}{\mu_{11}(\hat{\theta})}.$$

Let $x^* > \bar{x}$. For any ω let $\underline{T}_t(\omega) = \{\tau \leq t : \mu_{1\tau}(\theta) < x^*\}$ and $\bar{T}_t(\omega) = \{\tau \leq t : \mu_{1\tau}(\theta) \geq x^*\}$. Therefore,

$$\log \frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} \leq \sum_{\tau \in \underline{T}_t(\omega)} D(\mu_{1\tau}(\theta), \omega_\tau) + \sum_{\tau \in \bar{T}_t(\omega)} D(x^*, \omega_\tau) + \log \frac{\mu_{11}(\theta)}{\mu_{11}(\hat{\theta})},$$

where we are using that fact that D is decreasing in its first argument, as established in Lemma B.11. Therefore,

$$\log \frac{\mu_{1t+1}(\theta)}{\mu_{1t+1}(\hat{\theta})} \leq \log \frac{\mu_{11}(\theta)}{\mu_{11}(\hat{\theta})} + \sum_{\tau \in \underline{T}_t(\omega)} [D(\mu_{1\tau}(\theta), \omega_\tau) - D(x^*, \omega_\tau)] + \sum_{\tau=1}^t D(x^*, \omega_\tau) \quad (37)$$

By the strong law of large numbers, there exists a set $\tilde{\Omega}$ of full $\mathbb{P}^\theta$ -measure such that

$$\frac{1}{t} \sum_{\tau=1}^t D(x^*, \omega_\tau) \rightarrow \mathbb{E}^\theta D(x^*, \omega_t) < 0,$$

for all sample paths $\omega \in \tilde{\Omega}$ as $t \rightarrow \infty$, which implies that $\sum_{\tau=1}^t D(x^*, \omega_\tau) \rightarrow -\infty$.

We now show that there exists no sample path $\omega \in \tilde{\Omega}$ over which $\mu_{1t}(\theta)$ converges to one. Suppose to the contrary that for some $\omega \in \tilde{\Omega}$, it is the case that $\mu_{1t}(\theta) \rightarrow 1$. This means that over such a sample path, the left-hand side of (37) converges to $+\infty$ as $t \rightarrow \infty$. It also implies that the set $\underline{T}_t(\omega)$ remains finite for all values of t , which guarantees that the first two terms on the right-hand side of (37) remain finite as $t \rightarrow \infty$. Given that the last term goes to $-\infty$, it is then immediate that the right-hand side of (37) goes to $-\infty$, leading to a contradiction. Therefore, $\mu_{1t}(\theta)$ converges to 1 with $\mathbb{P}^\theta$ -probability zero. $\square$

Proof of Theorem 7

Let θ denote the realized state of the world and fix a state $\hat{\theta} \neq \theta$. Recall from Corollary 1 that

$$x_t = Ax_{t-1} + y(\omega_t),$$

where $x_{it} = \log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta}))$ and $y_i(\omega_{it}) = \log(\ell_i^\theta(\omega_{it})/\ell_i^{\hat{\theta}}(\omega_{it}))$. As a result,

$$x_t = A^t x_0 + \sum_{\tau=1}^t A^{t-\tau} y(\omega_\tau),$$

which in turn implies that

$$\mathbb{E}^\theta x_t = A^t x_0 + \sum_{\tau=1}^t A^{t-\tau} h(\theta, \hat{\theta}). \quad (38)$$

Let $\bar{x}_t = x_t - \mathbb{E}^\theta x_t$ and $\bar{y}(\omega_t) = y(\omega_t) - h(\theta, \hat{\theta})$ denote the corresponding demeaned random variables. It is then immediate that

$$\bar{x}_t = \sum_{\tau=1}^t z_{t,\tau},$$

where $z_{t,\tau} = A^{t-\tau} \bar{y}(\omega_\tau)$. The facts that spectral radius of A is equal to 1 and agents' signal space is finite implies that these random vectors are bounded, with second and fourth moments that are bounded uniformly in t and τ by some constants M_2 and M_4 , respectively. The fact that $\bar{x}_{it}$ is the sum of t independent random variables then implies that $\mathbb{E} \bar{x}_{it}^4 \leq tM_4 + 3(t^2 - t)M_2^2 \leq Ct^2$, for some finite constant C . Therefore, by Chebyshev's lemma,

$$\mathbb{P}^\theta \left(\frac{1}{t} |\bar{x}_{it}| > \epsilon \quad \text{for all } i \right) \leq \frac{\mathbb{E}^\theta \bar{x}_{it}^4}{(t\epsilon)^4} \leq \frac{C}{t^2 \epsilon^4}.$$

Summing over all t and using the Borel-Cantelli lemma guarantees that

$$\frac{1}{t} \bar{x}_{it} \rightarrow 0 \quad (39)$$

as $t \rightarrow \infty$ almost surely, for all i . On the other hand, (38) implies that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}^\theta \bar{x}_t = \lim_{t \rightarrow \infty} \left(\frac{1}{t} \sum_{\tau=0}^{t-1} A^\tau \right) h(\theta, \hat{\theta}) = wv' h(\theta, \hat{\theta})$$

where v and w are the left and right eigenvectors of matrix A , respectively, normalized such that $v'w = 1$. Combining the above with (39) thus guarantees that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mu_{it}(\hat{\theta}) = -w_i \sum_{j=1}^n v_j h_j(\theta, \hat{\theta}) \quad \mathbb{P}^\theta\text{-a.s.}$$

In other words, the belief assigned to any state $\hat{\theta} \neq \theta$ decays exponentially fast with the above rate. Noting that the rate at which e_t converges to zero is determined by state $\hat{\theta}$ whose rate of decay is the slowest completes the proof. $\square$

References

- Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar (2014), “Dynamics of information exchange in endogenous social networks.” *Theoretical Economics*, 9, 41–97.
- Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), “Bayesian learning in social networks.” *The Review of Economic Studies*, 78, 1201–1236.
- Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi (2010), “Spread of (mis) information in social networks.” *Games and Economic Behavior*, 70, 194–227.
- Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken (2016), “Network structure and the aggregation of information: Theory and evidence from Indonesia.” *American Economic Review*, 106, 1663–1704.
- Ali, S. Nageeb (2016), “Social learning with endogenous information.” Working Paper.
- Bala, Venkatesh and Sanjeev Goyal (1998), “Learning from neighbours.” *The Review of Economic Studies*, 65, 595–621.
- Banerjee, Abhijit (1992), “A simple model of herd behavior.” *The Quarterly Journal of Economics*, 107, 797–817.
- Banerjee, Abhijit, Emily Breza, Arun G. Chandrasekhar, and Markus Möbius (2016), “Naive learning with uninformed agents.” Working paper.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson (2013), “The diffusion of microfinance.” *Science*, 341.
- Berman, Abraham and Robert J. Plemmons (1979), *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York.
- 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.
- Breza, Emily (2016), “Field experiments, social networks, and development.” In *The Oxford Handbook on the Economics of Networks* (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.), chapter 16, 412–439, Oxford University Press.
- 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.
- Choi, Syngjoo, Edoardo Gallo, and Shachar Kariv (2016), “Networks in the laboratory.” In *The Oxford Handbook on the Economics of Networks* (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.), chapter 17, Oxford University Press.

- Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer (2015), "Asset pricing when 'This Time is Different'." Columbia Business School Research Paper No. 14-8.
- Cover, Thomas M. and Joy A. Thomas (1991), *Elements of Information Theory*. John Wiley & Sons, New York, NY.
- 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.
- Eyster, Erik, Matthew Rabin, and Georg Weizsäcker (2015), "An experiment on social mislearning." Working paper.
- Gale, Douglas and Shachar Kariv (2003), "Bayesian learning in social networks." *Games and Economic Behavior*, 45, 329–346.
- 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.
- Grimm, Veronika and Friederike Mengel (2015), "Experiments on belief formation in networks." Working Paper.
- Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz (2015), "The speed of social learning." Working Paper.
- Hartfiel, Darald J. (2002), *Nonhomogenous Matrix Products*. World Scientific Publishing, Singapore.
- Isenberg, Daniel J. (1986), "Group polarization: A critical review and meta-analysis." *Journal of Personality and Social Psychology*, 50, 1141–1151.
- Jackson, Mathew O. (2008), *Social and Economic Networks*. Princeton University Press, Princeton, NJ.

- Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012), “Non-Bayesian social learning.” *Games and Economic Behavior*, 76, 210–225.
- LaSalle, Joseph P (2012), *The Stability and Control of Discrete Processes*, volume 62. Springer Science & Business Media.
- 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 (2015), “Perceptions of correlation, ambiguity and communication.” Working paper.
- Levy, Gilat and Ronny Razin (2016), “Information diffusion in networks with the BPI heuristic: A note.” Working paper.
- Li, Wei and Xu Tan (2016), “Learning in local networks.” Working paper.
- Lobel, Ilan and Evan Sadler (2015), “Information diffusion in networks through social learning.” *Theoretical Economics*, 10, 807–851.
- Lobel, Ilan and Evan Sadler (2016), “Preferences, homophily, and social learning.” *Operations Research*, 64, 564–584.
- Möbius, Markus, Tuan Phan, and Adam Szeidl (2015), “Treasure hunt: Social learning in the field.” NBER Working Paper No. 21014.
- Mossel, Elchanan, Allan Sly, and Omer Tamuz (2015), “Strategic learning and the topology of social networks.” *Econometrica*, 83, 1755–1794.
- Mueller-Frank, Manuel (2013), “A general framework for rational learning in social networks.” *Theoretical Economics*, 8, 1–40.
- Mueller-Frank, Manuel and Claudia Neri (2015), “A general model of boundedly rational observational learning: Theory and evidence.” Working paper.
- Mueller-Frank, Manuel and Mallesh M. Pai (2016), “Social learning with costly search.” *American Economic Journal: Microeconomics*, 8, 83–109.
- Rabin, Matthew and Joel L. Schrag (1999), “First impressions matter: A model of confirmatory bias.” *The Quarterly Journal of Economics*, 114, 37–82.
- Rahimian, Mohammad Amin, Pooya Molavi, and Ali Jadbabaie (2014), “(Non-)Bayesian learning without recall.” Working paper.
- Smith, Lones and Peter Norman Sørensen (2000), “Pathological outcomes of observational learning.” *Econometrica*, 68, 371–398.
- Sunstein, Cass R. (2000), “Deliberative trouble? Why groups go to extremes.” *The Yale Law Journal*, 110, 71–119.