Foundations of non-Bayesian Social Learning*

Pooya Molavi† Alireza Tahbaz-Salehi‡ Ali Jadbabaie§

October 2015

Abstract

In this paper, we study the problem of non-Bayesian learning over social networks by taking an axiomatic approach. 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. 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 approach also enables us to provide a taxonomy of behavioral assumptions that underpin various non-Bayesian models of learning, including the canonical model of DeGroot. We then show that for a fairly large class of learning rules, the form of bounded rationality represented by imperfect recall is not an impediment to asymptotic learning, as long as agents assign weights of equal orders of magnitude to every independent piece of information. Finally, we show how the dispersion of information among different individuals in the social network determines the rate of learning.

Keywords: Non-Bayesian learning, social networks, bounded rationality.
JEL Classification: D83, D85, Z13.

*We are grateful to Daron Acemoglu, Kostas Bimpikis, Emily Breza, Arun Chandrasekhar, Matt Elliott, Paul Glasserman, Ben Golub, Sanjeev Goyal, Ilan Lobel, Mohamed Mostagir, Pietro Ortoleva, Andy Postlewaite, Amin Rahimian, Kamiar Rahnama-Rad, Alvaro Sandroni, Andrea Vedolin and seminar participants at Cambridge, Columbia Business School, Penn, and IESE Workshop on learning in social networks 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). This paper is partially based on the previously unpublished results in the working paper of Jadbabaie, Molavi, and Tahbaz-Salehi (2013).

†Department of Economics, Massachusetts Institute of Technology.
‡Columbia Business School, Columbia University.
§Department of Electrical and Systems Engineering, University of Pennsylvania.

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 specific functional forms on agents' learning rules, thus allowing them to capture the richness of the network interactions while maintaining computational tractability. Such heuristic non-Bayesian models in turn can be challenged on two grounds. First, although Bayesian learning is a well-defined concept, deviations from the Bayesian benchmark (such as the ones imposed by DeGroot's model) are bound to be ad hoc and arbitrary. Second, in many instances, the suggested heuristics are at best only loosely connected to the behavioral assumptions that motivated them.

In this paper, we address these challenges by taking an axiomatic approach towards social learning. In particular, rather than assuming a specific functional form for agents' social learning rules as in most of the literature, we impose a set of restrictions on how they incorporate their neighbors' information and characterize learning rules that satisfy those restrictions. This approach enables us to not only provide a systematic way of capturing deviations from Bayesian inference, but also to determine the behavioral assumptions that underpin various non-Bayesian models of social learning.

We consider an environment in which agents obtain information about an underlying state through private signals and communication with their neighbors. 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.

In addition to imperfect recall, we impose three other restrictions on agents' social learning rules, all of which 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*. Finally, 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 main 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. Crucially, given that IIA, LN, and monotonicity are satisfied by Bayesian learning, 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 axiom. 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 axiom, which we refer to as *separability*. According to this axiom, which serves as an alternative notion of independence to IIA, the posterior belief that each agent assigns to any given state is 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 time. Nevertheless, we show that both models result in asymptotic learning as long as they satisfy a condition we refer to as *unanimity*, according to which each agent adopts the shared beliefs of her neighbors whenever they all agree with one another. This condition guarantees that every independent piece of information is eventually taken into account with roughly the same weight, thus ensuring that agents' private signals are neither discarded nor amplified over time.

The juxtaposition of our representation theorems for the log-linear and DeGroot models reveal that the nature of social learning depends on the underlying notion of independence satisfied by the learning rule. We leverage this observation and obtain a fairly general class of learning rules by relaxing the IIA and separability axioms. In particular, we show that replacing these axioms with a weaker notion of independence, which we refer to as *weak separability*, results in a general class of learning rules that encompasses log-linear and DeGroot models as special cases. We then establish that all weakly separable learning rules result in asymptotic learning as long as a strong form of unanimity is satisfied. This general result illustrates that the form of deviation from Bayesian rationality represented by imperfect recall, in and of itself, is not an impediment to asymptotic aggregation of information.

We end the paper with a discussion on how the interplay between the dispersion of information and the structure of the social network determines the rate at which information gets aggregated. We show that this rate has a simple analytical characterization in terms of the relative entropy of agents' signal structures and different notions of network centrality. Our characterization illustrates that the way information is dispersed throughout the social network has non-trivial implications for the rate of learning. In particular, we show that when the informativeness of different agents' signal structures are comparable in the sense of Blackwell (1953), then a positive assortative matching of signal qualities and centralities maximizes the rate of learning. This result formalizes the idea that information diffuses faster if agents who get more attention from other individuals have access to higher quality signals. On the other hand, we also show that if information structures are such that each individual possesses some information crucial for learning, then the rate of learning is higher when agents with the best signals are located at the periphery of the network. The intuition behind this result is as follows: if the information required for distinguishing between the pair of states that are hardest to tell apart is only available to agents that receive very little effective attention from others, then it would take a long time for (i) those agents to collect enough information to distinguish between the two states; and (ii) for this information to be diffused throughout the network. A negative assortative matching of signal qualities and centralities guarantees that these two events happen in parallel, leading to a faster convergence rate.

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 from this approach in one significant way: instead of assuming functional forms for how agents incorporate their neighbors' opinions into their beliefs, we take an axiomatic approach towards social learning and obtain representations of learning rules that satisfy those axioms. This alternative 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; and (ii) provide a taxonomy of behavioral assumptions that underpin various non-Bayesian learning models.

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) and Gagnon-Bartsch and Rabin (2015) 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

[1] Some of the more recent contributions in this literature include Acemoglu, Ozdaglar, and ParandehGheibi (2010), Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012), and Banerjee, Breza, Chandrasekhar, and Möbius (2015).

[2] Other works in this literature include 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 (2015a,b), Mossel, Sly, and Tamuz (2015), and Mueller-Frank and Pai (2015). See Golub and Sadler (2015) for a thorough survey of the social learning literature, covering both Bayesian and non-Bayesian paradigms.

neighbors' beliefs are in turn affected by their own social interactions, whereas Bala and Goyal (1998) depart from the Bayesian framework by assuming that, in updating her beliefs, an agent does not make inferences concerning the experience of unobserved agents (such as some of the neighbors of her neighbors).

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 (2015) conduct a lab experiment in rural India to test whether a variant of DeGroot model can outperform the Bayesian framework in describing learning in social networks. Relatedly, Möbius, Phan, and Szeidl (2015) test between DeGroot learning and an alternative non-Bayesian model in which individuals "tag" information by describing its origin.[3] Despite its theoretical take, the axiomatic approach taken in our paper contributes to this debate by illustrating the behavioral assumptions that underpin DeGroot and other non-Bayesian learning models. More specifically, it suggests that testing the underlying axioms of a given learning rule 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 which 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. In Section 3, we introduce the notion of imperfect recall as our main behavioral assumption and present our main representation theorem. Section 4 is dedicated to the DeGroot learning model and its axiomatic foundations. We study a more general class of social learning rules in Section 5 and characterize the rate of learning in Section 6. Section 7 contains our conclusions. All proofs and some additional mathematical details are provided in the Appendix.

[3] Also see the recent works of Banerjee, Chandrasekhar, Duflo, and Jackson (2013), Grimm and Mengel (2014) and Mueller-Frank and Neri (2015). See Breza (2015) and Choi, Gallo, and Kariv (2015) for a survey of field and lab experiments on social networks.

2 Setup

Consider a collection of n individuals, denoted by $N = \{1, 2 \ldots, n\}$, 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 over 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 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, which we refer to as her *neighbors*. More specifically, at the beginning of time period t, and before observing the realization of her private signal ω_{it}, agent i observes the beliefs held by her neighbors at the previous time period.[4] This form of social interactions can be represented by a directed graph on n vertices, which we refer to as the *social network*. Each vertex of this graph corresponds to an agent and a directed edge (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.

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

$$\mu_{it+1} = \mathrm{BU}\big(f_{it}(\mu_i^t); \omega_{it+1}\big), \tag{1}$$

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 $\mathrm{BU}(\mu; \omega)$ denotes the Bayesian update of μ conditional on the observation of signal ω. The function $f_{it} : \Delta\Theta^{|N_i|(t+1)} \to \Delta\Theta$, 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]

[4] The 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) for a discussion of observational learning models in information-rich settings.

[5] With some abuse of notation, we represent agent i's social learning rule with $f_{it} : \Delta\Theta^{n(t+1)} \to \Delta\Theta$, with the understanding that it does not depend on the beliefs of agents who are not i's neighbors.

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 potential disparity between the ways agents process their private and social information in (1) is imposed for two reasons. First, it is natural to expect that agents find it easier to rationally process their private signals compared to the information provided by other individuals. Whereas each agent's private signals are distributed according to a distribution known to her, her neighbors' beliefs may encompass multiple pieces of potentially redundant information, which she may find hard to disentangle without complete knowledge of the social network or other agents' signal structures. 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.

3 Social Learning under Imperfect Recall

In this section, we impose a set of restrictions on how agents incorporate their neighbors' information into their beliefs and introduce the notion of "imperfect recall" as our main behavioral assumption. According to this notion, 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. We then show that as long as imperfect recall represents the only point of departure from Bayesian inference, our restrictions yield a representation that uniquely identifies agents' social learning rules up to a set of constants.

3.1 Axiomatic Foundations

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

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

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

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

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, agents' social learning rules are label neutral.

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 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 $\mathrm{cond}_{\bar{\Theta}} : \Delta\Theta \to \Delta\Theta$ denote the operator that maps a belief to the corresponding belief conditional on the subset of states $\bar{\Theta} \subseteq \Theta$; that is, $\mathrm{cond}_{\bar{\Theta}}(\mu)(\theta) = \mu(\theta|\bar{\Theta})$.

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

$$\mathrm{cond}_{\bar{\Theta}}\left(f_{it}(\mu_i^t)\right) = f_{it}\left(\mathrm{cond}_{\bar{\Theta}}(\mu_i^t)\right), \tag{2}$$

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

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.[6]

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 $\theta \in \bar{\Theta}$.

[6]Note 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 states.

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 updating.[7] 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.

We now turn to our main behavioral assumption on agents' social learning rules:

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.

In other words, agent i only relies on her neighbors' most recent opinions to update her belief, while ignoring the history of their beliefs prior to the current time period. Under imperfect recall, agents are essentially only concerned with *what* others believe as opposed to *how* those opinions are formed.

The key observations is that unlike LN, IIA, and monotonicity, the restriction imposed by imperfect recall represents a fundamental departure from Bayesian rationality. To see this, note that 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 IR. 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, all 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.

Note 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. 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 (relative to the Bayesian benchmark).

We end this discussion by noting that the deviation from Bayesian rationality captured by imperfect recall is a fairly standard notion of bounded rationality that is (implicitly or explicitly) imposed

[7] The exception is the pathological case in which beliefs on irrelevant states help agents gain additional information by allowing them to disentangle different observations that lead to similar beliefs. A sufficient condition to rule out such cases is that agents achieve impartial inference as defined by Eyster and Rabin (2014).

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. (2015)) all rely on imperfect recall by assuming that agents only use the last period beliefs of their neighbors.

3.2 Representation Theorem

With our four main restrictions in place, we now provide a characterization of agents' social learning rules:

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

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

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

The significance of this characterization is twofold. First, it shows that the restrictions imposed by LN, monotonicity, IIA, and IR 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 LN, monotonicity, IIA, and IR, then*

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

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

Thus, at every period, agent i linearly combines the log-likelihood ratios of her private signal with the log-likelihood 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}.[8] This observation highlights that the assumption of imperfect recall is orthogonal to whether agents under- or over-weight their neighbors' opinions (relative to the Bayesian benchmark), as the belief update rule in (4) is consistent with IR for all a_{ij}.

As a final remark, we 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 the first term on the right-hand side of (4), while keeping the log-linear structure of the social learning rule intact.

3.3 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 study whether 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

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

for all agents i.[9] 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.

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

[9]Note 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 $f_i : \Delta\Theta^n \to \Delta\Theta$ that map a belief profile to a single belief.

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*. A sufficient (but not necessary) condition for non-polarization is for agents' learning rules to be *unanimous*, in the sense that $f_i(\mu, \ldots, \mu) = \mu$ 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.

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

(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 $t \to \infty$. 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.[10]

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.

To see the intuition underlying this result, note that as long as IIA is satisfied, a Bayesian agent assigns in any period, on net, a weight of 1 to any independent piece of information that has reached her.[11] This property, however, is clearly violated by the learning rule in (4), implying that agents'

[10] In 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.

[11] In particular, under Bayesian updating,

$$\log \frac{\mu_{it+1}(\theta)}{\mu_{it+1}(\hat{\theta})} = \log \frac{\mu_{i0}(\theta)}{\mu_{i0}(\hat{\theta})} + \sum_{d=0}^{n} \sum_{j \in N_i^d} \sum_{\tau=0}^{t-d} \log \frac{\ell_j^{\theta}(\omega_{j\tau+1})}{\ell_j^{\hat{\theta}}(\omega_{j\tau+1})}$$

for all $\theta, \hat{\theta} \in \Theta$, where $N_i^d \subseteq N$ is the set of agents who are distance d from agent i on the social network and $N_i^0 = \{i\}$. For more, see Eyster and Rabin (2014).

beliefs depart from the prescriptions of Bayesian learning at any given finite time t. Nevertheless, 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 $t \to \infty$. Therefore, despite imperfect recall, each piece of information is eventually accounted for in the long-run (even if with a weight different from 1). 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 $t \to \infty$.

Before concluding this section, it is worth emphasizing once again that the four axioms that underpin the log-linear learning rule in (4) do not impose any restrictions — besides positivity — on weights a_{ij} that agents assign to their neighbors' beliefs. This role is instead played by the assumptions imposed in different parts of Theorem 2. As we establish in Lemma A.1 in the appendix, group polarization and depolarization are essentially restrictions on the spectral radius — i.e., the largest eigenvalue in absolute value — of matrix of weights $A = [a_{ij}]$.[12] In particular, if social learning rules are non-polarizing, then the spectral radius of A has to be equal to one. This observation also highlights that agents can learn the underlying state of the world even if their learning rules are not unanimous (for example, when $\sum_{j \in N_i} a_{ij} \neq 1$ for some or all agents i). Rather, for learning to be successful, it is sufficient for the agents to neither underweight nor overweight different pieces of information *collectively*.

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 $\hat{\theta} \neq \theta$.

[12] We adopt the convention that $a_{ij} = 0$ if agent j is not in the neighborhood of agent i.

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 LN, monotonicity, IR, and separability. Then, there exists a set of constants* $b_{ij} > 0$ *and* $b_{i0} \geq 0$ *such that*

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

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.[13] Thus, in addition to providing the axiomatic foundations 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 states that DeGroot learning results in the long-run aggregation of information dispersed throughout the social network:[14]

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

Contrasting this result with Theorem 2(c) highlights that under unanimity, both DeGroot learning and the log-linear learning rule (4) 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. 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 according to which agents set their beliefs as a weighted average of 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. (2015). At the same time, some other papers, such

[13] Lehrer and Wagner (1981) provide a characterization similar to ours, albeit under a different set of restrictions.
[14] See Jadbabaie et al. (2012) for a proof.

as DeMarzo et al. (2003), use the term DeGroot learning to refer to a model according to which 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.[15] This observation also illustrates that if signals and/or beliefs are not normally distributed (for example, when Θ is discrete as in our setting), a learning rule that is based on averaging of point estimates may violate both IIA and separability. Irrespective of terminology, our characterization results clarify the distinction in the behavioral assumptions that underpin each model.

5 A General Class of Learning Rules

Theorems 1 and 3 reveal that the nature of social learning depends crucially on the learning rules' 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. For example, learning rules take the log-linear form of (3) whenever IIA is satisfied, whereas agents follow the linear learning rule of DeGroot if IIA is replaced by separability. In this section, we provide a weaker notion of independence that encompasses both IIA and separability as special cases and obtain a fairly large class of learning rules. For this class, we then establish that imperfect recall is no impediment to learning as long as a strong form of unanimity is satisfied.

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

Weak Separability (WS)**.** There exists a smooth function $\psi_i : [0,1]^n \to \mathbb{R}_+$ such that (i) the elasticity of substitution between any of its two arguments is greater than or equal to one; and (ii) for any belief profile $\mu \in \Delta\Theta^n$

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

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

In other words, to determine the relative likelihoods of two given states, agent i only relies on her neighbors' opinions about those same states, with the only requirement being that she treats the information provided by her neighbors as gross substitutes.[16] Thus, any such function ψ_i results in a

[15] See Appendix B, 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.

[16] Even though this restriction is imposed on ψ_i, it is easy to verify that the elasticity of substitution between any two arguments of ψ_i coincides with that of f_i.

potentially distinct social learning rule of the form

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

for all $\theta \in \Theta$. For instance, any constant elasticity of substitution (CES) function ψ_i results in a weakly separable learning rule as long as the elasticity parameter is greater than or equal to 1.[17]

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 not only 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 $\psi_i(x) = b_{i0} + \sum_{j\in N_i} b_{ij}x_j$, respectively — but also that they represent its two extremes, with corresponding elasticities of substitution given by 1 and ∞, respectively.

Given its fairly 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. However, our next result shows that, under fairly general conditions, all weakly separable learning rules result in the long-run aggregation of information.

Theorem 5. *Suppose agents' social learning rules satisfy LN, IR, and monotonicity and are weakly separable. Then, all agents learn the underlying state almost surely if ψ_i is homogenous of degree 1 for all i.*

This result thus generalizes Theorems 2 and 4 to the entire class of weakly separable learning rules. Furthermore, Theorem 5 establishes that information aggregation under the log-linear or DeGroot learning rules is not due to their 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) each agent treats the information provided by her various neighbors as substitutes; and (iii) the corresponding ψ_i's are homogeneous of degree 1. The latter assumption, in particular, is key to asymptotic learning, as it plays a role similar to unanimity and non-polarization in our earlier results:[18] it guarantees that the effective weights that any given agent assigns to each independent piece of information are of the same order of magnitude as $t \to \infty$. As such, when ψ_i's are homogenous, signals are neither discarded nor amplified over time. In contrast — and parallel to our results in Theorem 2(a) and 2(b) — in the absence of homogeneity, agents may either mislearn the state or remain uncertain forever.

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

[17]Though note that WS does not require the elasticity parameter to be constant across ψ_i's domain.

[18]In fact, if ψ_i is CES, as is the case in log-linear and DeGroot learning rules, homogeneity of degree 1 is equivalent to the unanimity of the learning rule.

attention to the log-linear learning rule characterized in (4). However, the rate of learning is robust to the intricate details of the learning rule up to a first-order approximation and is solely determined by (i) how information is dispersed throughout the society and (ii) the structure of the social network.[19]

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 Preliminary Definitions

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.[20] 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 large 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 a 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

$$v_i = \sum_{j=1}^{n} v_j a_{ji}, \tag{9}$$

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

[19]Indeed, as we show in an earlier draft of the paper (Jadbabaie, Molavi, and Tahbaz-Salehi, 2013, Proposition 2), the rate of learning of the log-linear learning model characterized in Theorem 6 is a first-order approximation to that of the DeGroot model. This result generalizes to all weakly separable learning rules. A proof is available from the authors upon request.

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

out-central agents put a large weight on her opinion. Similarly, we define the *in-centrality* of agent i as

$$w_i = \sum_{j=1}^{n} w_j a_{ij}. \tag{10}$$

Parallel to our earlier notion, an agent's in-centrality captures the extent of her relience (directly or indirectly) on the information provided by other agents. Finally, note that equations (9) and (10) 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 (Lemma A.1). This condition also guarantees that in- and out-centralities are uniquely determined up to a scaling. We normalize these values by setting $\sum_{i=1}^{n} v_i w_i = 1$.[21]

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 $e_{it}^{\theta} \to 0$ almost surely as $t \to \infty$ whenever learning rules are non-polarizing. We define agent i's *rate of learning* as

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

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

$$\lambda_i^{\theta} = \min_{\hat{\theta}\neq\theta} w_i \sum_{j=1}^{n} v_j h_j(\theta, \hat{\theta}) \tag{11}$$

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 6, 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 (11) 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,

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

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 (11) 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 2. *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 (11) 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 a second bottleneck effect in the learning process, which we refer to as the *identification bottleneck*: the rate of learning is determined by the pair of states $(\theta, \hat{\theta})$ that are hardest to distinguish from one another.

6.3 Information Dispersion and Learning

We now use Theorem 6 and Corollary 2 to study how the interplay between the structural properties of the social network and the dispersion of information determines the rate of learning. Given that different agents learn the underlying state with potentially different rates, we restrict our comparative statics results to the common component of their learning rate, $\min_{\hat{\theta} \neq \theta} \sum_{j=1}^n v_j h_j(\theta, \hat{\theta})$. As highlighted by Corollary 2, this is equivalent to focusing on a social network consisting of agents with unanimous learning rules.

Recall that improving the quality of information available to any given agent can only facilitate learning. In particular, providing agent i with more precise signals increases her corresponding relative entropies on the right-hand side of (11), and as a result leads to faster information aggregation. Therefore, capturing the role of information dispersion on the rate of learning is meaningful only if the aggregate amount of information available throughout the society is kept constant. To capture this idea formally, we define the following concept, with $\ell_i = \{\ell_i^\theta\}_{\theta \in \Theta}$ denoting the signal structure of agent i:

Definition 1. The collection of signal structures $(\ell'_1, \ldots, \ell'_n)$ is a *reallocation* of $(\ell_1, \ldots, \ell_n)$ if there exists a permutation $\sigma : N \to N$ on the set of agents such that $\ell'_i = \ell_{\sigma(i)}$ for all i.

We next introduce an ordering for the informativeness of different agents' signal structures:

Definition 2. Agent i's signal structure is *uniformly more informative* than j's, denoted by $\ell_i \succeq \ell_j$, if $h_i(\theta, \hat{\theta}) \geq h_j(\theta, \hat{\theta})$ for all $\theta, \hat{\theta} \in \Theta$.

When $\ell_i \succeq \ell_j$, agent i has access to signals that are more discriminating between any pairs of states than those available to j. Note that uniform informativeness does not provide a complete ordering over agents' signal structures.[22] However, if agents' signal structures are comparable with respect to $\succeq$, there is an ordering of individuals such that agents who are ranked higher can distinguish between the underlying state θ and any other alternative $\hat{\theta}$ with fewer observations, regardless of the value of θ.

Corollary 3. *Suppose that agents' signal structures are comparable with respect to the uniform informativeness order and that $\ell_i \succeq \ell_j$ if and only if $v_i \geq v_j$. Then, no reallocation of signal structures increases the rate of learning.*

Thus, if the agents' signal structures can be ordered, the rate of learning is highest when the effective attention individuals receive from others is non-decreasing in the informativeness of their signals. In this sense, the rate of social learning is maximized under a "positive assortative matching" of signals and out-centralities. The intuition behind this result is that if an information structure is uniformly more informative than another pair, then by definition, it requires fewer number of observations to distinguish between any pair of states. As a result, allocating such an information structure to a more out-central agent guarantees that, irrespective of the underlying state, the high quality information receives a higher effective attention from the rest of individuals in the network.

Given that uniform informativeness is a partial order over the set of all signal structures, there are scenarios in which the conditions of Corollary 3 are not satisfied. In particular, if say, agent i is better than agent j in distinguishing between a pair of states (measured in terms of relative entropy) but is worse in distinguishing between another, then signal structures ℓ_i and ℓ_j are not comparable with respect to $\succeq$. The next example illustrates why the insights of Corollary 3 do not generalize to environments consisting of such "experts" — i.e., agents who are particularly well-informed about a subset of states but not necessarily about others.

Example 1. Consider a social network consisting of n agents and suppose that the set of states and observations are $\Theta = \{\theta_0, \theta_1, \ldots, \theta_n\}$ and $S = \{\text{Head}, \text{Tail}\}$, respectively. Furthermore, suppose that agents' signal structures are given by

$$\ell_i^\theta(\omega) = \begin{cases} \pi_i & \text{if } \theta = \theta_i, \omega = \text{Head}, \\ \pi_i & \text{if } \theta \neq \theta_i, \omega = \text{Tail}, \\ 1 - \pi_i & \text{otherwise}, \end{cases}$$

where $\pi_i > 1/2$. Thus, the signal structure of agent i enables her to distinguish θ_i from any other state $\theta \neq \theta_i$, whereas the rest of the states are observationally equivalent from her perspective.

[22] This notion is weaker than Blackwell's (1953) well-known criterion of informativeness, according to which a signal structure is more informative than another if any decision-maker prefers the former to the latter in all decision problems. Hence, if ℓ_i is more informative than ℓ_j in the sense of Blackwell, then $\ell_i \succeq \ell_j$, but not vice versa. See Appendix C for a formal derivation of this claim.

Given that θ_0 is observationally equivalent to θ_i from the point of view of all agents $j \neq i$, agent i is effectively the "expert" in learning θ_i. Furthermore, it is easy to see that the ability of agent i to distinguish θ_i from other states is increasing in π_i. In fact, the relative entropy corresponding to agent i's signal structure is

$$h_i(\theta, \hat{\theta}) = \begin{cases} H_i & \text{if } \theta = \theta_i, \hat{\theta} \neq \theta_i, \\ H_i & \text{if } \theta \neq \theta_i, \hat{\theta} = \theta_i, \\ 0 & \text{otherwise,} \end{cases} \tag{12}$$

where $H_i = (2\pi_i - 1)\big[\log \pi_i - \log(1-\pi_i)\big]$ is an increasing function of π_i. Given that H_i is the expected rate at which agent i accumulates evidence in favor of θ_i against any other state $\theta \neq \theta_i$, it essentially captures the level of "expertise" of agent i: a higher value of H_i (or equivalently, π_i) means a greater discrepancy between $\ell_i^{\theta_i}$ and other distributions in i's signal structure.

Equation (12) shows that, conditional on the realization of θ_0, the rate of learning is equal to $\lambda^{\theta_0} = \min_i v_i H_i$. Therefore, among all possible allocations of signal structures to agents, the negative assortative matching of agents' expertise and out-centralities maximizes the rate of learning; that is, speed of learning is maximized if $H_i \geq H_j$ whenever $v_i \leq v_j$ for all pairs of agents i and j. On the other hand, the positive assortative matching of the two leads to the slowest rate of learning.

Example 1 shows that if all agents are experts — in the sense that information structures are such that each individual possesses information crucial for learning — then, the "negative assortative matching" of expertise and out-centralities leads to the fastest rate of learning among all possible allocations of signal structures. Put differently, learning is obtained faster if the least out-central agents have access to signals of the highest quality (π_i closer to 1).

The intuition underlying this result is as follows. Recall from the discussion following Theorem 6 that two distinct bottleneck effects may emerge as impediments to social learning. On the one hand, due to the network bottleneck effect, the information available to the peripheral agents receives less attention from other individuals. On the other hand, the identification bottleneck effect means that the asymptotic rate of learning is determined by the state that is the hardest to distinguish from the realized state. As a result, if the information structures are such that each individual possesses some information that is crucial for learning, positive assortative matching of signal qualities and centralities minimizes the speed of learning. In such a scenario, the two bottleneck effects reinforce one another: the information required for distinguishing the pair of states that are hardest to tell apart is only available to agents that receive very little effective attention from others. As a result, learning the underlying state would take a long time. More concretely, in Example 1, the speed of learning under a positive assortative matching is equal to $v_{\min} H_{\min}$; the smallest value $\lambda^{\theta_0} = \min_i v_i H_i$ can obtain. On the contrary, the speed of learning is maximized under the negative assortative matching of signal qualities and centralities, as such a condition guarantees that the two bottlenecks are as far away from one another as possible.

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 the Bayesian benchmark 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 show that for a fairly large class of learning rules, the form of bounded rationality represented by imperfect recall is not an impediment to asymptotic learning, with a form of unanimity serving as the only requirement on the learning rules. Our characterization results also illustrate 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.

Despite its theoretical take, our paper contributes to the literature that focuses on identifying individual's learning behavior in the laboratory experiments and real-world settings. In particular, by illustrating the behavioral assumptions that underpin DeGroot and other non-Bayesian learning models, it suggests that testing the underlying axioms of a given learning rule can serve as an alternative (and complementary) approach to testing the model's predictions for agents' entire path of actions.

A 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 Bayes' rule,

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

On the other hand, IIA implies that

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

Note that $\mathrm{cond}_{\bar{\Theta}}(\mu)$ 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 $g_i : \mathbb{R}^n \to \mathbb{R}$

such that

$$\log \frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} = g_i \left(\log \frac{\mu_1(\theta)}{\mu_1(\hat{\theta})}, \dots, \log \frac{\mu_n(\theta)}{\mu_n(\hat{\theta})} \right) \tag{13}$$

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 θ, $\hat{\theta}$ and $\tilde{\theta}$. Given that (13) 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 (13) 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 A.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

$$\frac{f_i(\mu)(\theta)}{f_i(\mu)(\hat{\theta})} \geq \frac{\mu_i(\theta)}{\mu_i(\hat{\theta})} \geq 1 \tag{14}$$

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 (14) implies that $(Ay)_i/y_i \geq 1$ for all i. On the other hand, Corollary 6.1.8 of Horn and Johnson (1985, p. 348) 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. □

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,

$$x_t = \rho^t x_0 + \sum_{\tau=1}^{t} \rho^{t-\tau} r(\omega_\tau), \tag{15}$$

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 A.1 guarantees that $\rho > 1$. Therefore, equation(15) implies that

$$\begin{aligned} \rho^{-t} x_t &\leq x_0 + \sum_{\tau=1}^{T-1} \rho^{-\tau} r(\omega_\tau) + \sum_{\tau=T}^{t} \rho^{-\tau} r_{\max} \\ &\leq x_0 + \sum_{\tau=1}^{T-1} \rho^{-\tau} r(\omega_\tau) + \frac{\rho^{-T}}{1-\rho^{-1}} r_{\max} \end{aligned} \tag{16}$$

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 (16) is strictly negative, thus guaranteeing that $\limsup_{t\to\infty} \rho^{-t} x_t < 0$. Therefore,

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

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, 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. □

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, $\lim_{t\to\infty}\mu_{it}(\hat{\theta})=0$. 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 (15) 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 A.1 guarantees that $\rho<1$. Consequently,

$$\limsup_{t\to\infty}\sum_{i=1}^{n} v_i \log\frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} \leq \frac{r_{\max}}{1-\rho}. \tag{17}$$

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 (17) 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 A.1, it is immediate that $\rho=1$ whenever agents' social learning rules are non-polarizing. Consequently, equation (15) implies that

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

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

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

almost surely, and as a result,

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

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

$$\lim_{t\to\infty}\log\left(\frac{\mu_{jt}(\theta)}{\mu_{jt}(\hat{\theta})}\right)=\infty \tag{18}$$

with probability one. Now (4) guarantees that any agent i who has agent j in their neighborhood needs to also satisfy $\log(\mu_{it}(\theta)/\mu_{it}(\hat{\theta}))\to\infty$. Given that the social network is strongly connected, an inductive argument then guarantees that (18) is satisfied for all agents with probability one. This in turn implies that

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

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 $g_i : [0,1]^n \to [0,1]$ such that

$$f_i(\mu)(\theta) = g_i(\mu_1(\theta), \dots, \mu_n(\theta)) \tag{19}$$

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 (19) 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

$$g_i(\mu(\theta)) + g_i(\mu(\hat{\theta})) = 1 - \sum_{\tilde{\theta} \notin \{\theta, \hat{\theta}\}} g_i(\mu(\tilde{\theta})). \tag{20}$$

Note that changing $\mu(\theta)$ and $\mu(\hat{\theta})$ does not impact the right-hand side of (20) as long as $\mu(\tilde{\theta})$ is kept unchanged for all $\tilde{\theta} \notin \{\theta, \hat{\theta}\}$. Consequently,

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

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

$$h_i(x) + h_i(y) = h_i(x+y) \tag{21}$$

for any arbitrary $x, y \in \mathbb{R}^n$ such that $x, y \geq 0$ and $x + y \leq \mathbf{1}$, where $h_i(z) = g_i(z) - g_i(0)$. But (21) 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 b_{ij} such that

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

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

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

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

Proof of Theorem 5

We present the proof by stating and proving a sequence of lemmas. Throughout, we assume that functions $\psi_i : [0,1]^n \to \mathbb{R}$ satisfying (7) are such that $\psi_i(1,\ldots,1) = 1$ for all i. Note that since the social learning rule of agent i can be rewritten as (8), this is simply a normalization assumption and hence is without loss of generality.

To simplify the proofs, we also assume that ψ_i is twice differentiable in all arguments and 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. Finally, we define

$$\sigma_i^{(rj)}(x) = \frac{\psi_i^{(r)}(x)\psi_i^{(j)}(x)}{\psi_i(x)\psi_i^{(rj)}(x)}, \tag{22}$$

where $\psi_i^{(rj)}(x)$ to denote the second-order partial derivate of $\psi_i(x)$ with respect to x_r and x_j. The assumption that the elasticity of substitution between any two arguments of ψ_i is greater than equal to one guarantees that $\sigma_i^{(rj)}(x) \geq 1$ for all x and all $r, j \in N_i$.

We start by establishing a few properties of function ψ_i.

Lemma A.2. *Suppose agent i's learning rule satisfies weak separability and substitutability. Then,*

(a) ψ_i is increasing and jointly concave in all arguments.

(b) $q_i(x) = \log \psi_i \exp(x)$ is jointly convex in all arguments.

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

Proof of part (a) 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 $\sum_{r=1}^n \sum_{j=1}^n y_r y_j \psi_i^{(rj)}(x) \leq 0$ for all $x \in [0,1]^n$ and all $y \in \mathbb{R}^n$. Since ψ_i is homogenous of degree 1, $\psi_i^{(r)}$ is homogenous of degree zero for all r, which in turn implies that $\sum_{j=1}^n x_j \psi_i^{(rj)}(x) = 0$ for all $x \in [0,1]^n$. As a result,

$$\psi_i^{(rr)}(x) = -\frac{1}{x_r}\sum_{j\neq r} x_j \psi_i^{(rj)}(x).$$

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

$$\begin{aligned}
\sum_{r=1}^n \sum_{j=1}^n y_r y_j \psi_i^{(rj)}(x) &= \sum_{r=1}^n \sum_{j\neq r} y_r y_j \psi_i^{(rj)}(x) + \sum_{r=1}^n y_r^2 \psi_i^{(rr)}(x) \\
&= \sum_{r=1}^n \sum_{j\neq r} y_r y_j \psi_i^{(rj)}(x) - \sum_{r=1}^n \sum_{j\neq r} \frac{x_j}{x_r} y_r^2 \psi_i^{(rj)}(x) \\
&= -\frac{1}{2}\sum_{r=1}^n \sum_{j\neq r} \frac{1}{x_r x_j}(y_r x_j - y_j x_r)^2 \psi_i^{(rj)}(x).
\end{aligned}$$

As a result,

$$\sum_{r=1}^{n}\sum_{j=1}^{n} y_r y_j \psi_i^{(rj)}(x) = -\frac{1}{2}\sum_{r=1}^{n}\sum_{j\neq r}\frac{1}{x_r x_j}(y_r x_j - y_j x_r)^2 \frac{\psi_i^{(r)}(x)\psi_i^{(j)}(x)}{\psi_i(x)\sigma_i^{(rj)}(x)},$$

where $\sigma_i^{(rj)}$ is defined in (22). Now part (a) of the lemma and the fact that $\sigma_i^{(rj)}(x) \geq 0$ establish that the right-hand side of the above equation is non-positive, completing the proof. □

Proof of part (b) To show this, we prove that the Hessian of q_i is positive semi-definite. First, note that

$$\frac{\partial q_i}{\partial x_r} = e^{x_r}\frac{\psi_i^{(r)}(\exp(x))}{\psi_i(\exp(x))},$$

and as a result,

$$\frac{\partial^2 q_i}{\partial x_r^2} = \tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{\psi_i(\tilde{x})} - \left(\tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{\psi_i(\tilde{x})}\right)^2 + \tilde{x}_r^2 \frac{\psi_i^{(rr)}(\tilde{x})}{\psi_i(\tilde{x})},$$

where $\tilde{x}_r = \exp(x_r)$. Furthermore, the fact that ψ_i is homogenous of degree 1 implies that $\psi_i^{(r)}$ is homogenous of degree zero, which means that $\sum_{j=1}^{n}\tilde{x}_j\psi_i^{(rj)}(\tilde{x}) = 0$. Consequently,

$$\begin{aligned}\frac{\partial^2 q_i}{\partial x_r^2} &= \tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{\psi_i(\tilde{x})} - \left(\tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{\psi_i(\tilde{x})}\right)^2 - \tilde{x}_r \sum_{j\neq r}\tilde{x}_j \frac{\psi_i^{(rj)}(\tilde{x})}{\psi_i(\tilde{x})} \\ &= \tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{\psi_i(\tilde{x})} - \left(\tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{\psi_i(\tilde{x})}\right)^2 - \tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{[\psi_i(\tilde{x})]^2}\sum_{j\neq r}\tilde{x}_j \frac{\psi_i^{(j)}(\tilde{x})}{\sigma_i^{(rj)}(\tilde{x})} \\ &= \tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{[\psi_i(\tilde{x})]^2}\left[\psi_i(\tilde{x}) - \tilde{x}_r\psi_i^{(r)}(\tilde{x}) - \sum_{j\neq r}\tilde{x}_j \frac{\psi_i^{(j)}(\tilde{x})}{\sigma_i^{(rj)}(\tilde{x})}\right],\end{aligned}$$

where $\sigma_i^{(ij)}$ is given by (22). Now using the homogeneity of ψ_i one more time guarantees that $\psi_i(\tilde{x}) = \sum_{j=1}^{n}\tilde{x}_j\psi_i^{(j)}(\tilde{x})$, leading to

$$\frac{\partial^2 q_i}{\partial x_r^2} = \tilde{x}_r \frac{\psi_i^{(r)}(\tilde{x})}{[\psi_i(\tilde{x})]^2}\sum_{j\neq r}\tilde{x}_j\psi_i^{j}(\tilde{x})\left(1 - \frac{1}{\sigma_i^{(rj)}(\tilde{x})}\right). \tag{23}$$

On the other hand, a similar derivation implies that

$$\frac{\partial^2 q_i}{\partial x_r \partial x_j} = -\tilde{x}_r\tilde{x}_j \frac{\psi_i^{(r)}(\tilde{x})\psi_i^{(j)}(\tilde{x})}{[\psi_i(\tilde{x})]^2}\left(1 - \frac{1}{\sigma_i^{(rj)}(\tilde{x})}\right). \tag{24}$$

Now, the assumption that $\sigma_i^{(rj)}(\tilde{x}) \geq 1$ for all $r, j \in N_i$ and part (a) of the lemma guarantee that $\partial^2 q_i/\partial x_r^2 \geq 0$ for all r whereas $\partial^2 q_i/\partial x_r \partial x_j \leq 0$ for all $r \neq j$. Furthermore, (23) and (24) imply that

$$\frac{\partial^2 q_i}{\partial x_r^2} = -\sum_{j\neq r}\frac{\partial^2 q_i}{\partial x_r \partial x_j}$$

Putting these observations together implies that the Hessian of q_i is a symmetric, diagonally dominant Z-matrix, guaranteeing that it is also an M-matrix. Thus, by Exercise 4.15 of Berman and Plemmons (1979, p. 156), it is positive semi-definite. □

Proof of part (c) Part (a) of the lemma establishes that ψ_i is concave. Therefore, 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 A.3. $\mathbb{E}_t^\theta[\log \mu_{i,t+1}(\theta)] \geq q_i(\log \mu_t(\theta))$, *where* $q_i(x) = \log \psi_i \exp(x)$.

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

$$\mu_{i,t+1}(\theta) = \frac{\ell_i^\theta(\omega_{it+1}) f_i(\mu_t)(\theta)}{\sum_{\hat{\theta}} \ell_i^{\hat{\theta}}(\omega_{it+1}) f_i(\mu_t)(\hat{\theta})}. \tag{25}$$

Therefore,

$$\begin{aligned}\mathbb{E}_t^\theta[\log \mu_{i,t+1}(\theta)] &= \log f_i(\mu_t)(\theta) + \mathbb{E}_t^\theta \log\left(\frac{\ell_i^\theta(\omega_{it+1})}{\sum_{\hat{\theta}} \ell_i^{\hat{\theta}}(\omega_{it+1}) f_i(\mu_t)(\hat{\theta})}\right)\\ &\geq \log f_i(\mu_t)(\theta),\end{aligned}$$

where we are using Jensen's inequality. Since the learning rule satisfies weak separability, it is immediate that

$$\mathbb{E}_t^\theta[\log \mu_{i,t+1}(\theta)] \geq \log \psi_i(\mu_t(\theta)) - \log\left(\sum_{\hat{\theta}\in\Theta} \psi_i(\mu_t(\hat{\theta}))\right).$$

Now, the fact that $\sum_{\hat{\theta}\in\Theta} \mu_t(\hat{\theta}) = \mathbf{1}$ and part (c) of Lemma A.2 imply that $\sum_{\hat{\theta}\in\Theta} \psi_i(\mu_t(\hat{\theta})) \leq 1$, and as a result,

$$\begin{aligned}\mathbb{E}_t^\theta[\log \mu_{i,t+1}(\theta)] &\geq \log \psi_i(\mu_t(\theta))\\ &= q_i(\log \mu_t(\theta)),\end{aligned}$$

where last equality follows from the definition of q_i. $\square$

For the next lemma, we define the function $Q : \mathbb{R}_-^n \to \mathbb{R}_-^n$ as

$$Q(x) = \lim_{t\to\infty} \underbrace{q \circ q \circ \cdots \circ q}_{t \text{ times}}(x), \tag{26}$$

where $q : \mathbb{R}_-^n \to \mathbb{R}_-^n$ is the mapping obtained from concatenating functions $q_i(x) = \log \psi_i \exp(x)$. Note that function Q is well-defined only if the limit in (26) exists for all $x \in \mathbb{R}_-^n$, a statement we prove in Appendix D. We have the following lemma:

Lemma A.4. *Let Q be defined as in (26).*

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

(b) $Q(\alpha\mathbf{1} + x) = \alpha\mathbf{1} + Q(x)$ for all $\alpha > 0$ and all $x \in \mathbb{R}^n_-$.

(c) $Q(\log \mu_t(\theta)) \to Q^*$ as $t \to \infty$ with $\mathbb{P}^\theta$-probability one, where Q^* is finite almost surely.

Proof of part (a) Recall from Lemma A.2 that ψ_i is non-decreasing in all arguments, which implies that q_i is also non-decreasing for all i. The fact that Q is a obtained by composing function q with itself in turn guarantees that Q is also non-decreasing.

To show that Q is convex, note that by part (b) of Lemma A.2, q_i is jointly convex for all i. This, coupled with the fact that q_i are also non-decreasing guarantees that Q is also convex. Convexity of Q then guarantees that Q is also continuous. □

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

Proof of part (c) Recall from Lemma A.3 that $\mathbb{E}^\theta_t[\log \mu_{t+1}(\theta)] \geq q(\log \mu_t(\theta))$. Furthermore, the fact that Q is non-decreasing guarantees that

$$\begin{aligned} Q\left(\mathbb{E}^\theta_t[\log \mu_{t+1}(\theta)]\right) &\geq Q\left(q(\log \mu_t(\theta))\right) \\ &= Q\left(\log \mu_t(\theta)\right), \end{aligned}$$

where the equality is a consequence of the definition of Q. On the other hand, convexity of Q implies that $\mathbb{E}^\theta_t Q\left(\log \mu_{t+1}(\theta)\right) \geq Q\left(\mathbb{E}^\theta_t[\log \mu_{t+1}(\theta)]\right)$. As a result,

$$\mathbb{E}^\theta_t Q\left(\log \mu_{t+1}(\theta)\right) \geq Q\left(\log \mu_t(\theta)\right),$$

which guarantees that $Q\left(\log \mu_t(\theta)\right)$ is an n-dimensional (upper bounded) submartingale and hence, converges almost surely to some vector Q^* almost surely. □

Lemma A.5. *If the underlying state of the world is θ, then $\mu_{it}(\theta)$ converge to zero almost never.*

Proof. Note that the 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 and that $\psi_i^{(j)}$ is positive and uniformly bounded away from zero for all $j \in N_i$, an inductive argument guarantees that $\mu_{it}(\theta) \to 0$ for all agents i. As a consequence, it is immediate that $Q(\log \mu_t(\theta)) \to -\infty$ as $t \to \infty$ on any such path. Now part (c) of Lemma A.4 that such paths have measure zero on the true probability distribution, as the limit Q^* is finite almost surely. □

Lemma A.6. *Suppose θ denotes the underlying state of the world. For almost all paths, there exists a vector $u \in \mathbb{R}^n_{++}$ such that $\lim_{t\to\infty} \mu_{it}(\theta)/\|\mu_t(\theta)\| = u_i$ for all i.*

Proof. Recall from (25) that

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

Therefore, 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

$$\mu_{t+1}(\theta) = C(\omega_{t+1}, \mu_t(\theta))\mu_t(\theta), \tag{27}$$

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 A.5 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

$$\lim_{t\to\infty} \tau_B\left(P_t(\omega^t)\right) = 0$$

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 (27) then implies the result immediately. $\square$

Proof of Theorem 5 We are now ready to prove the theorem. Let θ denote the underlying state of the world and recall from part (c) of Lemma A.4 that $Q(\log \mu_t(\theta)) \to Q^*$ as $t \to \infty$ on almost all paths. On the other hand, Lemma A.6 guarantees that for any given path, there exists a positive vector u such that

$$\log\left(\frac{\mu_{it}(\theta)}{\|\mu_t(\theta)\|}\right) - \log u_i \to 0$$

for all i. As a result, the continuity of V guarantees that

$$Q\Big(\log u + (\log \|\mu_t(\theta)\|)\mathbf{1}\Big) \to Q^*$$

with $\mathbb{P}^\theta$-probability one. Thus, by part (b) of Lemma A.4,

$$(\log \|\mu_t(\theta)\|)\mathbf{1} + Q(\log u) \to Q^* \qquad \mathbb{P}^\theta\text{-a.s.},$$

implying that $\log \|\mu_t(\theta)\|$ converges to some finite limit almost surely. Thus, Lemma A.6 guarantees that $\mu_{it}(\theta) \to \mu_i^*(\theta)$ as $t \to \infty$ 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 now implies that $\mu_i^*(\theta) = 1$ for all i, establishing that all agents learn the underlying state of the world almost surely. □

Proof of Theorem 6

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

$$\mathbb{E}^\theta x_t = A^t x_0 + \sum_{\tau=1}^{t} A^{t-\tau} h(\theta, \hat{\theta}). \tag{28}$$

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,

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

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

$$\frac{1}{t}\bar{x}_{it} \to 0 \tag{29}$$

as $t \to \infty$ almost surely, for all i. On the other hand, (28) implies that

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

where v and w are the left and right eigenvectors of matrix A, respectively, normalized such that $v'w = 1$. Combining the above with (29) thus guarantees that

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

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. $\square$

Proof of Corollary 3

Without loss of generality, suppose that agents are indexed such that $v_1 \geq v_2 \geq \cdots \geq v_n$. Therefore, by assumption, $h_1(\theta,\hat{\theta}) \geq h_2(\theta,\hat{\theta}) \geq \cdots \geq h_n(\theta,\hat{\theta})$ for all $\theta,\hat{\theta} \in \Theta$. Also suppose that $(\ell'_1,\ldots,\ell'_n)$ is a reallocation of $(\ell_1,\cdots,\ell_n)$. Then, by the rearrangement inequality (Hardy, Littlewood, and Pólya, 1952, Theorem 368):

$$\sum_{i=1}^{n} v_i h_i(\theta,\hat{\theta}) \geq \sum_{i=1}^{n} v_i h'_i(\theta,\hat{\theta}),$$

for any given pair of states θ and $\hat{\theta}$, and as a result,

$$\min_{\hat{\theta}\neq\theta}\sum_{i=1}^{n} v_i h_i(\theta,\hat{\theta}) \geq \min_{\hat{\theta}\neq\theta}\sum_{i=1}^{n} v_i h'_i(\theta,\hat{\theta}),$$

completing the proof. $\square$

B Log-Linear Learning under Normality

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 $\hat{\theta}$,

$$\log\frac{\mu_{it}(\theta)}{\mu_{it}(\hat{\theta})} = \frac{\kappa_{it}}{2}(\hat{\theta}-\theta)(\hat{\theta}+\theta-2m_{it}). \tag{30}$$

On the other hand, recall from Theorem 1 that under LN, monotonicity, IIA, and IR, 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 (30) 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

$$m_{it+1} = \sum_{j=1}^n \left(\frac{a_{ij}\kappa_{jt}}{\sum_{r=1}^n a_{ir}\kappa_{rt}} \right) m_{jt} \tag{31}$$

$$\kappa_{it+1} = \sum_{j=1}^n a_{ij}\kappa_{jt}. \tag{32}$$

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 IIA and imperfect recall 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 (31) to update the means of their beliefs depend on t. This observation thus reveals 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 IIA and imperfect recall.

C Blackwell's Ordering and Uniform Informativeness

Blackwell (1953) defines a decision-theoretic notion of what it means for a signal structure to be more informative than another. According to this notion, a signal structure is more informative than another if a decision-maker with any utility function would prefer to use the former over the latter when facing any decision problem.

It is well-known that Blackwell's requirement for the ordering of signal structures is very strong and that most signal structure pairs are not comparable in the sense of Blackwell. One can define a related and weaker notion of informativeness (Jewitt, 2007): an information structure is said to be *Blackwell more informative than another on dichotomies* if the former is Blackwell more informative than the latter on all dichotomous subsets $\{\theta, \hat{\theta}\} \subset \Theta$.

The next result shows that our notion of uniform informativeness defined in Section 6 provides a more complete order over the set of signal structures than either of the notions above.

Proposition C.1. *Suppose that ℓ is Blackwell more informative than ℓ' on dichotomies. Then, ℓ is uniformly more informative than ℓ'.*

Proof. By the theorem of Blackwell and Girshick (1954, p.328), ℓ is Blackwell more informative than ℓ' on dichotomies, if and only if

$$\sum_{s\in S} \ell^{\theta}(s)\phi\left(\frac{\ell^{\hat{\theta}}(s)}{\ell^{\theta}(s)}\right) \geq \sum_{s\in S} \ell'^{\theta}(s)\phi\left(\frac{\ell'^{\hat{\theta}}(s)}{\ell'^{\theta}(s)}\right), \tag{33}$$

for all $\theta, \hat{\theta} \in \Theta$ and all convex functions ϕ. Given that $\phi(x) = -\log(x)$ is convex, this immediately guarantees that informativeness in the sense of Blackwell (on dichotomies) implies informativeness in the uniform sense. □

Notice that the above proof also establishes that the inverse of Proposition C.1 does not hold in general. In particular, for a signal structure to be more informative than another on dichotomies in the sense of Blackwell, inequality (33) should hold for all convex functions ϕ, whereas for uniform informativeness, it is sufficient that (33) is satisfied for $\phi(x) = \log(x)$. Thus, unlike most other information orders (such as Lehmann's (1988)), uniform informativeness does not coincide with Blackwell's order on dichotomies.

D Technical Appendix

Proposition D.1. *The mapping* $Q : \mathbb{R}^n_- \to \mathbb{R}^n_-$ *defined in* (26) *is well-defined.*

Proof. First we prove that

$$\min_{j\in N_i}\{x_j\} \leq \psi_i(x) \leq \max_{j\in N_i}\{x_j\} \tag{34}$$

for all $x \in [0,1]^{N_i}$. Furthermore, both inequalities hold strictly if there exist j and k such that $x_j \neq x_k$. This is a consequence of the assumption that $\psi_i(\mu(\theta))$ is strictly increasing in $\mu_j(\theta)$ with the derivative uniformly bounded away from zero.

Next, define the sequence of functions $\Psi^m : [0,1]^n \to [0,1]^n$ as $\Psi^0(x) = x$ and

$$\Psi^{m+1}(x) = \psi(\Psi^m(x)),$$

where $\psi : [0,1]^n \to [0,1]^n$ is the mapping obtained by concatenating functions ψ_i for all i. From (34) it is immediate that

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

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

We next show that $\underline{x} = \bar{x}$. Let $\Xi : [0,1]^n \to [0,1]$ be

$$\Xi(x) = \max_i x_i - \min_i x_i.$$

This is clearly a continuous function. On the other hand, by (34),

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

with equality if and only if $\max_i x_i = \min_i x_i$, which means that Ξ is a Lyapunov function for ψ.[23] 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 $\Xi(\psi(x)) - \Xi(x) = 0$ whenever $x \in X$. Since $\Xi(\psi(x)) - \Xi(x) = 0$ only if $\max_i x_i = \min_i x_i$, the above argument shows that $\max_i \Psi_i^m(x) - \min_i \Psi_i^m(x)$ must go to zero as m goes to infinity.

We can now define Q as

$$Q(x) = \log \lim_{m\to\infty} \Psi^m(\exp(x)),$$

where the existence of the limit is guaranteed by the above argument. □

[23] For a definition, see, LaSalle (2012, p. 8).

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.

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 (2015), "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.

Blackwell, David (1953), "Equivalent comparisons of experiments." *The Annals of Mathematical Statistics*, 24, 265–272.

Blackwell, David and M. A. Girshick (1954), *Theory of Games and Statistical Decisions*. John Wiley & Sons, New York, NY.

Breza, Emily (2015), "Field experiments, social networks, and development." In *The Oxford Handbook on the Economics of Networks* (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.), Oxford University Press.

Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri (2015), "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 (2015), "Networks in the laboratory." In *The Oxford Handbook on the Economics of Networks* (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.), 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, 221–43.

Eyster, Erik and Matthew Rabin (2014), "Extensive imitation is irrational and harmful." *The Quarterly Journal of Economics*, 129, 1861–1898.

Gagnon-Bartsch, Tristan and Matthew Rabin (2015), "Naive social learning, mislearning, and unlearning." 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 (2015), "Learning in social networks." In *The Oxford Handbook on the Economics of Networks* (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.), Oxford University Press.

Grimm, Veronika and Friederike Mengel (2014), "An experiment on learning in a multiple games environment." Working Paper.

Hardy, Godfrey Harold, John Edensor Littlewood, and George Pólya (1952), *Inequalities*, second edition. Cambridge University Press, London and New York.

Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz (2015), "When more information reduces the speed of learning." Working Paper.

Hartfiel, Darald J. (2002), *Nonhomogenous Matrix Products*. World Scientific Publishing Co. Pte. Ltd., Singapore.

Horn, Roger A. and Charles R. Johnson (1985), *Matrix Analysis*. Cambridge University Press, New York, NY.

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.

Jadbabaie, Ali, Pooya Molavi, and Alireza Tahbaz-Salehi (2013), "Information heterogeneity and the speed of learning in social networks." Columbia Business School Working Paper No. 13-28.

Jewitt, Ian (2007), "Information order in decision and agency problems." Working Paper.

LaSalle, Joseph P (2012), *The stability and control of discrete processes*, volume 62. Springer Science & Business Media.

Lehmann, Erich Leo (1988), "Comparing location experiments." *The Annals of Statistics*, 16, 521–533.

Lehrer, Keith and Carl Wagner (1981), *Rational Consensus in Science and Society, A Philosophical and Mathematical Study*. D. Reidel Publishing Company, Dordrecht, Holland.

Lobel, Ilan and Evan Sadler (2015a), "Information diffusion in networks through social learning." *Theoretical Economics*, 10, 807–851.

Lobel, Ilan and Evan Sadler (2015b), "Preferences, homophily, and social learning." *Operations Research*.

Möbius, Markus, Tuan Phan, and Adam Szeidl (2015), "Treasure hunt." 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 (2015), "Social learning with costly search." *American Economic Journal: Microeconomics*.

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.