Essays On Learning in Social Networks

Pooya Molavi

A Dissertation

in

Electrical and Systems Engineering

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2013

Supervisor of Dissertation

____
Ali Jadbabaie, Alfred Fitler Moore Professor of Electrical and Systems Engineering

Graduate Group Chair

____
Saswati Sarkar, Professor of Electrical and Systems Engineering

Dissertation Committee
George Pappas, Joseph Moore Professor of Electrical and Systems Engineering
Andrew Postlewaite, Harry P. Kamen Professor of Economics and Professor of Finance
Alejandro Ribeiro, Assistant Professor of Electrical and Systems Engineering
Alireza Tahbaz-Salehi, Assistant Professor of Decision, Risk, & Operations, Columbia University

ACKNOWLEDGMENTS

I am grateful to my adviser Ali Jadbabaie who has been extremely supportive during these years and has taught me a lot. I would also like to thank the members of my committee, Professors George Pappas, Andrew Postlewaite, Alejandro Ribeiro, and Alireza Tahbaz-Salehi, without whose guidance and help this thesis would not have been possible. I also take the opportunity to thank Alvaro Sandroni and Alireza Tahbaz-Salehi for their contributions in Chapters 2 and 3 and Ceyhun Eksin and Alejandro Ribeiro for their contributions in Chapter 4.

Last but not least, I would like to thank my parents, Akram and Hassan, and my brother, Nima, who gave me their unconditional love and support throughout my life.

ABSTRACT

ESSAYS ON LEARNING IN SOCIAL NETWORKS

Pooya Molavi
Ali Jadbabaie

Over the past few years, online social networks have become nearly ubiquitous, reshaping our social interactions as in no other point in history. The preeminent aspect of this social media revolution is arguably an almost complete transformation of the ways in which we acquire, process, store, and use information. In view of the evolving nature of social networks and their increasing complexity, development of formal models of social learning is imperative for a better understanding of the role of social networks in phenomena such as opinion formation, information aggregation, and coordination. This thesis takes a step in this direction by introducing and analyzing novel models of learning and coordination over networks. In particular, we provide answers to the following questions regarding a group of individuals who interact over a social network: 1) Do repeated communications between individuals with different subjective beliefs and pieces of information about a common true state lead them to eventually reach an agreement? 2) Do the individuals efficiently aggregate through their social interactions the information that is dispersed throughout the society? 3) And if so, how long does it take the individuals to aggregate the dispersed information and reach an agreement? This thesis provides answers to these questions given three different assumptions on the individuals' behavior in response to new information. We start by studying the behavior of a group of individuals who are fully rational and are only concerned with discovering the truth. We show that communications between rational individuals with access to complementary pieces of information eventually direct everyone to discover the truth. Yet in spite of its axiomatic appeal, fully rational agent behavior may not be a realistic assumption when dealing with large societies and complex networks due to the extreme computational complexity of Bayesian inference. Motivated by this observation, we next explore the implications of bounded rationality by introducing biases in the way agents interpret the opinions of others while at the same time maintaining the assumption that agents interpret their private observations rationally. Our analysis yields the result that when faced with overwhelming evidence in favor of the truth even biased agents will eventually learn to discover the truth. We further show that the rate of learning has a simple analytical characterization in terms of the relative entropy of agents' signal structures and their eigenvector centralities and use the characterization to perform comparative analysis. Finally, in the last chapter of the thesis, we introduce and analyze a novel model of opinion formation in which agents not only seek to discover the truth but also have the tendency to act in conformity with the rest of the population. Preference for conformity is relevant in scenarios ranging from participation in popular movements and following fads to trading in stock market. We argue that myopic agents who value conformity do not necessarily fully aggregate the dispersed information; nonetheless, we prove that examples of the failure of information aggregation are rare in a precise sense.

Contents

List of Figures

Chapter 1

Overview

Over the past few years, online social networks have become nearly ubiquitous, reshaping our social interactions as in no other point in history. The preeminent aspect of this "social media revolution" is arguably an almost complete transformation of the ways in which we acquire, process, store, and use information. News is no longer gathered exclusively by journalists and reported by traditional media, but emerges from information exchanges in a complex ecosystem comprising sources, journalists, and viewers. Increasingly, the distinction between providers and users of information is blurred as more individuals participate in creation and curation of content. A Pew Research Center survey published in March 2010 found that 37% of American internet users, or 29% of the population, had "contributed to the creation of news, commented about it or disseminated it via postings on social-media sites like Facebook or Twitter."[1]

The increasing pervasiveness of online social networks has had significant real world repercussions as far-reaching as in catalyzing the recent wave of popular protests in Egypt, Iran, Tunisia, and more recently Brazil and Turkey. These large-scale protests are widely believed to have been impossible without the help of social media—to such an extent that they are dubbed "Twitter revolutions" by some. Social networks act as conduits of information about the time and location of protests and accounts of the events. They counter inflammatory or complacent official sources of information. They also help coordinate the protests by providing individuals with information on participation decisions of others or presence and forcefulness of police at a given location.[2]

Understanding the role of networks in the dissemination and aggregation of information and facilitation of coordination among individuals is important not only in providing insights on social phenomena such as uprisings, but also from a normative point of view. For example, various development or public health projects in the developing world rely on the power of offline social networks in spreading information.[3] Similarly, successful decision making and collaboration in modern, complex organizations relies heavily on decentralized information sharing mechanisms among different divisions within the organization. A better understanding of the interplay of individual interactions and outcomes such as learning and coordination may thus be crucial in designing effective public policy or organizational structure.

In view of the evolving nature and role of social networks and their increasing complexity, development of formal models of "social learning" is imperative for a better understanding of the

[1] "The people formerly known as the audience," *The Economist*, July 7th 2011.

[2] "The digital demo," *The Economist*, June 29th 2013.

[3] Examples include education programs on deworming (Miguel and Kremer (2003)) and introduction of biofortified agricultural technologies (McNiven and Gilligan (2011)) or microfinance programs (Banerjee, Chandrasekhar, Duflo, and Jackson (2012)).

role of social networks in phenomena such as opinion formation, information aggregation, and coordination. This thesis takes a step in this direction by introducing and analyzing novel models of learning and coordination over networks. In particular, we are interested in providing answers to the following questions regarding a group of individuals who interact over a social network: 1) Do repeated communications between individuals with different subjective beliefs and pieces of information about a common true state lead them to eventually reach an agreement? 2) Do the individuals necessarily agree on what they would have agreed on if they all had knowledge of every piece of information available to every other member of the society? Said differently, do the individuals efficiently aggregate through their social interactions the information that is dispersed throughout the society? 3) And if so, how long does it take the individuals to aggregate the dispersed information and reach an agreement?

The answers to these questions crucially depend on the way agents behave in response to new information. This thesis explores the implications of several behavioral assumptions. We start by studying the behavior of a group of individuals who are fully rational and who are only concerned with discovering the truth and show that communications between rational individuals with access to complementary pieces of information eventually direct everyone to discover the truth. Yet in spite of its axiomatic appeal, fully rational agent behavior may not be a realistic assumption when dealing with large societies and complex networks, due to the extreme computational complexity of rational reasoning.[4] Motivated by this observation, we next explore the implications of bounded rationality by introducing biases in the way agents interpret the opinions of others while at the same time maintaining the assumption that agents interpret their private observations rationally. Our analysis yields the result that when faced with overwhelming evidence in favor of the truth even biased agents will eventually learn to discover the truth. Finally, in the last chapter of the thesis, we introduce and analyze a novel model of opinion formation in which agents not only seek to discover the truth but also have the tendency to act in conformity with the rest of the population. Preference for conformity is relevant in scenarios ranging from participation in popular movements and following fads to trading in a stock market. We argue that, unlike as in the models previously mentioned, myopic agents who value conformity do not necessarily aggregate the dispersed information efficiently; nonetheless, we prove that examples of the failure of information aggregation are rare, in a sense that is to be formalized later.

Although the details of the environment—beyond connectivity of the network and identifiability of the unknown parameter—are irrelevant for asymptotic aggregation of information, the rate of information aggregation crucially depends on them. To capture these details, we introduce two novel orders that are pertinent to problems in learning and networks: the first one ranks signal structures in terms of informativeness, and the second one ranks networks in terms of symmetry. We use these orders—in addition to a well-known measure of an individual's centrality in a network—to characterize the rate of learning for boundedly rational individuals. The analysis yields interesting insights on the implications of the interplay between information and network structures for social learning. For instance, we show that whether a symmetric network is conducive to learning is contingent upon the informativeness of individuals' signal structures.

[4] For a discussion of the complexities of fully rational social learning, see, for instance, DeMarzo, Vayanos, and Zwiebel (2003).

Related Literature

The first formal models of learning and agreement in presence of dispersed information are arguably due to DeGroot (1974) and Aumann (1976). In his seminal paper, DeGroot considers a group of individuals who must act together as a team, and such that each individual in the group has her own subjective probability distribution for the unknown value of some parameter. He presents a model which describes how the group can reach agreement on a common subjective probability distribution for the parameter: each individual assigns a fixed weight to the opinion of any individual in the group and repeatedly updates her subjective probability by computing the weighted average probability distribution across the individuals in the group. It is shown that this process leads the individuals in the group to asymptotically reach an agreement in their subjective probabilities of the parameter. In this model, it is assumed that there is no possibility of learning whether the opinion of one individual is closer to the truth than that of another. It is also assumed that, at the beginning, each individual chooses the weights that she is going to use and she then continues to use these weights throughout the process. The individuals in the DeGroot model are thus naïve: they do not use the new information obtained through the updating process to learn the identity of the individuals whose beliefs are closer to the truth and to modify the weights assigned to them accordingly.

Aumann (1976) takes the alternative approach of modeling individuals as perfectly rational. He considers two agents who are endowed with equal prior beliefs about the unknown parameter but form posteriors based on different pieces of information. Aumann argues that if the individuals (truthfully) inform each other of these posteriors and if the posteriors are not equal, a revision in the posteriors may be called for. He shows that if the agents are rational, the process of exchanging information on the posteriors will continue until these posteriors are equal. Said differently, individuals with equal priors "cannot agree to disagree," even if they base their beliefs on different information. Geanakoplos and Polemarchakis (1982) explicitly model the process of exchange and revision of posterior beliefs for two individuals with equal priors and show that the individuals eventually converge to a common posterior equilibrium. They further argue that the agents generically agree on a posterior they would have agreed on, had information been directly exchanged.

Both DeGroot's naïve learning model and the fully rational model of Aumann have been substantially extended to model opinion formation, learning, and consensus over social networks. DeMarzo, Vayanos, and Zwiebel (2003) introduce a DeGroot-based model of persuasion bias, the failure of agents to account for possible repetition in the information they receive, and use the model to explain phenomena such as social influence and unidimensional opinions. Acemoglu, Ozdaglar, and ParandehGheibi (2010) use a variant of DeGroot's model with forceful agents, who influence others without being influenced back, to explain the spread of misinformation in social networks. Golub and Jackson (2010) argue that, absent disproportionately influential agents, all agents' beliefs in large societies converge to the truth, a phenomenon they refer to as "wisdom of crowds." In Golub and Jackson (2012b), they show that homophily, the tendency of agents to associate disproportionately with those having similar traits, can slow the rate of convergence to a consensus. Bala and Goyal (1998) consider an alternative model of bounded rationality where agents do not make inferences about the information of the agents they do not directly observe. They show that, in a strongly connected network, local interactions ensure that all agents asymptotically obtain the same payoffs. Furthermore, they develop conditions on the distribution of prior beliefs, the structure of the social network, and the informativeness of observations under which the agents' limit actions are optimal.

The fully Bayesian learning mode has also been extended to a network setting by a number of papers among which are Gale and Kariv (2003), Rosenberg, Solan, and Vieille (2009), and Mueller-Frank (2013). Building on the work of Gale and Kariv, Rosenberg *et al.* consider a general setting in which a number of agents repeatedly observe signals that may be informative about an unknown parameter or the actions of others. They prove that—with forward-looking agents—eventually all motives for experimentation disappear. They also provide tight conditions on the network and information structures under which agents eventually reach a consensus in their payoffs. Mueller-Frank (2013) provides a framework for studying learning over social networks given general behavioral assumptions. He provides conditions on the agents' choice correspondences and their information structures under which rational learning leads to global consensus, local indifference, and local disagreement. He also provides a counterpart to the result of Geanakoplos and Polemarchakis (1982) (on optimal aggregation of information for agents with finite prior partitions of the state space) to a setting where agents interact over a social network.

Contributions

We start by introducing a model of Bayesian social learning in Chapter 2. Consider a group of individuals who are endowed with a common prior belief about an unknown state of the world that they attempt to learn. Agents repeatedly observe private signals that may be informative about the underlying state. They also observe the posteriors of their neighbors in a social network. The agents are assumed to have incomplete knowledge of the realized social network. However, they are endowed with a common prior over the set of all possible social networks and each is informed of the identities of her neighbors. In Section 2.2, we argue that if the realized social network is strongly connected—that is, if there exists a directed path for the flow of information from any agent to any other one—and if the agents do not face a global identification problem, then all agents asymptotically almost surely learn the realized state of the world. In other words, agents with common priors not only "can't disagree forever" but also agree on the posterior they would have agreed upon had information been directly exchanged. This result goes beyond the conclusions of Geanakoplos and Polemarchakis (1982) and Mueller-Frank (2013) in two important ways. First, we show that for the agents to reach an agreement it is not necessary for their private information to correspond to a finite partition of the state space. Rather, agents reach an agreement even if they continuously receive new private signals. Second and more importantly, the agents in our model reach an agreement in spite of having incomplete knowledge of the realized network and simply by communicating their posterior beliefs over the state space—and not their beliefs about the topology of the social network. In contrast, Mueller-Frank (2013) assumes the underlying social network to be commonly known by the agents.

The Bayesian framework introduced in Chapter 2 serves mainly as a benchmark that characterizes the behavior of fully rational individuals. The need for Bayesian agents to reason about the source and quality of the information obtained by their neighbors without having complete knowledge of the network structure significantly complicates the required calculations for Bayesian updating of beliefs. The complications with Bayesian learning persist even when individuals have complete information about the network structure, as they still need to perform deductions about the information of every other individual in the network while only observing the evolution of opinions of their neighbors.

Motivated by the complexities of Bayesian social learning, in Chapter 3, we study the evolution of beliefs in a non-Bayesian model of social learning introduced by Tahbaz-Salehi (2009) and Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012). The model is similar to the Bayesian model introduced in Chapter 2 with the exception that, instead of performing Bayesian updates, agents apply a simple learning rule to incorporate the views of individuals in their social neighborhoods: Agents first incorporate their private signals in a fully Bayesian manner as an intermediate step. They then combine this Bayesian posterior with the beliefs of their neighbors in a naïve way by updating their beliefs to a convex combination of the beliefs of their neighbors and their Bayesian posteriors. The model can thus be seen as a natural generalization of the DeGroot (1974) model to the case that agents make repeated observations over time. This is most evident by considering a scenario where signals are uninformative beyond the first time period—in which case the update rule reduces to DeGroot's.

Our analysis of Tahbaz-Salehi's model yields several important insights. As we argue in Section 3.2, the evolution of the individuals' beliefs in the non-Bayesian model asymptotically coincides with that of the Bayesian benchmark introduced in Chapter 2. More specifically, given a strongly connected social network and an identifiable state of the world, each agent eventually learns the true underlying state as if she were aware of the observations of all agents and updated her beliefs according to Bayes' rule. Thus, with a constant flow of new information, the key condition for social learning is that individuals take their personal observations into account in a Bayesian way. Repeated communications over the social network ensures that the idiosyncratic differences eventually disappear and learning is obtained. In Section 3.3, we show that the rate of learning has a simple analytical characterization in terms of the relative entropy of agents' signal structures and their eigenvector centralities. Our characterization establishes that the way information is dispersed throughout the social network has nontrivial implications for the rate of learning. In particular, we introduce a novel partial order on signal structures and show that when the informativeness of different agents' signal structures is comparable given this order, then a positive assortative matching of signal qualities and eigenvector centralities maximizes the rate of learning. On the other hand, 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. In Section 3.5, we introduce a novel notion of network symmetry and use it to argue that the extent of asymmetry in the structure of the social network plays a key role in the long run dynamics of the beliefs.

In Chapter 4, we extend the benchmark Bayesian learning framework to a setting where agents face a payoff externality—in addition to the informational externality which is characteristic of social learning models—that stems from a coordination motive. Consider a group of individuals who attempt to take actions which are close to an unknown state of the world while, at the same time, trying to choose actions which are similar to the average action across the rest of the population. Such agents face a trade-off between acting according to their best estimates of the state and trying to coordinate with other agents. Such trade-offs are important in decisions whether to participate in popular protests, trade decisions in financial markets (Morris and Shin (2002)), consumption decisions (Bramoullé, Kranton, and D'Amours (2009)), and organizational coordination (Calvó-Armengol and Beltran (2009)). The decisions of traders in a stock market, for example, depend on their beliefs about the fundamental stock values; nonetheless, traders also tend to consider how other traders will behave as their decisions could directly affect the gains from trade. In all of these scenarios, agents make decisions by attempting to second-guess the decisions of others while also trying to guess the value of an unknown.

We use the framework of dynamic games of incomplete information to model the agents' problem. Agents' preferences are represented by payoff functions, similar to the ones used by Morris and Shin (2002) and Calvó-Armengol and Beltran (2009), that have two components: an estimation term capturing the agents' preference to choose actions that are close to the realized state and a coordination term capturing their tendency to act in conformity with others. The game is played over multiple stages. At each stage of the game, each agent takes an action and observes a private signal and the previous choices made by others in her social neighborhood. An agent's action may reveal some information to her neighbors that was previously unknown to them. The neighbors can use this information to reevaluate their beliefs about the underlying parameter and their predictions of others' future behavior. These reevaluations may, in turn, lead agents to revise their actions over time.

We define the (weak perfect Bayesian) equilibrium assuming that agents are myopic and proceed to prove formal results regarding the agents' asymptotic equilibrium behavior. First, we show that each agent's action asymptotically converges to some limit action. We then use this result to prove that if the social network is sufficiently connected over time, agents asymptotically receive similar payoffs and choose similar actions. That is, agents eventually coordinate on the same action. Second, we show that if the agents' private observations are only functions of the unknown state (and not their own actions), then generically the agents eventually coordinate on the action on which they would have coordinated if they had directly exchanged their private signals. This result extends the main theorem of Chapter 2, as well as Theorem 4 of Mueller-Frank (2013), on optimal aggregation of information, to the cases where the state space is not finite and the agents face payoff externalities induced by a coordination motive.

CHAPTER 2

BAYESIAN SOCIAL LEARNING WITH INCREASING INFORMATION

In this chapter, we study the problem of aggregation of information on a strongly connected social network when individuals incorporate the views of their neighbors in a fully Bayesian manner. Consider a group of agents who repeatedly observe private signals that are potentially informative about an underlying state of the world. We assume that agents' private observations are drawn from different distributions and that they may be only partially informative about the state. We also assume that individuals cannot observe the beliefs held by all members of the society. Rather, they only have access to the opinions of the agents in their social neighborhood. Agents are not informed of the structure of the underlying social network beyond their immediate neighborhood. The agents use the information contained in these observations to update their beliefs according to Bayes' rule.

We show that if the realized social network is almost surely strongly connected and the underlying state of the world is globally identifiable, agents asymptotically almost surely learn the realized state of the world. Thus, in spite of the agents' incomplete information about the underlying social network and the local nature of their interactions, the information content of the agents' private signals is eventually fully aggregated. In other words, each agent's belief at the end of the learning process is the same as what it would have been if she had had direct access to the private observations of everyone in the society.

The work in this chapter is related to the literature on Bayesian learning over networks. The focus of the social learning literature is on modeling the way agents use their observations to update their beliefs and characterizing the outcomes of the learning process. Examples include, Bikhchandani, Hirshleifer, and Welch (1992), Banerjee (1992), and Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) that study sequential decision problems; and Borkar and Varaiya (1982), Gale and Kariv (2003), Rosenberg, Solan, and Vieille (2009), Lobel and Sadler (2012), Mueller-Frank (2013), and Acemoglu, Bimpikis, and Ozdaglar (forthcoming) that study repeated and simultaneous interactions. Our model belongs to the latter class of models. Borkar and Varaiya (1982), Gale and Kariv (2003), and Rosenberg, Solan, and Vieille (2009) study models of repeated interactions similar to the one introduced in this chapter. Their focus is, however, on providing conditions under which agents reach consensus in their actions. Lobel and Sadler (2012) and Acemoglu, Bimpikis, and Ozdaglar (forthcoming) study social learning under aggregate network uncertainty and endogenous costly network formation, respectively. The work most closely related to ours is by Mueller-Frank (2013) who studies optimal social learning under general choice correspondences, given that the agents' private information can be represented by a finite partition of the state space. In contrast, in our model, agents continuously make new private observations enabling them to form increasingly finer partitions of the state space.

2.1 Model

2.1.1 Agents and Observations

Consider a collection of n individuals, denoted by $N = \{1, 2 \ldots, n\}$, who are attempting to learn an underlying state of the world $\theta \in \Theta$. The underlying state is drawn at $t = 0$ according to probability distribution $\nu \in \Delta\Theta$ with full support over Θ, which we assume to be finite.

Even though the realized state of the world remains unobservable to the individuals, they make repeated noisy observations about θ over discrete time. In particular, at each time period $t \in \mathbb{N}$ and conditional on the realization of state θ, agent i observes a private signal $s_{it} \in S_i$ which is drawn according to distribution $\ell_i^\theta(\cdot) \in \Delta S_i$, denoted the signal structure of agent i. We assume that the signal space S_i is finite. The realized signals are independent across individuals and over time. Let $s_t = (s_{1t}, \ldots, s_{nt}) \in S = \times_{i=1}^n S_i$ denote the signal profile realized at time t and denote the infinite signal profile sequence by $s = (s_1, s_2, \ldots) \in S^{\mathbb{N}}$.

We do not require the observations to be informative about the state. In fact, each agent may face an identification problem, in the sense that she might not be able to distinguish between two states. We say two states are observationally equivalent from the point of view of an agent if the conditional distributions of her signals under the two states coincide. More specifically, the elements of the set $\Theta_i^\theta = \{\tilde{\theta} \in \Theta : \ell_i^{\tilde{\theta}} = \ell_i^\theta\}$ are observationally equivalent to state θ from the point of view of agent i. However, we impose the following global identifiability assumption on the agents' signal structures.

Assumption 2.1. There are no two states that are observationally equivalent from the point of view of all agents, that is, $\Theta_1^\theta \cap \cdots \cap \Theta_n^\theta = \{\theta\}$ for all $\theta \in \Theta$.

The above assumption guarantees that even though each agent may face an identification problem in isolation, there is no global identification problem: agents' observations are informative enough to allow them to collectively learn the underlying state.

2.1.2 Social Network

At every time period, in addition to her private observation, each agent also has access to the beliefs of a subset of other agents, denoted by $N_i \subseteq N$ and called her neighbors. The interactions between agents can be summarized by a directed network $g = [g_{ji}] \in G = \{0, 1\}^{n \times n}$ where $g_{ji} = 1$ if and only if agent j is a neighbor of agent i, that is, if $j \in N_i$. The social network is generated at $t = 0$ according to some probability distribution $\psi \in \Delta G$ independently of other random variables in the model. Each agent i is informed of her social neighborhood N_i at time $t = 0$.

A directed path in the social network from vertex i to vertex j is a sequence of vertices starting with i and ending with j such that each vertex is a neighbor of the next vertex in the sequence. We say the social network is strongly connected if there exists a directed path from each vertex to any other.

Assumption 2.2. The underlying social network is strongly connected with ψ-probability one.

The above assumption simply guarantees that information can flow from any agent in the network to any other.

2.1.3 Learning Rule

Let $(\Omega, \mathscr{B})$ be a measurable space, where $\Omega = \Theta \times S^{\mathbb{N}} \times G$ is the space containing the realizations of the underlying state of the world, the sequence of signal profiles over time, and the realization of the social network, and $\mathscr{B}$ is the σ-algebra of subsets of Ω. Conditional on the realization of state θ, sample paths $s \in S^{\mathbb{N}}$ are realized according to the probability distribution

$$P^{\theta} = (\ell_1^{\theta} \times \cdots \times \ell_n^{\theta})^{\mathbb{N}},$$

where $\mathbb{N}$ denotes the set of natural numbers. We use $\mathbb{P}$ to denote the probability distribution over Ω induced by ν, $\{\ell_i^{\theta}\}_{\theta \in \Theta, i \in N}$ and ψ, i.e., the probability distribution defined as

$$\mathbb{P}(\theta, B, g) = \int_B \nu(\theta) P^{\theta}(ds)\psi(g),$$

for any $\theta \in \Theta$, $g \in G$, and measurable set $B \subseteq S^{\mathbb{N}}$.

Agents incorporate their private signals and views of their neighbors in a fully Bayesian manner. The posterior belief of agent i on Θ and the σ-algebra generated by all her observations up to time period t—which consists of her private signals and the sequence of beliefs of her neighbors up to time t—are defined recursively as follows:

$$\begin{aligned} \mathscr{H}_{it} &= \sigma\left(s_{i1}, \ldots, s_{it}, (\mu_{j1}, \ldots, \mu_{jt-1})_{j \in N_i}\right) \\ \mu_{it}(\theta) &= \mathbb{P}(\theta | \mathscr{H}_{it}), \end{aligned}$$

with $\mu_{i0} = \nu$. We also use $\mathscr{H}_i$ to denote the smallest σ-algebra that contains $\mathscr{H}_{it}$ for all t.

Finally, we assume that the description of the model, as presented in this section, is common knowledge.

2.2 Bayesian Social Learning

Given the Bayesian model described in the previous section, we are interested in the long run behavior of the agents' beliefs and in particular whether they learn the true realized state. The following theorem shows that each agent asymptotically assigns probability of one to the realized state even if no agent can identify the realized state on her own—as long the agents do not face a global identification problem.

Theorem 2.1. *Suppose that Assumptions 2.1 and 2.2 hold. Then,*

$$\mu_{it}(\cdot) - \mathbf{1}_{\theta}(\cdot) \longrightarrow 0 \qquad \mathbb{P}\text{-a.s.},$$

for all i, where $\mathbf{1}_{\theta}(\cdot)$ is the degenerate probability measure that assigns a unit mass on the underlying state θ.

Thus, Bayesian agents will eventually learn the true parameter as long as the network is strongly connected and no two states are observationally equivalent from the point of view of all agents. We remark that it is critical for agents to assign a nonzero prior belief on all states. Clearly, if a Bayesian agent considers a state θ to be impossible then no new information would convince her that θ is the underlying state of the world.

The intuition behind Theorem 2.1 is simple. Consider an arbitrary agent i in the social network. Given any $\theta \in \Theta$, agent i can asymptotically learn—using only her private signals—whether the realized state belongs to the set Θ_i^θ. Furthermore, this fact is common knowledge. Therefore, any agent j, who is a neighbor of agent i and observes i's beliefs, also learns whether the realized state belongs to Θ_i^θ. A similar argument shows that the agents who observe j and, inductively by the strong connectivity of the network, all other agents can asymptotically learn whether the event Θ_i^θ is realized. The same is true for any such event Θ_i^θ with $i \in N$ and $\theta \in \Theta$. Therefore, by the global identifiability assumption, for each $\theta \in \Theta$, every agent asymptotically learns whether θ is realized.

2.3 Proof of Theorem 2.1

We define $\mathscr{P}_{it}$ as the σ-algebra generated by the private signals observed by agent i up to period t,

$$\mathscr{P}_{it} = \sigma(s_{i1}, \ldots, s_{it}),$$

and $\mathscr{P}_i$ as the smallest σ-algebra containing $\mathscr{P}_{it}$ for all t. We also define $\mathscr{U}_{it}$ as the σ-algebra generated by the sequence of beliefs of agent i up to period t,

$$\mathscr{U}_{it} = \sigma(\mu_{i1}, \ldots, \mu_{it}),$$

and $\mathscr{U}_i$ as the smallest σ-algebra containing $\mathscr{U}_{it}$ for all t. Finally, we let $l_{it}(\cdot) \in \Delta S_i$ be the empirical distribution of the signals observed by agent i up to time t and let $B_i^\theta \in \mathscr{P}_i$ be the event

$$B_i^\theta = \left\{\omega \in \Omega : l_{it} \to \ell_i^\theta \quad \text{as} \quad t \to \infty \quad \text{given } \omega\right\}.$$

By the strong law of large numbers,

$$\mathbb{1}_{\Theta_i^\theta} = \mathbb{1}_{B_i^\theta} \qquad \mathbb{P}\text{-a.s.}$$

Therefore,

$$\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{P}_i\right] = \mathbb{E}\left[\mathbb{1}_{B_i^\theta} \middle| \mathscr{P}_i\right] = \mathbb{1}_{B_i^\theta} = \mathbb{1}_{\Theta_i^\theta} \qquad \mathbb{P}\text{-a.s.,}$$

where the second equality is a consequence of the fact that $B_i^\theta \in \mathscr{P}_i$. Hence, the fact that $\mathscr{P}_i \subseteq \mathscr{H}_i$ implies that

$$\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_i\right] = \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{P}_i\right] \middle| \mathscr{H}_i\right] = \mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{P}_i\right] = \mathbb{1}_{\Theta_i^\theta} \qquad \mathbb{P}\text{-a.s.}$$

Thus, by Lévy's zero-one law, for any $\theta \in \Theta$,

$$\mu_{it}\left(\Theta_i^\theta\right) = \mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_{it}\right] \longrightarrow \mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_i\right] = \mathbb{1}_{\Theta_i^\theta} \qquad \mathbb{P}\text{-a.s.} \tag{2.1}$$

In other words, agent i asymptotically rules out states that are not observationally equivalent to the underlying state of the world from her point of view.

Next, consider agent j such that $i \in N_j$. For any $\theta \in \Theta$, we have,

$$\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{U}_{it}\right] = \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_{it}\right] \middle| \mathscr{U}_{it}\right] = \mathbb{E}\left[\mu_{it}\left(\Theta_i^\theta\right) \middle| \mathscr{U}_{it}\right] = \mu_{it}\left(\Theta_i^\theta\right),$$

where the first equality is a consequence of $\mathscr{U}_{it} \subseteq \mathscr{H}_{it}$,[5] and the third equality is due to the fact that μ_{it} is measurable with respect to $\mathscr{U}_{it}$. Therefore, by (2.1),

$$\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{U}_{it}\right] \longrightarrow \mathbb{1}_{\Theta_i^\theta} \qquad \mathbb{P}\text{-a.s.}$$

On the other hand, by Lévy's zero-one law, $\mathbb{E}[\mathbb{1}_{\Theta_i^\theta}|\mathscr{U}_{it}]$ converges to $\mathbb{E}[\mathbb{1}_{\Theta_i^\theta}|\mathscr{U}_i]$ with $\mathbb{P}$-probability one, which implies that $\mathbb{E}[\mathbb{1}_{\Theta_i^\theta}|\mathscr{U}_i] = \mathbb{1}_{\Theta_i^\theta}$ with $\mathbb{P}$-probability one. Hence, for any agent j such that $i \in N_j$, the fact that $\mathscr{U}_i \subseteq \mathscr{H}_j$ implies that

$$\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_j\right] = \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{U}_i\right] \middle| \mathscr{H}_j\right] = \mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{U}_i\right] = \mathbb{1}_{\Theta_i^\theta} \qquad \mathbb{P}\text{-a.s.}$$

Thus, by Lévy's zero-one law,

$$\mu_{jt}\left(\Theta_i^\theta\right) = \mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_{jt}\right] \longrightarrow \mathbb{E}\left[\mathbb{1}_{\Theta_i^\theta} \middle| \mathscr{H}_j\right] = \mathbb{1}_{\Theta_i^\theta},$$

with $\mathbb{P}$-probability one. Therefore, not only agent j would be able to eventually distinguish between any two states that are not observationally equivalent from her own point of view, but also between those that are not observationally equivalent from the points of view of any of her neighbors.

Finally, the fact that the social network is strongly connected guarantees that every individual would be able to eventually distinguish between any two states that are observationally equivalent from the point of view of some other agent. This observation coupled with the assumption that no two states are observationally equivalent from the point of view all agents imply

$$\mu_{it}(\theta) \longrightarrow \mathbb{1}_{\{\theta\}} \qquad \mathbb{P}\text{-a.s.},$$

completing the proof. □

2.4 Concluding Remarks

This chapter of the thesis studies a model of social learning where agents incorporate their private signals and the beliefs of their neighbors in a fully Bayesian manner. We showed that as long as the realized social network is strongly connected and agents do not face a global identification problem, they asymptotically almost surely learn the realized state. This result is true even if no agent can identify the realized state on her own and in spite of the agents' incomplete information about the realized social network and the sources of others' information.

The Bayesian learning framework mainly serves as a benchmark. The need for Bayesian agents to reason about the source and quality of the information obtained by their neighbors without having complete knowledge of the network structure significantly complicates the required calculations for Bayesian updating of beliefs, well beyond agents' regular computational capabilities. Gale and Kariv (2003) illustrate the complications that can arise due to repeated Bayesian deductions in a simple network. Also, as DeMarzo, Vayanos, and Zwiebel (2003) point out, in order to disentangle old information from new, a Bayesian agent needs to recall the information she received from her neighbors in the previous communication rounds, and therefore, "[w]ith multiple communication rounds, such calculations would become quite laborious, even

[5] Note that by definition, $\mu_{it} \in \mathscr{H}_{it}$, implying that $\mathscr{U}_{it} \subseteq \mathscr{H}_{it}$.

if the agent knew the entire social network." The necessary information and the computational burden of these calculations are simply prohibitive for adopting Bayesian learning, even in relatively simple networks.[6]

Motivated by these observations on the complexity of Bayesian updating, in the next chapter, we study a tractable model of social learning that asymptotically agrees with the Bayesian benchmark and perform explicit comparative analysis on the effect of the signal and network structures on the rate of learning.

[6]An exception, as shown by Mossel and Tamuz (2010), is the case in which agents' signal structures, their prior beliefs, and the social network are common knowledge and all signals and priors are normally distributed.

Chapter 3

Information Heterogeneity and the Speed of Social Learning

This chapter examines how social interactions in the presence of dispersed information determine the long run dynamics of the beliefs. To this end, we utilize a simple model of opinion formation introduced by Tahbaz-Salehi (2009) and Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012), which in turn is a variant of the seminal model of DeGroot (1974). Agents are endowed with a sequence of private signals that are potentially informative about an underlying state of the world. We capture the dispersion and heterogeneity of information by assuming that agents' private observations are drawn from different distributions and that they may be only partially informative about the state. We also assume that individuals cannot observe the beliefs held by all members of the society. Rather, they only have access to the opinions of the agents in their social clique. The key feature of the model is that each agent linearly combines her private observations with the opinions of her neighbors. Agents first incorporate their private signals in a fully Bayesian manner as an intermediate step. They then combine this Bayesian posterior with the beliefs of their neighbors in a naïve way by updating their beliefs to a convex combination of the beliefs of their neighbors and their Bayesian posteriors.

We argue that the evolution of the individuals' beliefs asymptotically coincides with that of the Bayesian benchmark introduced in Chapter 2. More specifically, given a strongly connected social network and an identifiable state of the world, each agent eventually learns the true underlying state as if she were aware of the observations of all agents and updated her beliefs according to Bayes' rule. Thus, with a constant flow of new information, the key condition for social learning is that individuals take their personal observations into account in a Bayesian way. Repeated communications over the social network ensures that the idiosyncratic differences eventually disappear and learning is obtained.

We also characterize the speed of social learning as a function of the primitives of the environment, namely, the structure of the social network and agents' information structures. We show that the structural features of the network which are relevant for the long run dynamics of the beliefs are summarized via its *eigenvector centrality*, a recursively defined notion of importance of different agents in the network. We also show that the rate of learning depends on the information structure of each agent via the expected log-likelihood ratios of her signal distributions, a quantity known as *relative entropy*. More specifically, our characterization result establishes that conditional on the realization of a given state, the rate at which agents rule out another state as a possibility is equal to the convex combination of the relative entropies, where the weights are given by the agents' eigenvector centralities.

Our characterization of the rate of learning enables us to analyze how the dispersion of in-

This chapter is based partly on Jadbabaie, Molavi, and Tahbaz-Salehi (2013).

formation throughout the network affects the long run dynamics of the beliefs. We show that if the vectors of pairwise relative entropies of all agents are comparable with respect to the product order—what we refer to as *uniform informativeness* ordering of signal structures—then, assigning signals of the highest quality to the more central agents increases the speed of social learning. Put differently, the positive assortative matching of signal qualities and eigenvector centralities leads to the fastest convergence of the beliefs to the truth. The intuition behind this result is simple: given that more central individuals receive more effective attention from other agents in the social network, assigning the information structure that is uniformly more informative to the more central agents guarantees that information is diffused faster.

This result is in line with the empirical observations of Banerjee, Chandrasekhar, Duflo, and Jackson (2012), who study how variations in the set of individuals exogenously informed about a microfinance program determine the eventual participation rate. By analyzing the social network and participation data in 43 rural villages in southern India, they find that the long run rate of participation is higher when the "injection points" have higher eigenvector centralities. Even though based on somewhat different premises, the theoretical implications of our model match Banerjee et al. (2012)'s empirical observations that (i) the network properties and locations of the exogenously informed individuals can substantially impact the diffusion of information; (ii) the extent of diffusion is significantly higher when the information injection points have higher eigenvector centralities; and more generally, (iii) beyond eigenvector centrality, other measures of network connectivity (such as average degree, average path length, clustering, etc.) do not play a substantial role in the long run dynamics of the beliefs. To the best of our knowledge, our work is the first to study the role of information heterogeneity and injection points from a theoretical point of view.

We also use our characterization of the rate of learning to analyze the long run dynamics of the beliefs under conditions that essentially correspond to the polar opposite of uniform informativeness ordering of signal structures. More specifically, we show that if every agent possesses some information that is crucial for learning—in the sense that other agents cannot learn the underlying state without her participation in the network—then the positive assortative matching of signal qualities and eigenvector centralities no longer maximizes the rate of learning. Rather, learning is obtained more rapidly if the least central agents receive signals of the highest quality. 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. On the other hand, a negative assortative matching of signal qualities and eigenvector centralities guarantees that these two events happen in parallel, leading to a faster convergence rate.

We also provide a comparative analysis of how the structural properties of the social network determine the speed at which agents' beliefs concentrate around the truth. We define a novel partial order over the set of social networks which captures the extent of asymmetry in the network structure. In particular, we define a network to be *more regular* than another if the eigenvector centrality of the former is majorized by that of the latter. According to this notion, eigenvector centralities in a more regular social network are more evenly distributed among different individuals.

We show that if the agents' signal structures are comparable in the uniform sense, then the rate of learning is decreasing in the social network's regularity. This is a consequence of the fact that, under positive assortative matching of signal qualities and eigenvector centralities, more dispersion in centralities guarantees that higher quality signals receive more effective attention, thus speeding up the learning process. In contrast, if all agents posses some information crucial for learning, then the speed of learning is higher in more regular networks, as in such a scenario an informational bottleneck effect plays a dominant role in determining the long run dynamics of the beliefs: learning is complete only when the information uniquely available to the most marginal agent is diffused throughout the society.

The juxtaposition of our results suggests that, in general, the exact role played by social interactions in the dynamics of the beliefs does not disentangle from the informational content of such communications. Rather, as our analysis highlights, the long run dynamics of the beliefs is sensitive to the specifics of how information is dispersed throughout the social network.

Our work belongs to the large body of works that study learning over social networks. One of the main strands of this literature focuses on simple, non-Bayesian rule-of-thumb updating processes. These works, for the most part, rely on the well-known opinion formation model of DeGroot (1974), according to which, agents update their beliefs as the average of their neighbors' opinions. DeMarzo, Vayanos, and Zwiebel (2003), Acemoglu, Ozdaglar, and ParandehGheibi (2010), and Golub and Jackson (2010, 2012a,b) are among the papers that use DeGroot's model to study opinion dynamics. Golub and Jackson (2010), for example, provide conditions on the structures of a growing sequence of social networks under which asymptotic opinions of all agents converge to the truth. The DeGroot model is generalized by Jadbabaie et al. (2012) to allow for constant arrival of new information over time. We utilize this latter model to study the speed of social learning in the presence of information heterogeneity.

The paper most closely related to the work presented in this chapter is the recent work of Golub and Jackson (2012b), who study the speed of convergence of the benchmark DeGroot learning model. However, the two works focus on different questions. Whereas Golub and Jackson study the role of homophily—the tendency of agents to associate disproportionately with those having similar traits—in belief dynamics, the focus of our analysis is to understand how the distribution of information over the social network affects the rate at which agents learn the truth. Furthermore, unlike the benchmark DeGroot model, we explicitly model agents' signal structures and analyze the role of information heterogeneity in their long run beliefs.

Finally, our work is related to the smaller and much more recent collection of empirical studies that focus on learning and the diffusion of information over social networks. As already mentioned, Banerjee et al. (2012) study the diffusion of a microfinance program in rural India and find that the eventual participation rate depends on the eigenvector centrality of the information injection points. Chandrasekhar, Larreguy, and Xandri (2011) conduct a unique lab experiment to test several models of learning over social networks. They observe that the evolution of choices made by the subjects are better explained by a variant of the DeGroot model rather than Bayesian learning models. Among other related empirical studies of learning and information aggregation in social networks are Choi, Gale, and Kariv (2005, 2012), Alatas, Banerjee, Chandrasekhar, Hanna, and Olken (2012), and Corazzini, Pavesi, Petrovich, and Stanca (2012).

3.1 Model

3.1.1 Agents and Observations

Consider a collection of n individuals, denoted by $N = \{1, 2 \ldots, n\}$, who are attempting to learn an underlying state of the world $\theta \in \Theta$. The underlying state is drawn at $t = 0$ according to probability distribution $\nu \in \Delta\Theta$ with full support over Θ, which we assume to be finite.

Even though the realized state of the world remains unobservable to the individuals, they make repeated noisy observations about θ over discrete time. In particular, 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 $\ell_i^\theta(\cdot) \in \Delta S$.[7] We assume that the signal space S is finite and that $\ell_i^\theta(\cdot)$ has full support over S for all i and all $\theta \in \Theta$. The realized signals are independent across individuals and over time. We refer to the collection of conditional probability distributions $\{\ell_i^\theta(\cdot)\}_{\theta \in \Theta}$ as the *signal structure* of agent i and for simplicity denote it by ℓ_i.

An individual's observations may not be informative regarding the underlying state. Rather, each agent may face an identification problem in the sense that she may not be able to distinguish between two states. We say two states are observationally equivalent from the point of view of an agent if the conditional distributions of her signals under the two states coincide. More specifically, the elements of the set $\Theta_i^\theta = \{\tilde{\theta} \in \Theta : \ell_i^{\tilde{\theta}} = \ell_i^\theta\}$ are observationally equivalent to state θ from the point of view of agent i. However, we impose the following global identifiability assumption on the agents' signal structures.

Assumption 3.1. There are no two states that are observationally equivalent from the point of view of all agents, that is, $\Theta_1^\theta \cap \cdots \cap \Theta_n^\theta = \{\theta\}$ for all $\theta \in \Theta$.

The above assumption guarantees that even though each agent may face an identification problem in isolation, there is no global identification problem: agents' observations are informative enough to allow them to collectively learn the underlying state.

Let $\omega_t = (\omega_{1t}, \ldots, \omega_{nt})$ denote the signal profile realized at time t and denote the set of infinite signal profile sequences by $\Omega = \{\omega : \omega = (\omega_1, \omega_2, \ldots)\}$. Conditional on the realization of state θ, sample paths $\omega \in \Omega$ are realized according to the probability distribution

$$\mathbb{P}^\theta = (\ell_1^\theta \times \cdots \times \ell_n^\theta)^{\mathbb{N}},$$

where $\mathbb{N}$ denotes the set of natural numbers. We use $\mathbb{P}$ to denote the probability distribution over $\Theta \times \Omega$ defined as $\mathbb{P}(\theta, \cdot) = \nu(\theta)\mathbb{P}^\theta(\cdot)$. The expectation operators with respect to probability distributions $\mathbb{P}^\theta$ and $\mathbb{P}$ are denoted by $\mathbb{E}^\theta$ and $\mathbb{E}$, respectively. Finally, we define $\mathscr{F}_{it}$ as the σ-algebra generated by the past history of agent i's observations up to time period t, and let $\mathscr{F}_t$ be the smallest σ-algebra containing all $\mathscr{F}_{it}$ for $i \in N$.

3.1.2 Learning Rule

At every time period, in addition to her private observation, each agent also has access to the beliefs of a subset of other agents, which we refer to as her neighbors. Agents apply a simple learning rule similar to the learning model of DeGroot (1974) to incorporate their private signals and the views of their neighbors. In particular, at every time period, each agent updates her

[7] All our results generalize to the case that agents' signal spaces are nonidentical.

belief as a convex combination of (i) the Bayesian posterior belief conditioned on her private signal; and (ii) the opinions of her neighbors. More precisely, if $\mu_{it}(\cdot) \in \Delta\Theta$ denotes the belief of agent i at time period t, then

$$\mu_{it+1} = a_{ii}\,\mathrm{BU}(\mu_{it};\omega_{it+1}) + \sum_{j\neq i} a_{ij}\mu_{jt}, \tag{3.1}$$

where $\mathrm{BU}(\mu_{it};\omega_{it+1})$ is the Bayesian update of agent i's opinion at time period t following the observation of the private signal ω_{it+1}, and $\{a_{ij}\}$ is a collection of nonnegative constants such that $a_{ii} > 0$ for all i and $a_{ij} > 0$ if and only if $j \neq i$ is a neighbor of i. Thus, the value of a_{ij} captures the weight that agent i assigns to the belief of agent j. For simplicity of exposition, we assume that $a_{ii} = \alpha$ for all i. By construction, individuals cannot directly incorporate the views of agents with whom they are not neighbors. Finally, note that for μ_{it+1} to be a well-defined probability distribution over Θ, the weights that each agent i assigns to her Bayesian posterior belief and the beliefs of her neighbors must add up to one, that is, $\sum_{j=1}^{n} a_{ij} = 1$.

3.1.3 Social Network

The extent of social interactions can be summarized by the matrix $A = [a_{ij}]$, which we refer to as the *social interaction matrix*. Equivalently, one can capture the social interactions between the agents by a directed, weighted graph on n vertices. Each vertex of this graph, which we refer to as the *social network*, corresponds to an agent and a directed edge (j,i) with weight $a_{ij} > 0$ is present from vertex j to vertex i if agent j is a neighbor of agent i. Thus, the social interaction matrix A is the (weighted) adjacency matrix of the underlying social network. Given this equivalence, we use the two concepts interchangeably.

A directed path in the social network from vertex i to vertex j is a sequence of vertices starting with i and ending with j such that each vertex is a neighbor of the next vertex in the sequence. We say the social network is strongly connected if there exists a directed path from each vertex to any other.

Assumption 3.2. The underlying social network is strongly connected.

The above assumption simply guarantees that information can flow from any agent in the network to any other. Expressed in terms of the social interaction matrix, Assumption 3.2 is equivalent to assuming that A is irreducible.[8]

3.2 Non-Bayesian Social Learning

As already mentioned, learning rule (3.1) is based on the learning model of DeGroot (1974) in which agents update their beliefs as a convex combination of their neighbors' beliefs at the previous time period. The feature that distinguishes our model from the benchmark DeGroot learning model is the constant arrival of new information over time. Whereas in the DeGroot model each agent has only a single observation, the individuals in our model receive information in small bits over time. In fact, if no individual receives any informative signals beyond time period

[8] Matrix A is *reducible*, if there exists a permutation matrix P such that $P'AP$ is block upper triangular. Otherwise, A is said to be *irreducible*. For more, see Berman and Plemmons (1979).

$t = 1$, learning rule (3.1) reduces to the DeGroot update. In this sense, our model is a natural generalization of the DeGroot learning models to the case that agents make repeated observations over time. This feature of the model is key for our results as it enables us to introduce information heterogeneity in a simple manner, and hence, study the role that different agents' signal structures play in the extent of information aggregation.

The model exhibits a number of other desirable features which make it suitable for the study of information aggregation over social networks. First, the evolution of the individuals' beliefs asymptotically coincides with those of Bayesian agents. Note that as in the benchmark DeGroot model, agents do not adjust the updating rule to account for the network structure and the differences in the precision of information that other agents may have learned over time. Nevertheless, the next result shows that the dispersed information in the social network is successfully aggregated.

Proposition 3.1. *Suppose that Assumptions 3.1 and 3.2 hold. Then,*

$$\mu_{it}(\cdot) - \mathbf{1}_\theta(\cdot) \longrightarrow 0 \qquad \mathbb{P}\text{-a.s.},$$

for all i, where $\mathbf{1}_\theta(\cdot)$ is the degenerate probability measure that assigns a unit mass on the underlying state θ.

Thus, the resulting learning process asymptotically coincides with Bayesian learning despite the fact that agents use a DeGroot-style update to incorporate the views of their neighbors. More specifically, each agent eventually learns the true underlying state of the world as if she were aware of the observations of all agents and updated her beliefs according to Bayes' rule.

Another feature of the model is its simplicity, which enables us to analytically characterize the asymptotic behavior of the agents' beliefs. In particular, unlike the Bayesian learning models—in which each agent needs to form and update beliefs about the observations made by all other agents while only observing the beliefs (or actions) of her neighbors—our model is tractable. Yet, as we show in the following sections, the model is rich enough to capture the nontrivial interplay between the informativeness of agents' private signal structures and the structure of the social network on the one hand and the extent and speed of learning on the other.

Finally, we remark that recent empirical evidence suggests that, in certain contexts, DeGroot-like learning models do a good job in explaining individuals' learning processes. More specifically, in a series of experiments conducted in rural India, Chandrasekhar et al. (2011) test several models of learning over social networks. They follow the evolution of choices made by subjects who repeatedly observe the actions of their neighbors. Their findings suggest that the observed paths of learning are better explained by a variant of the DeGroot model rather than Bayesian learning models.

3.3 Rate of Social Learning

In this section, we characterize the rate at which agents learn the underlying state as a function of the structure of the social network and the agents' signal structures. Before presenting the results, we define a few key concepts.

3.3.1 Relative Entropy

Given their heterogeneous signal structures, different individuals may have access to different information about the underlying state of the world. For example, some agents may receive more informative signals conditional on the realization of a specific state. Furthermore, the signal structure of each given individual may not be equally informative about all states, in the sense that the collection of her private signals may provide her with more information about a given state than another. To measure the extent of such heterogeneity, we borrow the concept of relative entropy, first introduced by Kullback and Leibler (1951), from the information theory literature.

Definition 3.1. Given two discrete probability distributions p and q with identical supports, the *relative entropy* of q with respect to p is

$$D(p\|q) = \sum_j p_j \log \frac{p_j}{q_j}$$

where p_j and q_j are the probabilities of the realization of the j-th outcome.

One can verify that $D(p\|q) \geq 0$ for all pairs of distributions p and q and that $D(p\|q) = 0$ if and only if $p = q$. In this sense, relative entropy is a nonsymmetric measure of the discrepancy between the two probability distributions. Alternatively, the relative entropy of q with respect to p is the expected value of the log-likelihood ratio test when p and q correspond to the null and alternative hypotheses distributions, respectively.[9]

The information content of agent i's signal structure can be measured in terms of the relative entropies of her marginal signal distributions. For any pair of states $\theta, \hat{\theta} \in \Theta$, let

$$h_i(\theta, \hat{\theta}) = D(\ell_i^{\theta} \| \ell_i^{\hat{\theta}}).$$

Thus, $h_i(\theta, \hat{\theta})$ is a measure of the expected information (per observation) in agent i's signal structure 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 the agent to rule out $\hat{\theta}$ with an arbitrarily large confidence. In particular, the number of observations required to reach a given prespecified 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 generated by her signal structure. On the other hand, if $h_i(\theta, \hat{\theta}) = 0$, then agent i would not be able to distinguish between the states based on her private signals alone, no matter how many observations she makes. In view of the above discussion, we define the following novel partial order over the set of signal structures:

Definition 3.2. Signal structure $\{\ell_i^{\theta}(\cdot)\}_{\theta \in \Theta}$ is *uniformly more informative* than $\{\ell_i'^{\theta}(\cdot)\}_{\theta \in \Theta}$, denoted by $\ell_i \succeq \ell_i'$, if the corresponding relative entropies satisfy

$$h_i(\theta, \hat{\theta}) \geq h_i'(\theta, \hat{\theta}) \qquad \text{for all} \quad \theta, \hat{\theta} \in \Theta.$$

[9] In information theory, the value of $D(p\|q)$ is interpreted as the expected number of extra bits required to code the information if the coder mistakenly assumes that the random variable is generated according to distribution q when in fact the true underlying distribution is p. For more on relative entropy and related concepts in information theory, see Cover and Thomas (1991).

In other words, if $\ell_i \succeq \ell'_i$, then ℓ_i is more discriminating between any pairs of states than ℓ'_i, and as a result, signals generated according to the former provide more information to agent i than the ones generated according to the latter. This notion of informativeness is weaker than Blackwell (1953)'s well-known criterion, 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.[10] Hence, if ℓ_i is more informative than ℓ'_i in the sense of Blackwell, then $\ell_i \succeq \ell'_i$, but not vice versa.[11] Finally, we remark that similar to Blackwell's ordering, uniform informativeness is a partial order on the signal structures, as not all signal structures are comparable in the sense of uniform informativeness.

3.3.2 Eigenvector Centrality

In addition to the individuals' signal structures, the detailed structure of the social network also plays a key role in the extent and speed of information aggregation. The following notion is a measure of the importance of different agents in the sense of information flow.

Definition 3.3. Given the matrix of social interactions A, the *eigenvector centrality* is a nonnegative vector v such that for all i,

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

and $\|v\|_1 = 1$. The i-th element of v is the eigenvector centrality of agent i.

The eigenvector centrality of agent i is thus a measure of her importance defined, recursively, as a function of the importance of the agents who are connected to her: an agent is more central if other more central agents put a large weight on her opinion. The Perron-Frobenius theorem guarantees that if the underlying social network is strongly connected, then the eigenvector centrality is a well-defined notion and is uniquely determined. Furthermore, $v_i > 0$ for all i.[12]

3.3.3 Speed of Learning

In the remainder of this section, we provide an analytical characterization of the speed of learning over the social network. As implied by Proposition 3.1, minimal connectivity and identifiability conditions are sufficient to ensure that agents learn the realized state *asymptotically*. However, after making any finite number of observations, agents remain uncertain about the underlying state of the world. The extent of this uncertainty at a given time period t can be measured via

$$e_t = \frac{1}{2} \sum_{i=1}^{n} \left\| \mu_{it}(\cdot) - \mathbf{1}_{\theta}(\cdot) \right\|_1. \tag{3.2}$$

[10] There exists an extensive literature in decision theory and information economics that provides foundations for different information orderings. For decision theoretic foundations of the concept of relative entropy see Sandroni (2000), Blume and Easley (2006), Lehrer and Smorodinsky (2000), and Cabrales, Gossner, and Serrano (2012).

[11] More discussions on the relationship between the two notions can be found in Appendix A.

[12] A detailed review of eigenvector and other notions of centrality is provided by Jackson (2008). For other applications of network centrality in economics, see Ballester, Calvó-Armengol, and Zenou (2006), Calvó-Armengol, de Martí, and Prat (2011), Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012) and Elliott and Golub (2013).

The above expression is the total variation distance between agents' beliefs at time t and the probability distribution that assigns a unit mass on the realized state of the world.[13] Thus, learning the underlying state implies that e_t converges to zero as $t \to \infty$. We define the *rate of learning* as

$$\lambda = \liminf_{t\to\infty} \frac{1}{t} \left| \log e_t \right|. \tag{3.3}$$

The above quantity is inversely proportional to the number of time periods it takes for the agents' uncertainty about the underlying state of the world to fall below some given threshold. In this sense, a higher value of λ implies that agents reach any given level of certainty about the state within a shorter time interval. The next proposition characterizes the rate of learning in terms of the primitives of the environment.

Proposition 3.2. *Suppose that Assumptions 3.1 and 3.2 are satisfied. Then, the following statements hold in a set of $\mathbb{P}$-probability one.*

(a) The rate of learning λ is finite and nonzero.

(b) Given any collection of signal structures, $\lambda \le r$, where

$$r = \alpha \min_{\theta\in\Theta} \min_{\hat\theta\neq\theta} \sum_{i=1}^{n} v_i h_i(\theta,\hat\theta) \tag{3.4}$$

and v_i is the eigenvector centrality of agent i.

(c) Furthermore, $\lambda = r + o\left(\max_{i,\theta,\hat\theta} \left\|\log \ell_i^{\theta}(\cdot) - \log \ell_i^{\hat\theta}(\cdot)\right\|\right)$.[14]

Part (a) of the above proposition establishes that agents learn the underlying state of the world (asymptotically) exponentially fast. In particular, the fact that $\lambda \in (0,\infty)$ means that, for large enough values of t, uncertainty e_t is proportional to $\exp(-\lambda t)$. The significance of Proposition 3.2, 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 that information is distributed among different agents. In particular, part (b) provides an upper bound on the rate of learning in terms of the relative entropies of agents' signal structures and their eigenvector centralities. Part (c) then shows that this upper bound is an arbitrarily good approximation to the rate of learning when the ratio $\ell_i^{\hat\theta}(s_i)/\ell_i^{\theta}(s_i)$ is close enough to 1 for all pairs of states, all signals s_i, and all agents. That is, the upper bound in (3.4) is arbitrarily tight when the information endowment of agents are small in the sense that no single private signal is very informative about the underlying state. In view of this result, throughout the rest of the current chapter, we use r as a proxy for the rate of learning.

Expression (3.4) for the rate of learning has an intuitive interpretation. Recall that relative entropy $h_i(\theta,\hat\theta)$ is the expected rate at which agent i 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

[13] The choice of the total variation distance between the probability distributions (or equivalently, the l_1-norm) in (3.2) is not essential. Given the equivalence of norms in finite dimensional spaces, replacing the l_1-norm with some other vector norm would simply lead to a rescaling of e_t by a positive constant, which does not affect the rate of learning.

[14] Given two real-valued functions $f(\cdot)$ and $g(\cdot)$ for which $\lim_{x\to 0} f(x) = \lim_{x\to 0} g(x) = 0$, the expression $f(x) = o(g(x))$ means that $\lim_{x\to 0} f(x)/g(x) = 0$.

in the informativeness of the agents' signals (in the uniform sense) cannot lead to a slower rate of learning. In fact, we have the following straightforward corollary to Proposition 3.2.

Corollary 3.1. *Suppose that $\ell_i \succeq \ell'_i$ for all agents i. Then, for any given social network, the rate of learning with signal structures $(\ell_1, \ldots, \ell_n)$ is no smaller than with signal structures $(\ell'_1, \ldots, \ell'_n)$.*

In addition to the signal structures, the rate of learning also depends on the structure of the social network. In particular, the relative entropy between distributions $\ell_i^\theta(\cdot)$ and $\ell_i^{\hat{\theta}}(\cdot)$ is weighed by agent i's eigenvector centrality, which measures the effective attention she receives from other agents in the social network. This characterization implies that in the presence of dispersed information, the process of learning exhibits a *network bottleneck effect*: the long run dynamics of the beliefs is less sensitive to changes in the information of agents located at the periphery of the network. On a broader level, this observation highlights the fact that even when the total amount of information available throughout the network is kept constant, the way this information is dispersed among different individuals may play a key role in determining the rate of social learning.

Another key observation is that learning is complete only if agents can rule out all incorrect states. More specifically, conditional on the realization of $\theta \in \Theta$, the speed of learning depends on the rate at which agents rule out the state $\hat{\theta} \neq \theta$ that is closest to θ in terms of relative entropy. Furthermore, the realization of the state itself affects the rate, as some states are easier to learn than others. Thus, as (3.4) suggests, the rate of learning is determined by minimizing the weighted sum of relative entropies over both the realized state θ and 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 (*ex ante*) rate of learning is determined by the pair of states $(\theta, \hat{\theta})$ that are hardest to distinguish from one another.

We end this discussion with a few remarks. First, note that, by definition, the rate of learning λ defined in (3.3) characterizes the agents' uncertainty about the underlying state asymptotically and does not capture the short term, transient dynamics of the beliefs. Thus, even though the structural properties of the social network other than its eigenvector centrality do not appear in the expression for the rate of learning, they play a role in how beliefs evolve in the short term. Finally, we remark that the rate of learning λ is in fact (the absolute value of) the top Lyapunov exponent of the dynamical system that describes the evolution of the agents' beliefs. Given that the top Lyapunov exponent of a dynamical system characterizes the rate of separation of infinitesimally close trajectories, it is not surprising that this quantity coincides with the rate of learning λ.

3.4 Information Allocation and Learning

Proposition 3.2 characterizes the long run dynamics of the beliefs in terms of the structural properties of the social network and the informativeness of each agent's observations. In this section, we study how the interplay between these two components may lead to nontrivial implications for the rate of social learning.

In view of Corollary 3.1, the rate of learning increases as agents receive more informative signals. Therefore, in order to capture the effect of dispersed information on the rate of learning in a meaningful way, we normalize the informativeness of the signal structures by keeping the total amount of information at the network level fixed.

Definition 3.4. The collection of signal structures $(\ell'_1,\ldots,\ell'_n)$ is a *reallocation* of $(\ell_1,\ldots,\ell_n)$ if there exists a permutation $\sigma : N \to N$ such that $\ell'_i = \ell_{\sigma(i)}$ for all i.

Thus, if $(\ell'_1,\ldots,\ell'_n)$ is a reallocation of $(\ell_1,\ldots,\ell_n)$, then the total amount of information available throughout the society is identical under the two information structures, even though the information available to any given individual may be different.

3.4.1 Learning Under Uniform Informativeness Ordering

We start our analysis by focusing on environments in which different agents' signal structures are comparable in the sense of uniform informativeness; that is, the collection of signal structures $(\ell_1,\ldots,\ell_n)$ are such that for any pair of agents i and j, either $\ell_i \succeq \ell_j$ or $\ell_j \succeq \ell_i$. Recall that by Definition 3.2, a signal structure is uniformly more informative than another, if the former is more discriminating between any pair of states than the latter. Thus, when the agents' signal structures are comparable, then there is an ordering of the 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.

Even though the notion of uniform informativeness only provides a partial ordering over the set of all signal structures, there are many real world scenarios in which the quality of exogenous information available to different agents can be unambiguously ranked in a natural way. For instance, in various marketing or public health campaigns only a subset of agents are exogenously informed (say, about a new product or the benefits of deworming). The rest of the individuals, on the other hand, do not have access to any exogenous sources of information. Rather, they can only obtain information via their interactions with one another or the exogenously informed agents. In such scenarios, the assumption of comparability of signal structures in the sense of uniform informativeness is naturally satisfied.

Proposition 3.3. *Suppose that the collection of signal structures $(\ell_1,\ldots,\ell_n)$ are comparable in the sense of uniform informativeness. Furthermore, suppose that for any pair of agents, $\ell_i \succeq \ell_j$ if and only if $v_i \geq v_j$, where v is the eigenvector centrality corresponding to the social network. 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 nondecreasing in the informativeness of their signals. In this sense, the positive assortative matching of signals and eigenvector centralities maximizes the rate of learning. The intuition behind this result is that if an information structure is uniformly more informative than another, 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 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.

The above result is in line with the empirical observations of Banerjee et al. (2012) who study how participation in a microfinance program diffuses through social networks. By focusing on the participation data from 43 villages in South India, Banerjee et al. (2012) find that the long run rate of microfinance participation is higher when the "injection points"—that is, individuals who were exogenously informed about the program—have higher eigenvector centralities in the social network. Proposition 3.3 shows a similar result for our model: if the information available to the agents can be ordered in the sense of uniform informativeness (as is indeed the case if agents are either informed about the program or not), then the speed of learning is maximized when agents with the highest eigenvector centralities are chosen as injection points. Finally, we remark that, as in the observations made by Banerjee et al. (2012), our result suggests that except for eigenvector centrality, other measures of network connectivity (such as average degree, average path length, clustering, etc.) do not play a role in the long run dynamics of the beliefs.

3.4.2 Experts and Learning Bottlenecks

Proposition 3.3 shows that as long as all signal structures are pairwise comparable in the uniform sense, positive assortative matching of signal qualities and eigenvector centralities maximizes the rate of social learning. However, given that uniform informativeness is a partial order over the set of all signal structures, there are many scenarios in which the conditions of Proposition 3.3 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 in the sense of Definition 3.2. In this subsection, we study how in the presence of such "experts"—i.e., agents who are particularly well-informed about a subset of states but not necessarily about others—the allocation of signal structures over the social network determines the rate of learning. Before presenting our general result, however, it is constructive to focus on a specific example.

Example 3.1. Consider a social network consisting of n agents and suppose that the set of states and observations are $\Theta = \{\theta_0, \theta_1, \dots, \theta_n\}$ and $S = \{\text{Head}, \text{Tail}\}$, respectively. Furthermore, suppose that agents' signal structures are given by

$$\ell_i^\theta(s) = \begin{cases} \pi_i & \text{if } \theta = \theta_i, s = \text{Head}, \\ \pi_i & \text{if } \theta \neq \theta_i, s = \text{Tail}, \\ 1 - \pi_i & \text{otherwise}, \end{cases} \tag{3.5}$$

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.

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

$$h_i(\theta, \hat{\theta}) = \begin{cases} H_i & \text{if } \theta = \theta_i, \hat{\theta} \neq \theta_i, \\ H_i & \text{if } \theta \neq \theta_i, \hat{\theta} = \theta_i, \\ 0 & \text{otherwise}, \end{cases} \tag{3.6}$$

where $H_i = (2\pi_i - 1)\left[\log \pi_i - \log(1 - \pi_i)\right]$ 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 $\ell_i^{\theta_i}(\cdot)$ and other distributions in i's signal structure.

Given (3.6), the rate of learning is equal to

$$r = \alpha \min_i v_i H_i,$$

where v is the eigenvector centrality. In view of the above expression, one can verify that among all possible allocations of signal structures to agents, the negative assortative matching of agents' expertise and eigenvector 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 3.1 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 eigenvector centralities leads to the fastest rate of learning among all possible allocations of signal structures. Put differently, it is best if the least central agents receive signals of the highest quality (π_i closer to 1).

The intuition behind this observation is as follows. Recall from the discussion following Proposition 3.2 that the process of learning exhibits two distinct bottlenecks, namely, the network and the identification bottleneck effects. 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 state pairs that are the hardest to distinguish from one another. 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 eigenvector 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 3.1, the speed of learning under a positive assortative matching is equal to $\alpha v_{\min} H_{\min}$; the smallest value $r = \alpha \min_i v_i H_i$ can obtain. On the contrary, the speed of learning is maximized under the negative assortative matching of signal qualities and eigenvector centralities, as such an allocation would guarantee that the two bottlenecks are as far away from one another as possible.

We remark that in the discussions following Proposition 3.3 and Example 3.1, we used terms such as "higher quality signals" and "better information" somewhat loosely. Whereas in the former case, such terms are used to refer to a signal structure that is more informative in the uniform sense, in the context of Example 3.1 similar terms are used to refer to the signal structure of an agent with a higher expertise about a state. Nevertheless, we emphasize that despite this apparent inconsistency, these two cases can indeed be unified in a consistent manner using a weaker notion of information ordering. In particular, we say a signal structure is *weakly more informative* than another, if it is uniformly more informative under some permutation of the states.[15] This definition immediately implies that in the context of Example 3.1, the signal structure of agent i is (weakly) more informative than agent j's if and only if $H_i \geq H_j$. Thus, in both cases, a signal of higher quality provides more information about the state and leads to a faster reduction of uncertainty.

We now proceed to show that the insights obtained from Example 3.1 remain valid in more

[15]Formally, $\ell_i \succeq_w \ell_j$ if there exists a permutation $\sigma : \Theta \to \Theta$ such that $\ell'_i \succeq \ell_j$ where $\ell'^{\theta}_i(\cdot) = \ell^{\sigma(\theta)}_i(\cdot)$ for all θ.

general settings. Given the (ordered) pair of states $(\theta,\hat{\theta})$, let

$$\gamma_i(\theta,\hat{\theta}) = \sup\{\beta : h_i(\theta,\hat{\theta}) \geq \beta h_j(\theta,\hat{\theta}) \quad \forall j \neq i\}$$

capture the extent to which the signal structure of agent i is more informative in distinguishing θ from $\hat{\theta}$ relative to other agents'. Thus, $\gamma_i(\theta,\hat{\theta}) \geq 1$ means that no other agent can rule out $\hat{\theta}$ when the underlying state is θ with fewer observations (in expectation) than agent i. Furthermore, let

$$E_i = \{(\theta,\hat{\theta}) : \theta \neq \hat{\theta} \text{ and } \gamma_i(\theta,\hat{\theta}) \geq 1\}$$

be the set of state pairs which agent i can tell apart better than any other agent.[16] Finally, we define the following notions of expertise:

Definition 3.5. Provided that $E_i \neq \emptyset$, the *relative* and *absolute expertise* of agent i are, respectively,

$$\begin{aligned}\gamma_i &= \min\{\gamma_i(\theta,\hat{\theta}) : (\theta,\hat{\theta}) \in E_i\},\\ \varepsilon_i &= \min\{h_i(\theta,\hat{\theta}) : (\theta,\hat{\theta}) \in E_i\}.\end{aligned}$$

We have the following result.

Proposition 3.4. *Suppose that $E_i \neq \emptyset$ for all i. Furthermore, suppose that $\varepsilon_i \leq \varepsilon_j$ if and only if $v_i \geq v_j$, where v is the eigenvector centrality corresponding to the social network. Then, no reallocation of signal structures increases the rate of learning by more than $\alpha(\max_i \varepsilon_i)/(\min_i \gamma_i)$.*

Thus, if for any agent i, there exists a pair of states $(\theta,\hat{\theta})$ for which she can accumulate evidence in favor of one versus the other at a higher rate than all other agents, then negative assortative matching of absolute levels of expertise and eigenvector centralities leads to rapid learning. In particular, no reallocation of signal structures can increase the rate of learning by more than $\alpha \max_i \varepsilon_i / \min_i \gamma_i$. Note that this constant is inversely proportional to $\min_i \gamma_i$. Therefore, *ceteris paribus*, increasing the relative expertise of all agents leads to a smaller upper bound. In fact, this upper bound can be arbitrarily small if agents' signal structures are such that their relative expertise are large enough.

3.5 Network Regularity and Learning

Our analysis thus far was focused on how the rate of learning changes with the reallocation of information among agents, while keeping the structure of the social network fixed. In this section, we provide a comparative analysis of how the structural properties of the social network determine the speed at which agents' beliefs concentrate around the truth.

We start by defining a partial order over the set of social networks. As before, let v denote the eigenvector centrality corresponding to the social interaction matrix (equivalently, social network) A.

[16] Given that relative entropy is a nonsymmetric measure of discrepancy between two distributions, it may be the case that $(\hat{\theta},\theta) \notin E_i$, even though $(\theta,\hat{\theta}) \in E_i$.

Definition 3.6. Social network A is *more regular* than A' if v' majorizes v, that is, if

$$\sum_{i=1}^{k} v_{[i]} \leq \sum_{i=1}^{k} v'_{[i]} \qquad \forall k \in \{1, \ldots, n\}, \tag{3.7}$$

where $x_{[i]}$ is the i-th largest element of vector x.[17]

Intuitively, a social network is more regular if the effective attention that different agents receive from the rest of the society is more evenly distributed. In particular, as (3.7) suggests, increasing the centralities of more central agents at the expense of more marginal ones would lead to a less regular social network. The following simple example illustrates how the notion of regularity defined above captures the extent of asymmetry in the structure of the social network.

Example 3.2. Consider the ring social network, depicted in Figure 3.1(a). Each agent i updates her belief as a function of her private observations and the opinion of a single other agent, namely agent $i-1$. Due to the full symmetry in the network structure, it is immediate that all agents have equal eigenvector centralities, that is, $v_i = 1/n$ for all i. Thus, no other social network is (strictly) more regular than the ring.[18] There are, however, other social networks that are as regular as the ring network. In particular, any social network for which the sum of the weights assigned to the opinion of each agent is equal across the society—that is, $\sum_{j\neq i} a_{ji}$ is equal for all i—is as regular as the ring network.[19]

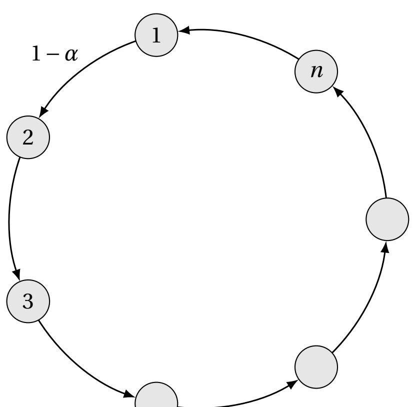

(a) The ring social network

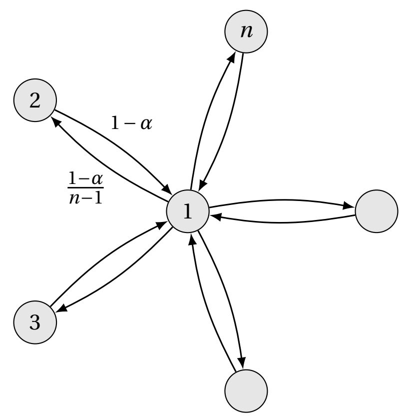

(b) The star social network

Figure 3.1. The ring and star social networks

At the opposite end of the spectrum lies the star social network, depicted in Figure 3.1(b). As the figure suggests, a single agent, namely agent 1, takes a a disproportionately more central

[17] Equivalently, A is more regular than A' if the probability distribution corresponding to probability mass function $v^{\downarrow}$ first-order stochastically dominates the probability distribution corresponding to $v'^{\downarrow}$, where $x^{\downarrow}$ denotes the vector with the same components of x sorted in the nonincreasing order.

[18] This is a consequence of the fact that $\sum_{i=1}^{k} v'_{[i]} \geq k/n$ for any stochastic vector v'.

[19] In the graph theory literature, such a network is simply referred to as a *regular graph*. Note that, by construction, $\sum_{j\neq i} a_{ji} = 1-\alpha$ in any regular network.

position in the network relative to all other agents. Such an asymmetry is also reflected in the agents' eigenvector centralities. In particular, $v_1 = 1/2$, whereas $v_i = 1/(2n-2)$ for $i \neq 1$. One can verify that for the eigenvector centrality v' corresponding to any other social network, $\sum_{i=1}^{k} v'_{[i]} \leq \sum_{i=1}^{k} v_{[i]}$ for all k. Thus, the highly asymmetric star network is indeed the least regular social network.

We end this discussion by contrasting our notion of regularity with an alternative notion of symmetry introduced by Acemoglu et al. (2012), who measure the symmetry of the network topology in terms of the standard deviation of the agents' centralities. Even though both notions capture the dispersion in the agents' centralities, the two are not identical. In particular, as we show in Appendix B, if a network is more regular than another, then it is also more symmetric as defined by Acemoglu et al. (2012). The converse, however, is not true as regularity only defines a partial order over the set of social networks. Finally, we remark that, under both notions, the ring and star networks correspond to the most and least symmetric network structures, respectively.

We now present the main result of this section. For a given social network and a collection of signal structures, let r^* denote the fastest rate of learning that can be obtained via reallocation of signals.

Proposition 3.5. *Suppose that A is more regular than A' and that $\ell_i = \ell'_i$ for all i. Also suppose that the collection of signal structures $(\ell_1, \ldots, \ell_n)$ are comparable in the sense of uniform informativeness. Then, $r^* \leq r'^*$.*

Thus, if the agents' signal structures can be ordered in the sense of uniform informativeness, the rate of learning under the optimal allocation of signals is decreasing in the regularity of the social network. The intuition behind the above result is simple: with uniform informativeness ordering of signal structures and under positive assortative matching of agents' centralities and signal qualities, a higher level of dispersion in centralities guarantees that higher quality signals receive more effective attention, thus speeding up the learning process. Proposition 3.5 also implies that the ring and star social networks correspond to the smallest and largest rates of learning, respectively. The next example shows that in large societies, the difference in the performance of the two social networks can be arbitrarily large.

Example 2 (continued). Consider the ring and star social networks depicted in Figures 3.1(a) and 3.1(b), respectively. Also suppose that the private observations are such that $m \leq n$ agents have access to identical signal structures that are informative about the underlying state of the world, whereas the remaining $n-m$ individuals do not observe any informative signals. In other words, for any given i, either $h_i(\theta,\hat{\theta}) = \bar{h}(\theta,\hat{\theta}) > 0$ for all distinct pairs $(\theta,\hat{\theta})$, or $h_i(\theta,\hat{\theta}) = 0$ for all $(\theta,\hat{\theta})$. Clearly, the signal structures can be ordered in the uniform sense. Proposition 3.5 thus implies that the (maximum) speed of learning in the star network is higher than in the ring social network. Furthermore, simple algebraic derivations imply that as $n \to \infty$,

$$r^*_{\text{ring}} / r^*_{\text{star}} = O(m/n). \tag{3.8}$$

Thus, if the number of informed agents m does not grow at the same rate as the network size n, learning in the ring and star networks occurs at diverging rates.

Proposition 3.5 establishes that if the signal structures are pairwise comparable in the uniform sense, then rate of learning is decreasing in the social network's regularity. Furthermore, the above example shows that in large networks, the role played by the structural properties of the networks can be significant. In view of the discussion in Section 3.4.2, however, one would expect that similar results would no longer hold in the presence of expert agents. The following simple example shows that this is indeed the case.

Example 3.3. Consider the collection of signal structures given in (3.5) and assume that $\pi_i = \pi$ for all i. Thus, each agent i is the expert in learning state θ_i. As we showed in Example 3.1, the rate of learning is simply equal to $r = \alpha v_{\min} H$. This immediately implies that if social network A is more regular than A', then $r \geq r'$, regardless of the allocation of the signal structures.

Thus, in stark contrast to the case in which signal structures are comparable in the uniform sense, more regularity in the social network implies a larger rate of convergence. This is due to the fact that when all agents are experts, the information bottleneck effect plays a dominant role in determining the long run dynamics of the beliefs: learning is complete only when the information uniquely available to the most marginal agent is diffused throughout the society. It is thus the centrality of the least central agent that determines the rate of information diffusion, and as a result, a more regular structure guarantees a faster convergence.

Despite this observation, our next result establishes that in the presence of experts, the rate of learning in large societies does not vary significantly as a function of the network structure. Contrasting it with (3.8), the result also highlights yet another way in which the long run dynamics of the beliefs crucially depend on the way information is dispersed throughout the network.

Proposition 3.6. *Consider a sequence of information structures $(\ell_{1,n}, \ldots, \ell_{n,n})$ parametrized by the number of agents n. Also suppose that*

(a) $E_{i,n} \neq \emptyset$ for all i and all n.

(b) There exists $c \geq 1$ such that $\max_i \varepsilon_{i,n} < c \min_i \varepsilon_{i,n}$ for all n.

(c) $\liminf_{n\to\infty} \gamma_{i,n}/n > 0$ for all i.

Then, for any two sequences of social networks A_n and A'_n and any allocation of signal structures,

$$0 < \liminf_{n\to\infty} \frac{r_n}{r'_n} \leq \limsup_{n\to\infty} \frac{r_n}{r'_n} < \infty.$$

Thus, in the presence of expert agents, the rate of learning is essentially of the same order for all network structures and signal allocations. Such a result is due to the fact that when all agents are experts, the information content in all agents' private signals are equally important for learning. Hence, regardless of the structure of the social network or the allocation of the signals, the rate of learning depends on the eigenvector centrality of the most marginal agent, which is always of order $1/n$. We remark that the assumptions required for the above proposition to hold are fairly weak. Assumption (a) simply means that all agents are experts, whereas (b) requires that the absolute expertise of the agents do not diverge from one another. Finally, (c) is a technical assumption guaranteeing that the relative expertise of the agents are large enough.

3.6 Proofs

Proof of Proposition 3.1

Before presenting the proofs of the results in the paper, we state and prove three lemmas which will later be used in the proof of the proposition.

Let $\mu_{it}^{\theta}(\cdot)$ be the restriction of $\mu_{it}(\cdot)$ to the event that the underlying state is θ. We prove the proposition by showing that, for all θ, all $\hat{\theta} \neq \theta$, and all i and with $\mathbb{P}^{\theta}$-probability one, $\mu_{it}^{\theta}(\hat{\theta}) \to 0$. The belief update rule (3.1) can thus be rewritten as

$$\mu_{it+1}^{\theta}(\hat{\theta}) = \alpha \frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})} \mu_{it}^{\theta}(\hat{\theta}) + \sum_{j \neq i} a_{ij} \mu_{jt}^{\theta}(\hat{\theta}), \tag{3.9}$$

where $m_{it}^{\theta}(s) = \sum_{\tilde{\theta} \in \Theta} \mu_{it}^{\theta}(\tilde{\theta}) \ell_i^{\tilde{\theta}}(s)$, called the time t one-step-ahead forecast of agent i, is the probability that agent i assigns at time t to the event that she observes signal s in the next time period. More generally, with some abuse of notation, we define the k-step-ahead forecast of agent i at time t as

$$m_{it}^{\theta}(s_1, \dots, s_k) = \sum_{\tilde{\theta} \in \Theta} \mu_{it}^{\theta}(\tilde{\theta}) \ell_i^{\tilde{\theta}}(s_1) \ell_i^{\tilde{\theta}}(s_2) \cdots \ell_i^{\tilde{\theta}}(s_k). \tag{3.10}$$

The following lemma, which is proved by Tahbaz-Salehi (2009), shows that agents' one-step-ahead forecasts are asymptotically almost surely correct.

Lemma 3.1. *Suppose that Assumption 3.2 holds. Then,*

$$m_{it}^{\theta}(\cdot) - \ell_i^{\theta}(\cdot) \longrightarrow 0 \qquad \mathbb{P}^{\theta}\text{-a.s.},$$

for all i.

We next present and prove a simple lemma which is later used in the proof of the proposition.

Lemma 3.2. *Suppose that Assumption 3.2 holds. Then, for all $\hat{\theta} \in \Theta$,*

$$\mathbb{E}^{\theta}\left[\mu_{t+1}^{\theta}(\hat{\theta}) \middle| \mathscr{F}_t\right] - A\mu_t^{\theta}(\hat{\theta}) \longrightarrow 0 \qquad \mathbb{P}^{\theta}\text{-a.s.},$$

where $\mu_t^{\theta}(\hat{\theta})$ is the n-dimensional column vector with the i-th element equal to $\mu_{it}^{\theta}(\hat{\theta})$.

Proof. Equation (3.9) can be written in the vector form as

$$\mu_{t+1}^{\theta}(\hat{\theta}) = A\mu_t^{\theta}(\hat{\theta}) + \alpha \operatorname{diag}\left(\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})} - 1\right)_{i \in N} \mu_t^{\theta}(\hat{\theta}). \tag{3.11}$$

Taking conditional expectations from both sides of the above equation implies

$$\mathbb{E}^{\theta}\left[\mu_{t+1}^{\theta}(\hat{\theta}) \middle| \mathscr{F}_t\right] - A\mu_t^{\theta}(\hat{\theta}) = \alpha \operatorname{diag}\left(\mathbb{E}^{\theta}\left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})} - 1 \middle| \mathscr{F}_t\right]\right)_{i \in N} \mu_t^{\theta}(\hat{\theta}).$$

On the other hand, we have

$$\mathbb{E}^{\theta}\left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})}\middle|\mathscr{F}_t\right]=\sum_{s\in S}\ell_i^{\hat{\theta}}(s)\frac{\ell_i^{\theta}(s)}{m_{it}^{\theta}(s)}\longrightarrow\sum_{s\in S}\ell_i^{\hat{\theta}}(s)\qquad\mathbb{P}^{\theta}\text{-a.s.,}$$

where the convergence is a consequence of Lemma 3.1. The fact that $\ell_i^{\hat{\theta}}$ is a probability measure on S implies $\sum_{s\in S}\ell_i^{\hat{\theta}}(s)=1$, completing the proof. □

The next lemma establishes that not only agents make accurate predictions about their private observations in the next period, but also make correct predictions about any finite time horizon in the future.

Lemma 3.3. *Suppose that Assumption 3.2 holds. Then,*

$$m_{it}^{\theta}(s_1,\ldots,s_k)-\ell_i^{\theta}(s_1)\ell_i^{\theta}(s_2)\cdots\ell_i^{\theta}(s_k)\longrightarrow 0\qquad\mathbb{P}^{\theta}\text{-a.s.,}$$

for all natural numbers k, all $(s_1,s_2,\ldots,s_k)\in S^k$, and all i.

Proof. We prove this lemma by induction. In Lemma 3.1, we established that the claim is true for $k=1$. For the rest of the proof, we assume that the claim is true for $k-1$ and show that $m_{it}^{\theta}(s_1,\ldots,s_k)$ converges to $\ell_i^{\theta}(s_1)\cdots\ell_i^{\theta}(s_k)$ for any arbitrary sequence of signals $(s_1,\ldots,s_k)\in S^k$.[20]

First, note that Lemma 3.2 and equation (3.9) imply that for all $\hat{\theta}\in\Theta$,

$$\mathbb{E}^{\theta}\left[\mu_{it+1}^{\theta}(\hat{\theta})\middle|\mathscr{F}_t\right]-\mu_{it+1}^{\theta}(\hat{\theta})+\alpha\left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})}-1\right]\mu_{it}^{\theta}(\hat{\theta})\longrightarrow 0\qquad\mathbb{P}^{\theta}\text{-a.s.}$$

Multiplying both sides by $\prod_{\tau=2}^{k}\ell_i^{\hat{\theta}}(s_\tau)$ for an arbitrary signal sequence $(s_2,\ldots,s_k)\in S^{k-1}$ and summing up over all $\hat{\theta}\in\Theta$ lead to

$$\sum_{\hat{\theta}\in\Theta}\left(\prod_{\tau=2}^{k}\ell_i^{\hat{\theta}}(s_\tau)\right)\left(\mathbb{E}^{\theta}\left[\mu_{it+1}^{\theta}(\hat{\theta})\middle|\mathscr{F}_t\right]-\mu_{it+1}^{\theta}(\hat{\theta})\right)+\alpha\sum_{\hat{\theta}\in\Theta}\left(\prod_{\tau=2}^{k}\ell_i^{\hat{\theta}}(s_\tau)\right)\left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})}-1\right]\mu_{it}^{\theta}(\hat{\theta})\longrightarrow 0$$

with $\mathbb{P}^{\theta}$-probability one. On the other hand,

$$\sum_{\hat{\theta}\in\Theta}\left(\prod_{\tau=2}^{k}\ell_i^{\hat{\theta}}(s_\tau)\right)\left(\mathbb{E}^{\theta}\left[\mu_{it+1}^{\theta}(\hat{\theta})\middle|\mathscr{F}_t\right]-\mu_{it+1}^{\theta}(\hat{\theta})\right)=\mathbb{E}^{\theta}\left[m_{it+1}^{\theta}(s_2,\ldots,s_k)\middle|\mathscr{F}_t\right]-m_{it+1}^{\theta}(s_2,\ldots,s_k),$$

where we have used the definition of $(k-1)$-step-ahead forecasts of agent i. The induction hypothesis and the dominated convergence theorem for conditional expectations imply that the right-hand side of the above equation converges to zero with $\mathbb{P}^{\theta}$-probability one. Therefore,

$$\sum_{\hat{\theta}\in\Theta}\left(\prod_{\tau=2}^{k}\ell_i^{\hat{\theta}}(s_\tau)\right)\left[\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})}-1\right]\mu_{it}^{\theta}(\hat{\theta})\longrightarrow 0\qquad\mathbb{P}^{\theta}\text{-a.s.,}$$

[20] Recall that we use $s\in S$ to denote a generic element of the signal space, whereas ω_{it} denotes the random variable corresponding to i's observation at time t.

for any arbitrary sequence of signals $(s_2,\ldots,s_k)\in S^{k-1}$, which is equivalent to

$$\frac{1}{m^\theta_{it}(\omega_{it+1})}m^\theta_{it}(\omega_{it+1},s_2,\ldots,s_k)-m^\theta_{it}(s_2,\ldots,s_k)\longrightarrow 0 \qquad \mathbb{P}^\theta\text{-a.s.}$$

Thus, once again by the induction hypothesis,

$$m^\theta_{it}(\omega_{it+1},s_2,\ldots,s_k)-m^\theta_{it}(\omega_{it+1})\prod_{\tau=2}^{k}\ell^\theta_i(s_\tau)\longrightarrow 0 \qquad \mathbb{P}^\theta\text{-a.s.}$$

The dominated convergence theorem for conditional expectations implies

$$\mathbb{E}^\theta\left[\left|m^\theta_{it}(\omega_{it+1},s_2,\ldots,s_k)-m^\theta_{it}(\omega_{it+1})\prod_{\tau=2}^{k}\ell^\theta_i(s_\tau)\right|\Big|\mathscr{F}_t\right]\longrightarrow 0 \qquad \mathbb{P}^\theta\text{-a.s.}$$

Rewriting the conditional expectation operator as a sum over all possible realizations of ω_{it+1} leads to

$$\sum_{\tilde{s}\in S}\ell^\theta_i(\tilde{s})\left|m^\theta_{it}(\tilde{s},s_2,\ldots,s_k)-m^\theta_{it}(\tilde{s})\prod_{\tau=2}^{k}\ell^\theta_i(s_\tau)\right|\longrightarrow 0$$

$\mathbb{P}^\theta$-almost surely, and therefore, guaranteeing

$$m^\theta_{it}(s_1,s_2,\ldots,s_k)-m^\theta_{it}(s_1)\prod_{\tau=2}^{k}\ell^\theta_i(s_\tau)\longrightarrow 0 \qquad \mathbb{P}^\theta\text{-a.s.},$$

for all $s_1\in S$.[21] Finally, the fact that $m^\theta_{it}(s_1)\to\ell^\theta_i(s_1)$ with $\mathbb{P}^\theta$-probability one (Lemma 3.1) completes the proof. □

We next show that for any agent i, there exists a finite sequence of private signals that is more likely to realize under the state θ than any other state $\hat{\theta}$, unless $\hat{\theta}$ is observationally equivalent to θ from the point of view of agent i.

Lemma 3.4. *For any agent i, there exists a positive integer $\hat{k}^\theta_i$, a sequence of signals $(\hat{s}^\theta_{i1},\ldots,\hat{s}^\theta_{i\hat{k}_i})\in S^{\hat{k}^\theta_i}$, and constant $\delta^\theta_i\in(0,1)$ such that*

$$\prod_{\tau=1}^{\hat{k}^\theta_i}\frac{\ell^{\hat{\theta}}_i(\hat{s}^\theta_{i\tau})}{\ell^\theta_i(\hat{s}^\theta_{i\tau})}<\delta^\theta_i \qquad \forall\hat{\theta}\notin\Theta^\theta_i. \tag{3.12}$$

Proof. By definition, for any $\hat{\theta}\notin\Theta^\theta_i$, the probability measures $\ell^{\hat{\theta}}_i$ and ℓ^θ_i are distinct. Therefore, by the Kullback-Leibler inequality, there exists some constant $\epsilon^\theta_i>0$ such that

$$\sum_{s\in S}\ell^\theta_i(s)\log\frac{\ell^\theta_i(s)}{\ell^{\hat{\theta}}_i(s)}>\epsilon^\theta_i,$$

[21] Recall that, by assumption, $\ell^\theta_i(s)>0$ for all s, all θ, and all i.

for all $\hat{\theta} \not\in \Theta_i^\theta$, which then implies

$$\prod_{s\in S}\left[\frac{\ell_i^{\hat{\theta}}(s)}{\ell_i^\theta(s)}\right]^{\ell_i^\theta(s)} < \varepsilon_i^\theta,$$

for $\varepsilon_i^\theta = \exp(-\epsilon_i^\theta)$. On the other hand, given the fact that rational numbers are dense on the real line, there exist strictly positive rational numbers $\{q^\theta(s)\}_{s\in S}$—with $q^\theta(s)$ chosen sufficiently close to $\ell_i^\theta(s)$—satisfying $\sum_{s\in S} q^\theta(s) = 1$, such that

$$\prod_{s\in S}\left[\frac{\ell_i^{\hat{\theta}}(s)}{\ell_i^\theta(s)}\right]^{q^\theta(s)} < \varepsilon_i^\theta \qquad \forall \hat{\theta} \not\in \Theta_i^\theta. \tag{3.13}$$

Therefore, the above inequality can be rewritten as

$$\prod_{s\in S}\left[\frac{\ell_i^{\hat{\theta}}(s)}{\ell_i^\theta(s)}\right]^{k^\theta(s)} < (\varepsilon_i^\theta)^{\hat{k}_i^\theta} \qquad \forall \hat{\theta} \not\in \Theta_i^\theta,$$

for some positive integers $k^\theta(s)$ and $\hat{k}_i^\theta$, satisfying $\hat{k}_i^\theta = \sum_{s\in S} k^\theta(s)$. Picking the sequence of signals of length $\hat{k}_i^\theta$, $(\hat{s}_{i1}^\theta, \ldots, \hat{s}_{i\hat{k}_i^\theta}^\theta)$, such that s appears $k^\theta(s)$ many times in the sequence and setting $\delta_i = (\varepsilon_i^\theta)^{\hat{k}_i^\theta}$ proves the lemma. $\square$

The above lemma shows that the sequence of private signals in which any signal $s \in S$ appears with a frequency close enough to $\ell_i^\theta(s)$ is more likely under the state θ than any other state $\hat{\theta}$ which is distinguishable from θ. We now proceed to the proof of Proposition 3.1.

Proof of Proposition 3.1 First, we prove that with $\mathbb{P}^\theta$-probability one agent i assigns an asymptotic belief of zero on states that are not observationally equivalent to θ from her point of view.

Recall that according to Lemma 3.3, the k-step-ahead forecasts of agent i are eventually correct for all positive integers k, guaranteeing that $m_{it}^\theta(s_1, \ldots, s_k) \to \prod_{\tau=1}^k \ell_i^\theta(s_\tau)$ with $\mathbb{P}^\theta$-probability one for any sequence of signals $(s_1, \ldots, s_k)$. In particular, the claim is true for the integer $\hat{k}_i^\theta$ and the sequence of signals $(\hat{s}_{i1}^\theta, \ldots, \hat{s}_{i\hat{k}_i^\theta}^\theta)$ satisfying (3.12) in Lemma 3.4:

$$\sum_{\hat{\theta}\in\Theta} \mu_{it}^\theta(\hat{\theta}) \prod_{\tau=1}^{\hat{k}_i^\theta} \frac{\ell_i^{\hat{\theta}}(\hat{s}_{i\tau}^\theta)}{\ell_i^\theta(\hat{s}_{i\tau}^\theta)} \longrightarrow 1 \qquad \mathbb{P}^\theta\text{-a.s.}$$

Therefore,

$$\sum_{\hat{\theta}\not\in\Theta_i^\theta} \mu_{it}^\theta(\hat{\theta}) \prod_{\tau=1}^{\hat{k}_i^\theta} \frac{\ell_i^{\hat{\theta}}(\hat{s}_{i\tau}^\theta)}{\ell_i^\theta(\hat{s}_{i\tau}^\theta)} + \sum_{\hat{\theta}\in\Theta_i^\theta} \mu_{it}^\theta(\hat{\theta}) - 1 \longrightarrow 0 \qquad \mathbb{P}^\theta\text{-a.s.},$$

leading to

$$\sum_{\hat{\theta}\not\in\Theta_i^\theta} \mu_{it}^\theta(\hat{\theta}) \left(1 - \prod_{\tau=1}^{\hat{k}_i^\theta} \frac{\ell_i^{\hat{\theta}}(\hat{s}_{i\tau}^\theta)}{\ell_i^\theta(\hat{s}_{i\tau}^\theta)}\right) \longrightarrow 0 \qquad \mathbb{P}^\theta\text{-a.s.}$$

The fact that $\hat{k}_i^\theta$ and $(\hat{s}_{i1}^\theta, \ldots, \hat{s}_{i\hat{k}_i^\theta}^\theta)$ were chosen to satisfy (3.12) implies that

$$1 - \prod_{\tau=1}^{\hat{k}_i^\theta} \frac{\ell_i^{\hat{\theta}}(\hat{s}_{i\tau}^\theta)}{\ell_i^\theta(\hat{s}_{i\tau}^\theta)} > 1 - \delta_i^\theta > 0 \qquad \forall \hat{\theta} \not\in \Theta_i^\theta,$$

and as a consequence, it must be the case that $\mu_{it}^\theta(\hat{\theta}) \to 0$ as $t \to \infty$ for any $\hat{\theta} \not\in \Theta_i^\theta$. Therefore, with $\mathbb{P}^\theta$-probability one, agent i assigns an asymptotic belief of zero on any state $\hat{\theta}$ that is not observationally equivalent to θ from her point of view.

Now consider the belief update rule for agent i given by equation (3.9), evaluated at some state $\hat{\theta} \not\in \Theta_i^\theta$:

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

We have already shown that $\mu_{it}^\theta(\hat{\theta}) \to 0$, $\mathbb{P}^\theta$-almost surely. However, this is not possible unless $\sum_{j \neq i} a_{ij} \mu_{jt}^\theta(\hat{\theta})$ converges to zero as well, which implies that $\mu_{jt}^\theta(\hat{\theta}) \to 0$ with $\mathbb{P}^\theta$-probability one for all j such that $a_{ij} > 0$. Note that this happens even if $\hat{\theta}$ is observationally equivalent to θ from the point of view of agent j; that is, even if $\hat{\theta} \in \Theta_j^\theta$. As a result, all neighbors of agent i will assign an asymptotic belief of zero to parameter $\hat{\theta}$ regardless of their signal structure. We can extend the same argument to the neighbors of neighbors of agent i, and by induction—since the social network is strongly connected—to all agents in the network. Thus, with $\mathbb{P}^\theta$-probability one,

$$\mu_{it}^\theta(\hat{\theta}) \longrightarrow 0 \qquad \forall i \in N, \qquad \forall \hat{\theta} \not\in \Theta_1^\theta \cap \cdots \cap \Theta_n^\theta,$$

implying that all agents asymptotically $\mathbb{P}^\theta$-almost surely assign a belief of zero on states that are not observationally equivalent to θ from the point of view of all individuals in the society. Therefore, Assumption 3.1 implies that $\mu_{it}^\theta(\hat{\theta}) \to 0$ for all $\hat{\theta} \neq \theta$, with $\mathbb{P}^\theta$-probability one, guaranteeing complete learning by all agents. □

Proof of Proposition 3.2

We first provide a proof for statement (b) of the proposition. We then proceed to prove parts (a) and (c).

Proof of Part (b) Let $\mu_{it}^\theta(\cdot)$ be the restriction of $\mu_{it}(\cdot)$ to the event that the underlying state is θ. The belief update rule (3.1) can thus be rewritten as

$$\mu_{it+1}^\theta(\hat{\theta}) = \alpha \frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^\theta(\omega_{it+1})} \mu_{it}^\theta(\hat{\theta}) + \sum_{j \neq i} a_{ij} \mu_{jt}^\theta(\hat{\theta}),$$

where $m_{it}^\theta(s) = \sum_{\tilde{\theta} \in \Theta} \mu_{it}^\theta(\tilde{\theta}) \ell_i^{\tilde{\theta}}(s)$ is the probability that agent i assigns at time t to the event that she observes signal s in the next time period. Taking logarithms of both sides of the above equation and using Jensen's inequality imply

$$\log \mu_{it+1}^{\theta}(\hat{\theta}) \geq \sum_{j=1}^{n} a_{ij} \log \mu_{jt}^{\theta}(\hat{\theta}) + \alpha \log\left(\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})}\right).$$

Let $x_t^{\theta}(\hat{\theta}) = \sum_{i=1}^{n} v_i \log \mu_{it}^{\theta}(\hat{\theta})$. Multiplying both sides of the above inequality by v_i and summing over i lead to

$$x_{t+1}^{\theta}(\hat{\theta}) \geq x_t^{\theta}(\hat{\theta}) + q_t^{\theta}(\hat{\theta}),$$

where

$$q_t^{\theta}(\hat{\theta}) = \alpha \sum_{i=1}^{n} v_i \log\left(\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{m_{it}^{\theta}(\omega_{it+1})}\right),$$

and as a result,

$$\frac{1}{T}\left(x_T^{\theta}(\hat{\theta}) - x_0^{\theta}(\hat{\theta})\right) \geq \frac{1}{T}\sum_{t=0}^{T-1} q_t^{\theta}(\hat{\theta}). \tag{3.14}$$

By Proposition 3.1, agent i asymptotically learns the underlying state of the world θ, that is, $m_{it}^{\theta}(s) \to \ell_i^{\theta}(s)$ as $t \to \infty$ with $\mathbb{P}^{\theta}$-probability one for all s. Hence, $q_t^{\theta}(\hat{\theta}) - p_t^{\theta}(\hat{\theta}) \to 0$ with $\mathbb{P}^{\theta}$-probability one, where

$$p_t^{\theta}(\hat{\theta}) = \alpha \sum_{i=1}^{n} v_i \log\left(\frac{\ell_i^{\hat{\theta}}(\omega_{it+1})}{\ell_i^{\theta}(\omega_{it+1})}\right).$$

Taking the limit of both sides of inequality (3.14) as $T \to \infty$ implies

$$\limsup_{T\to\infty} \frac{1}{T}[x_T^{\theta}(\hat{\theta})] \geq \lim_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T-1}\left(q_t^{\theta}(\hat{\theta}) - p_t^{\theta}(\hat{\theta})\right) + \lim_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T-1} p_t^{\theta}(\hat{\theta}).$$

Given that $q_t^{\theta}(\hat{\theta}) - p_t^{\theta}(\hat{\theta})$ converges to zero on $\mathbb{P}^{\theta}$-almost all paths, the Cesàro means theorem implies that the first term on the right-hand side of the above inequality is equal to zero $\mathbb{P}^{\theta}$-almost surely.[22] Furthermore, by the strong law of large numbers, the second term on the right-hand side is equal to $\mathbb{E}^{\theta}[p_t^{\theta}(\hat{\theta})]$. Therefore,

$$\limsup_{t\to\infty} \frac{1}{t}\sum_{i=1}^{n} v_i \log \mu_{it}^{\theta}(\hat{\theta}) \geq \mathbb{E}^{\theta}[p_t^{\theta}(\hat{\theta})] = -\alpha \sum_{i=1}^{n} v_i h_i(\theta, \hat{\theta}), \tag{3.15}$$

with $\mathbb{P}^{\theta}$-probability one, where we are using the definition of $h_i(\theta, \hat{\theta})$.

Let e_t^{θ} denote the restriction of e_t to the event that the underlying state is θ. By definition,

$$e_t^{\theta} = \sum_{i=1}^{n} \sum_{\hat{\theta}\neq\theta} \mu_{it}^{\theta}(\hat{\theta}),$$

[22] For a statement and proof of the Cesáro means theorem see, for example, Hardy (1992, pp. 100–102).

and as a result

$$
\begin{aligned}
\log e_t^\theta &\geq \log\left(\max_{\hat\theta\neq\theta}\max_i \mu_{it}^\theta(\hat\theta)\right)\\
&= \max_{\hat\theta\neq\theta}\max_i \log\mu_{it}^\theta(\hat\theta)\\
&\geq \max_{\hat\theta\neq\theta}\sum_{i=1}^n v_i\log\mu_{it}^\theta(\hat\theta),
\end{aligned}
$$

where in the last inequality we are using the fact that $\sum_{i=1}^n v_i = 1$. Thus, for all $\hat\theta\neq\theta$,

$$
\begin{aligned}
\limsup_{t\to\infty}\frac{1}{t}\log e_t^\theta &\geq \limsup_{t\to\infty}\frac{1}{t}\sum_{i=1}^n v_i\log\mu_{it}^\theta(\hat\theta)\\
&\geq -\alpha\sum_{i=1}^n v_i h_i(\theta,\hat\theta),
\end{aligned}
$$

$\mathbb{P}^\theta$-almost surely, where the second inequality is a consequence of (3.15). Consequently, with $\mathbb{P}^\theta$-probability one,

$$
\liminf_{t\to\infty}\frac{1}{t}\left|\log e_t^\theta\right| \leq \alpha\min_{\hat\theta\neq\theta}\sum_{i=1}^n v_i h_i(\theta,\hat\theta),
$$

and hence,

$$
\lambda = \liminf_{t\to\infty}\frac{1}{t}\left|\log e_t\right| \leq \alpha\min_\theta\min_{\hat\theta\neq\theta}\sum_{i=1}^n v_i h_i(\theta,\hat\theta) = r, \tag{3.16}
$$

$\mathbb{P}$-almost surely. □

Proof of Part (a) The fact that $\lambda\leq r$ immediately implies that λ is finite with $\mathbb{P}$-probability one. Thus, we only need to show that the rate of learning λ is strictly positive. Define $\phi_t^\theta : \mathbb{R}^{n(|\Theta|-1)} \to \mathbb{R}^{n(|\Theta|-1)}$ as the mapping that maps $(\mu_{it-1}^\theta(\hat\theta))_{i,\hat\theta\neq\theta}$ to $(\mu_{it}^\theta(\hat\theta))_{i,\hat\theta\neq\theta}$. Also let

$$
M_t^\theta = \operatorname{diag}\left(M_t^\theta(\hat\theta)\right)_{\hat\theta\neq\theta}, \tag{3.17}
$$

where

$$
M_t^\theta(\hat\theta) = A + \alpha\operatorname{diag}\left(\frac{\ell_i^{\hat\theta}(\omega_{it})}{\ell_i^\theta(\omega_{it})} - 1\right)_{i\in N}.
$$

[23]

It is easy to verify that the block diagonal matrix M_t^θ is the Jacobian of ϕ_t^θ evaluated at the origin. Thus, the linear dynamical system generated by $\{M_t^\theta\}_{t\in\mathbb{N}}$ is a linearization of the nonlinear dynamical system $\{\phi_t^\theta\}_{t\in\mathbb{N}}$ describing the evolution of the beliefs conditional on the realization of state θ. Finally, we let ζ^θ be the top Lyapunov exponent (TLE) corresponding to the sequence of matrices M_t^θ, that is,

$$
\zeta^\theta = \lim_{t\to\infty}\frac{1}{t}\log\left\|M_t^\theta M_{t-1}^\theta\cdots M_1^\theta\right\|. \tag{3.18}
$$

[23] Note that both ϕ_t^θ and M_t^θ are defined conditional on the realization of a given path $\omega\in\Omega$. However, we suppress this dependency for notational simplicity.

Furstenberg and Kesten (1960) show that ζ^θ exists, is independent of the choice of the matrix norm, and is equal to a deterministic constant with $\mathbb{P}^\theta$-probability one. We have the following lemma.

Lemma 3.5. *$\zeta^\theta < 0$ with $\mathbb{P}^\theta$-probability one.*

Proof. By Theorems 1 and 2 of Furstenberg and Kesten (1960), the constant ζ^θ defined in (3.18) is equal to

$$\zeta^\theta = \lim_{t\to\infty} \frac{1}{t}\mathbb{E}^\theta\left[\log\left\|M_t^\theta M_{t-1}^\theta \cdots M_1^\theta\right\|_1\right]$$

and therefore,

$$\begin{aligned}
\zeta^\theta &\leq \lim_{t\to\infty} \frac{1}{t}\log\mathbb{E}^\theta\left\|M_t^\theta M_{t-1}^\theta \cdots M_1^\theta\right\|_1 \\
&\leq \lim_{t\to\infty} \frac{1}{t}\log\mathbb{E}^\theta\left[\mathbf{1}' M_t^\theta M_{t-1}^\theta \cdots M_1^\theta \mathbf{1}\right] \\
&= \lim_{t\to\infty} \frac{1}{t}\log\left(\mathbf{1}'[\mathbb{E}^\theta M_1^\theta]^t\mathbf{1}\right),
\end{aligned}$$

where the first inequality is a consequence of Jensen's inequality and the second is due to the fact that $\|X\|_1 \leq \mathbf{1}'X\mathbf{1}$ for any nonnegative matrix X. Furthermore, given that $\mathbb{E}^\theta M_1^\theta = A$, the Perron-Frobenius theorem implies that $[\mathbb{E}^\theta M_1^\theta]^t \to \mathbf{1}v'$ as $t\to\infty$ where v is the left eigenvector of A. Hence,

$$\zeta^\theta \leq \lim_{t\to\infty} \frac{1}{t}\log n = 0.$$

On the other hand, Assumption 3.1 guarantees that $\rho(M_t^\theta) \neq 1$ with $\mathbb{P}^\theta$-positive probability, where $\rho(X)$ denotes the spectral radius of matrix X. Thus, Theorem 2 of Kesten and Spitzer (1984) implies that $\zeta^\theta \neq 0$, which completes the proof. □

The above lemma thus shows that the linear dynamical system generated by $\{M_t^\theta\}_{t\in\mathbb{N}}$ is $\mathbb{P}^\theta$-almost surely exponentially stable with exponent ζ^θ. Next, we have the following lemma:

Lemma 3.6. *For all i and all $\hat\theta \neq \theta$,*

$$\limsup_{t\to\infty} \mu_{it}^\theta(\hat\theta) \leq \zeta^\theta,$$

with $\mathbb{P}^\theta$-probability one.

The above lemma, the proof of which is provided at the end of the chapter, follows from the fact that, under the condition of Lyapunov regularity, exponential stability of the linearization of a nonlinear dynamical system guarantees that the original nonlinear system is also exponentially stable with the same exponent.[24] Lemma 3.6 implies that for $\mathbb{P}^\theta$-almost all ω,

$$\limsup_{t\to\infty} \frac{1}{t}\log e_t^\theta \leq \zeta^\theta,$$

and therefore,

$$\lambda = \liminf_{t\to\infty} \frac{1}{t}\left|\log e_t\right| \geq \min_\theta |\zeta^\theta| > 0 \tag{3.19}$$

[24]For more on the Lyapunov exponents, Lyapunov regularity, and stability of random dynamical systems, see Appendix C.

with $\mathbb{P}$-probability one, which completes the proof of part (a). $\square$

Proof of Part (c) Since M_t^θ is block diagonal,

$$\|M_t^\theta M_{t-1}^\theta \cdots M_1^\theta\| = \max_{\hat{\theta}\neq\theta} \|M_t^\theta(\hat{\theta}) M_{t-1}^\theta(\hat{\theta}) \cdots M_1^\theta(\hat{\theta})\|,$$

and hence,

$$\zeta^\theta = \max_{\hat{\theta}\neq\theta} \zeta^\theta(\hat{\theta}),$$

where $\zeta^\theta(\hat{\theta})$ is the TLE of $M_t^\theta(\hat{\theta})$. Also recall that (3.16) and (3.19) provide upper and lower bounds on the rate of learning λ, respectively. Combining the above immediately implies that

$$\min_\theta \min_{\hat{\theta}\neq\theta} |\zeta^\theta(\hat{\theta})| \le \lambda \le r, \tag{3.20}$$

$\mathbb{P}$-almost surely. In the remainder of the proof, we provide a lower bound on $|\zeta^\theta(\hat{\theta})|$ and show that this lower bound is arbitrarily close to r if $\ell_i^{\hat{\theta}}(s_i)/\ell_i^\theta(s_i)$ is close enough to one for all agents and all signals. More specifically, we use a result by Gharavi and Anantharam (2005) who provide an upper bound for the TLE of Markovian products of nonnegative matrices, expressed as the maximum of a nonlinear concave function over a finite-dimensional convex polytope of probability distributions. In order to apply their result, we first need to introduce some new notation.

Let $s = (s_1, s_2, \ldots, s_n)$ where $s_i \in S$. Define $\mathscr{S} = \{1, 2, \ldots, |S|^n\}$ and let $f : \mathscr{S} \to S^n$ be an enumeration of the set of signal profiles S^n.[25] For all $k \in \mathscr{S}$, we use $f_i(k)$ to denote the observation of agent i when signal profile $f(k)$ is realized; i.e., $f(k) = (f_1(k), \ldots, f_n(k))$. For all $k \in \mathscr{S}$, let $p^k = \mathbb{P}^\theta(f(k))$ and let Q^k be the realization of $M_t^\theta(\hat{\theta})$ given $\omega_t = f(k)$; that is, Q^k is the $n \times n$ matrix defined as

$$Q^k = A + \alpha \,\mathrm{diag}\left(\frac{\ell_i^{\hat{\theta}}(f_i(k))}{\ell_i^\theta(f_i(k))} - 1\right)_{i\in N}.$$

Let $H(p)$ be the entropy of p defined as

$$H(p) = -\sum_{k\in\mathscr{S}} p^k \log p^k.$$

Let $\mathscr{M}$ be the set of all probability measures on $(N \times \mathscr{S}) \times (N \times \mathscr{S})$. Finally, with a slight abuse of notation, for any $\eta \in \mathscr{M}$, let $H(\eta)$ be the entropy of η defined as

$$H(\eta) = -\sum_{\substack{i,j\in N\\ k,l\in\mathscr{S}}} \eta_{ij}^{kl} \log \frac{\eta_{ij}^{kl}}{\eta_{i*}^{k*}},$$

where

$$\eta_{i*}^{k*} = \sum_{\substack{j\in N\\ l\in\mathscr{S}}} \eta_{ij}^{kl}.$$ [26]

Gharavi and Anantharam (2005) show that the solution to the optimization problem below

[25] An enumeration of a finite set X is a bijective mapping from $\{1, 2, ..., |X|\}$ to X.

[26] More generally, an asterisk $*$ in lieu of an index indicates summation over that index.

is an upper bound on the TLE of the set of i.i.d. matrices $\{M_t^\theta(\hat\theta)\}_{t\in\mathbb{N}}$ when $M_t^\theta(\hat\theta)\in\{Q^k\}_{k\in\mathscr{S}}$ and Q^k is realized with probability p^k. That is, $\zeta^\theta(\hat\theta)\le\xi^\theta(\hat\theta)$, where

$$\begin{aligned}
\xi^\theta(\hat\theta) = \max_{\eta\in\mathscr{M}} \quad & H(\eta)+F(\eta)-H(p) && && (3.21)\\
\text{subject to} \quad & \eta_{**}^{kl}=p^kp^l && \forall\, k,l\in\mathscr{S}, && (3.22)\\
& \eta_{i*}^{k*}=\eta_{*i}^{*k} && \forall\, i\in N \quad \forall\, k\in\mathscr{S}, && (3.23)\\
& \eta_{ij}^{kl}=0 && \forall\, i,j\in N \quad \forall\, k,l\in\mathscr{S} \quad \text{s.t.} \quad Q_{ji}^k=0, && (3.24)
\end{aligned}$$

where Q_{ij}^k is the (i,j) element of Q^k and $F:\mathscr{M}\to\mathbb{R}$ is defined as

$$F(\eta)=\sum_{\substack{i,j\in N\\ k,l\in\mathscr{S}}}\eta_{ij}^{kl}\log Q_{ji}^k,$$

with the usual convention that $0\log 0=0$. Note that the diagonal elements of Q^k are positive, whereas its off-diagonal elements are equal to the corresponding elements of A. Hence, $Q_{ji}^k=0$ if and only if $a_{ji}=0$. Consequently, when solving for the optimal solution, we can let $\eta_{ij}^{kl}=0$ whenever $a_{ji}=0$ and drop constraint (3.24) altogether. Given that maximizing (3.21) subject to (3.22) and (3.23) is a strictly convex problem, the first-order conditions characterize the unique optimal solution. Using Lagrange multipliers ρ^{kl} and ν_i^k for the first and second set of constraints, the first-order condition with respect to η_{ij}^{kl} is

$$\log Q_{ji}^k-\log\eta_{ij}^{kl}+\log\eta_{i*}^{k*}+\rho^{kl}+\nu_i^k-\nu_j^l=0, \tag{3.25}$$

for $i,j\in N$ and $k,l\in\mathscr{S}$ such that $Q_{ji}^k\neq0$. Thus, any (η,ρ,ν) that solves (3.22), (3.23), and (3.25) simultaneously corresponds to an optimal solution.

In the case that $\ell_i^{\hat\theta}(\cdot)=\ell_i^\theta(\cdot)$ for all i, it is easy to verify that

$$\begin{aligned}
\hat\eta_{ij}^{kl}&=v_ja_{ji}p^kp^l\\
\hat\rho^{kl}&=\log p^k\\
\hat\nu_i^k&=-\log(v_ip^k)
\end{aligned}$$

satisfy optimality conditions (3.22), (3.23) and (3.25). Substituting $\hat\eta$ in (3.21) then implies that $\xi^\theta(\hat\theta)=0$.

If, on the other hand, the ratio $\ell_i^{\hat\theta}(s_i)/\ell_i^\theta(s_i)$ is close to one for all agents i and all signals, we can approximate $\xi^\theta(\hat\theta)$ by its first-order Taylor expansion around the point $\ell_i^{\hat\theta}(\cdot)=\ell_i^\theta(\cdot)$ as

$$\xi^\theta(\hat\theta)=\sum_{i=1}^n\sum_{s_i\in S}\left[\left.\frac{\partial\xi^\theta(\hat\theta)}{\partial\log\ell_i^{\hat\theta}(s_i)}\right|_{\ell_i^{\hat\theta}(\cdot)=\ell_i^\theta(\cdot)}\cdot\left(\log\ell_i^{\hat\theta}(s_i)-\log\ell_i^\theta(s_i)\right)\right]+o\left(\max_i\left\|\log\ell_i^{\hat\theta}(\cdot)-\log\ell_i^\theta(\cdot)\right\|\right), \tag{3.26}$$

where we are using the fact that $\xi^\theta(\hat\theta)$ evaluated at $\ell_i^{\hat\theta}(\cdot)=\ell_i^\theta(\cdot)$ is equal to zero.[27] Since none of the constraints of the optimization problem depends on $\ell_i^{\hat\theta}(\cdot)$, by the envelope theorem, the

[27]Note that (3.26) is the first-order Taylor expansion of $\xi^\theta(\hat\theta)$ with respect to $\log\ell_i^{\hat\theta}(s_i)$ rather than $\ell_i^{\hat\theta}(s_i)$.

derivative of $\xi^\theta(\hat\theta)$ with respect to $\log \ell_i^{\hat\theta}(\cdot)$ is simply equal to the partial derivative of $H(\eta)+F(\eta)-H(p)$ with respect to $\log \ell_i^{\hat\theta}(\cdot)$ holding η fixed and then evaluating the result at the optimal solution $\hat\eta$. The terms $H(p)$ and $H(\eta)$ do not explicitly depend on $\ell_i^{\hat\theta}(\cdot)$. On the other hand, $F(\eta)$ is given by

$$F(\eta) = \sum_{\substack{i,j\in N\\ k,l\in\mathscr{S}}} \eta_{ij}^{kl} \log a_{ji} + \sum_{\substack{j\in N\\ k,l\in\mathscr{S}}} \eta_{jj}^{kl} \log \frac{\ell_j^{\hat\theta}(f_j(k))}{\ell_j^{\theta}(f_j(k))}.$$

Therefore,

$$\begin{aligned}
\frac{\partial \xi^\theta(\hat\theta)}{\partial \log \ell_i^{\hat\theta}(s_i)} &= \sum_{\substack{j\in N\\ k,l\in\mathscr{S}}} \hat\eta_{jj}^{kl}\, \partial \left[\log \frac{\ell_j^{\hat\theta}(f_j(k))}{\ell_j^{\theta}(f_j(k))}\right] \Big/ \partial \log \ell_i^{\hat\theta}(s_i) \\
&= \sum_{k,l\in\mathscr{S}} \alpha v_i p^k p^l \frac{\partial}{\partial \log \ell_i^{\hat\theta}(s_i)} \left[\log \frac{\ell_i^{\hat\theta}(f_i(k))}{\ell_i^{\theta}(f_i(k))}\right] \\
&= \alpha v_i \ell_i^\theta(s_i),
\end{aligned}$$

where all the derivatives are evaluated at $\ell_i^{\hat\theta}(s_i) = \ell_i^{\theta}(s_i)$. Substituting the above in (3.26) thus implies

$$\begin{aligned}
\xi^\theta(\hat\theta) &= \alpha \sum_{i=1}^n \sum_{s_i\in S} v_i \ell_i^\theta(s_i) \left(\log \ell_i^{\hat\theta}(s_i) - \log \ell_i^{\theta}(s_i)\right) + o\left(\max_i \left\|\log \ell_i^{\hat\theta}(\cdot) - \log \ell_i^{\theta}(\cdot)\right\|\right) \\
&= -\alpha \sum_{i=1}^n v_i h_i(\theta,\hat\theta) + o\left(\max_i \left\|\log \ell_i^{\hat\theta}(\cdot) - \log \ell_i^{\theta}(\cdot)\right\|\right).
\end{aligned}$$

Therefore, by (3.20) and the fact that $\xi^\theta(\hat\theta) \ge \zeta^\theta(\hat\theta)$,

$$r + o\left(\max_{i,\theta,\hat\theta} \left\|\log \ell_i^{\hat\theta}(\cdot) - \log \ell_i^{\theta}(\cdot)\right\|\right) \le \lambda \le r,$$

which completes the proof. □

Proof of Proposition 3.3

Without loss of generality assume that agents are indexed such that $v_1 \ge v_2 \ge \cdots \ge v_n$. Therefore, by assumption, $h_1(\theta,\hat\theta) \ge h_2(\theta,\hat\theta) \ge \cdots \ge h_n(\theta,\hat\theta)$ 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 Hardy-Littlewood rearrangement inequality[28]

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

[28]The Hardy-Littlewood rearrangement inequality states that if $x_1 \ge x_2 \ge \cdots \ge x_n$ and $y_1 \ge y_2 \ge \cdots \ge y_n$, then, for any permutation $(y_1',\ldots,y_n')$ of $(y_1,\ldots,y_n)$, we have $\sum_{i=1}^n x_i y_i \ge \sum_{i=1}^n x_i y_i'$. For a reference, see, for example, Theorem 368 of Hardy, Littlewood, and Pólya (1952).

for any given pair of states θ and $\hat{\theta}$. Therefore,

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

for all $\theta \in \Theta$, and as a result

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

completing the proof. □

Proof of Proposition 3.4

We first state and prove a simple lemma.

Lemma 3.7. *Suppose that $E_i \neq \emptyset$. Then,*

$$\alpha v_i \varepsilon_i \leq r_i \leq \alpha(v_i + 1/\gamma_i)\varepsilon_i,$$

where $r_i = \alpha \min_{(\theta,\hat{\theta})\in E_i} \sum_{j=1}^{n} v_j h_j(\theta,\hat{\theta})$.

Proof. For any pair of states $(\theta,\hat{\theta})$, we have $\sum_{j=1}^{n} v_j h_j(\theta,\hat{\theta}) \geq v_i h_i(\theta,\hat{\theta})$. Thus,

$$r_i \geq \alpha v_i \min_{(\theta,\hat{\theta})\in E_i} h_i(\theta,\hat{\theta}) = \alpha v_i \varepsilon_i,$$

which establishes the lower bound. On the other hand, for all $(\theta,\hat{\theta}) \in E_i$,

$$\begin{aligned}\sum_{j=1}^{n} v_j h_j(\theta,\hat{\theta}) &\leq v_i h_i(\theta,\hat{\theta}) + \sum_{j\neq i} v_j h_i(\theta,\hat{\theta})/\gamma_i(\theta,\hat{\theta}) \\ &\leq (v_i + 1/\gamma_i) h_i(\theta,\hat{\theta}).\end{aligned}$$

Therefore,

$$r_i \leq \alpha(v_i + 1/\gamma_i) \min_{(\theta,\hat{\theta})\in E_i} h_i(\theta,\hat{\theta}) = \alpha(v_i + 1/\gamma_i)\varepsilon_i,$$

establishing the upper bound. □

Proof of Proposition 3.4 Throughout the proof, without loss of generality, we assume that agents are indexed such that $v_1 \geq \cdots \geq v_n$. The assumption of the proposition thus guarantees that $\varepsilon_1 \leq \cdots \leq \varepsilon_n$.

Since $E_i \neq \emptyset$,

$$r_i = \alpha \min_{(\theta,\hat{\theta})\in E_i} \sum_{j=1}^{n} v_j h_j(\theta,\hat{\theta})$$

is well-defined for all i. Furthermore, for any (ordered) pair of states $(\theta,\hat{\theta})$ such that $\theta \neq \hat{\theta}$, there exists some agent who is no worse than others in distinguishing between the two; or equivalently,

$\cup_i E_i = \{(\theta,\hat\theta) : \theta \neq \hat\theta\}$. Therefore,

$$r = \min_i r_i. \tag{3.27}$$

Now let $\gamma = \min_i \gamma_i$ and suppose that there exists a reallocation of signal structures with the corresponding rate of learning r' such that $r + (\alpha/\gamma)\max_i \varepsilon_i < r'$, which in view of (3.27) can be rewritten as

$$\min_i r_i + \frac{\alpha}{\gamma}\max_i \varepsilon_i < \min_i r'_i.$$

As a result,

$$r_k + \frac{\alpha}{\gamma}\max_i \varepsilon_i < r'_i \qquad \text{for all } i,$$

where $k \in \arg\min_i r_i$. Thus, by Lemma 3.7,

$$\alpha v_k \varepsilon_k + \frac{\alpha}{\gamma}\max_i \varepsilon_i < \alpha v_i \varepsilon'_i + \frac{\alpha}{\gamma}\varepsilon'_i \qquad \text{for all } i.$$

Note that the above inequality holds for all i only if $v_k\varepsilon_k < v_i\varepsilon'_i$ for all i. In particular,

$$v_k\varepsilon_k < v_i\varepsilon'_i \qquad \text{for all } i \in \{k,\dots,n\}.$$

Furthermore, recall that by assumption, $v_i \le v_k$ for $i \ge k$. Therefore, $\varepsilon_k < \varepsilon'_i$ for all $i \in \{k,\dots,n\}$. This, however, leads to a contradiction. In particular, given that $(\varepsilon'_1,\dots,\varepsilon'_n)$ is a permutation of $(\varepsilon_1,\dots,\varepsilon_n)$ there are at most $n-k$ indices j such that $\varepsilon_k < \varepsilon'_j$. Thus, no reallocation of signals can increase the rate of learning by more than $(\alpha/\gamma)\max_i \varepsilon_i$. □

Proof of Proposition 3.5

By assumption, the collection of signal structures $(\ell_1,\dots,\ell_n)$ are comparable in the sense of uniform informativeness. Without loss of generality, assume that signal structures are indexed such that $\ell_1 \succeq \cdots \succeq \ell_n$, that is, $h_1(\theta,\hat\theta) \ge \cdots \ge h_n(\theta,\hat\theta)$ for all $\theta,\hat\theta \in \Theta$. Thus, by Proposition 3.3, the optimal rates of learning in social networks A and A' are given by

$$r^* = \alpha \min_\theta \min_{\hat\theta\neq\theta} \sum_{i=1}^n v_{[i]} h_i(\theta,\hat\theta)$$
$$r'^* = \alpha \min_\theta \min_{\hat\theta\neq\theta} \sum_{i=1}^n v'_{[i]} h_i(\theta,\hat\theta),$$

respectively. On the other hand, given the assumption that A is more regular than A' and by Lemma B.1, we have,

$$\sum_{i=1}^n v_{[i]} h_i(\theta,\hat\theta) \le \sum_{i=1}^n v'_{[i]} h_i(\theta,\hat\theta),$$

which immediately implies $r^* \le r'^*$. □

Proof of Proposition 3.6

To simplify notation, we suppress the dependence on the size of the network n whenever there is no risk of confusion. Since $E_i \neq \emptyset$, by Lemma 3.7,

$$\alpha v_i \varepsilon_i \leq r_i \leq \alpha (v_i + 1/\gamma_i)\varepsilon_i.$$

Thus, by (3.27),

$$\begin{aligned} r &\leq \alpha \left(\min_i v_i + \max_i 1/\gamma_i \right) \max_i \varepsilon_i \\ &\leq \alpha \left(1/n + \max_i 1/\gamma_i \right) \max_i \varepsilon_i, \end{aligned}$$

where the second inequality is due to the fact that $\min_i v_i \leq 1/n$. On the other hand, the lower bound in Lemma 3.7 and (3.27) imply

$$\begin{aligned} r &\geq \alpha (\min_i v_i)(\min_i \varepsilon_i) \\ &\geq \alpha (\min_i \varepsilon_i)/(2n-2) \\ &\geq \alpha (\min_i \varepsilon_i)/(2n), \end{aligned}$$

where the second inequality is consequence of the fact that $\min_i v_i \geq 1/(2n-2)$. Hence, given any two sequences of social networks A_n and A'_n,

$$\begin{aligned} r_n/r'_n &\leq \left(\frac{1/n + \max_i 1/\gamma_{i,n}}{1/(2n)} \right) \frac{\max_i \varepsilon_{i,n}}{\min_i \varepsilon_{i,n}} \\ &\leq 2c(1 + \max_i n/\gamma_{i,n}). \end{aligned}$$

Assumption (c) of the proposition then immediately implies that there exists a uniform upper bound on the right-hand side of the above inequality for all n, and hence,

$$\limsup_{n\to\infty} r_n/r'_n < \infty.$$

A similar argument shows

$$\liminf_{n\to\infty} r_n/r'_n > 0,$$

completing the proof. □

Proof of Lemma 3.6

The proof relies on the following theorem, which is a corollary to Theorems 1 and 2 of Barreira and Valls (2007).

Theorem (Barreira and Valls)**.** *Consider the dynamical system $\varphi_t(x) = P_t x + f_t(x)$ with trajectory $x_t \in \mathbb{R}^k$, where $P_t \in \mathbb{R}^{k\times k}$ for all t. Also, suppose that the following hold:*

(a) The linear dynamical system generated by $\{P_t\}_{t\in\mathbb{N}}$ is Lyapunov regular.

(b) *The TLE corresponding to* $\{P_t\}_{t\in\mathbb{N}}$ *is negative, that is,* $\varsigma = \lim_{t\to\infty} \frac{1}{t}\log\|P_t P_{t-1}\dots P_1\| < 0$.

(c) $f_t(x)$ *is a continuous map with* $f_t(0) = 0$ *for all* t.

(d) *There are constants* $C, q > 0$ *such that* $\|f_t(x) - f_t(y)\| \le C\|x-y\|\big(\|x\|^q + \|y\|^q\big)$, *for all* t *and all* $x, y \in \mathbb{R}^k$.

For all t_0 *and* $\epsilon > 0$, *there exist a neighborhood* V *of the origin and a constant* K *such that if* $x_{t_0} \in V$, *then for all* $t \ge t_0$,

$$\|x_t\| \le K e^{(t-t_0)(\varsigma+\epsilon)+\epsilon t_0}\|x_{t_0}\|.$$

We use the above theorem to prove Lemma 3.6. Recall that $\phi_t^\theta : \mathbb{R}^{n(|\Theta|-1)} \to \mathbb{R}^{n(|\Theta|-1)}$ denotes the function that maps $\bar{\mu}_{t-1}^\theta = (\mu_{it-1}^\theta(\hat\theta))_{i,\hat\theta\neq\theta}$ to $\bar{\mu}_t^\theta = (\mu_{it}^\theta(\hat\theta))_{i,\hat\theta\neq\theta}$. Hence, the dynamical system $\{\phi_t^\theta\}_{t\in\mathbb{N}}$ with trajectory $\bar{\mu}_t^\theta \in \mathbb{R}^{n(|\Theta|-1)}$ describes the evolution of the beliefs conditional on the realization of state θ.

Given any $x \in \mathbb{R}^{n(|\Theta|-1)}$, the mapping ϕ_t^θ can be decomposed as

$$\phi_t^\theta(x) = M_t^\theta x + f_t^\theta(x),$$

where M_t^θ, defined in (3.17), is the Jacobian of ϕ_t^θ evaluated at the origin, and $f_t : \mathbb{R}^{n(|\Theta|-1)} \to \mathbb{R}^{n(|\Theta|-1)}$ is the higher-order residual given by

$$\Big[f_t^\theta(x)\Big]_{i,\hat\theta} = \alpha\left(\frac{\ell_i^{\hat\theta}(\omega_{it})}{g_{it}^\theta(x)} - \frac{\ell_i^{\hat\theta}(\omega_{it})}{\ell_i^\theta(\omega_{it})}\right)x_{i,\hat\theta},$$

with $g_{it}^\theta(x)$ defined as

$$g_{it}^\theta(x) = \ell_i^\theta(\omega_{it})\left(1 - \sum_{\tilde\theta\neq\theta} x_{i,\tilde\theta}\right) + \sum_{\tilde\theta\neq\theta}\ell_i^{\tilde\theta}(\omega_{it})x_{i,\tilde\theta}. \tag{3.28}$$

Thus, if M_t^θ and f_t^θ satisfy the conditions of above theorem, then for all $\epsilon > 0$, there exists a neighborhood V of the origin and a constant K such that if $\bar{\mu}_{t_0}^\theta \in V$, then for all $t \ge t_0$,

$$e_t^\theta \le K e^{(t-t_0)(\zeta^\theta+\epsilon)+\epsilon t_0} e_{t_0}^\theta,$$

where we are using the fact that $e_t^\theta = \|\bar{\mu}_t^\theta\|_1$. Furthermore, note that by Proposition 3.1, for almost all ω, there exists a t_0 such that $\bar{\mu}_{t_0}^\theta \in V$. Therefore, given that $\epsilon > 0$ is arbitrary,

$$\limsup_{t\to\infty}\frac{1}{t}\log e_t^\theta \le \zeta^\theta,$$

which proves Lemma 3.6. Thus, to complete the proof it is sufficient to verify that conditions (a)–(d) of the theorem of Barreira and Valls are satisfied for almost all ω.

Since the collection of matrices $\{M_t^\theta\}_{t\in\mathbb{N}}$ are independent and identically distributed, by the well-known Multiplicative Ergodic Theorem of Oseledets (1968), the linear dynamical system $\{M_t^\theta\}_{t\in\mathbb{N}}$ is Lyapunov regular for almost all ω.[29] This guarantees that assumption (a) is satisfied

[29] The original Multiplicative Ergodic Theorem due to Oseledets (1968) requires matrices $\{M_t^\theta\}_{t\in\mathbb{N}}$ to be invertible.

for almost all ω. Furthermore, recall that, as we showed in Lemma 3.5, the TLE corresponding to M_t^θ is strictly negative with $\mathbb{P}^\theta$-probability one, which implies that assumption (b) is also satisfied. To verify assumption (c), note that $f_t^\theta(0) = 0$. Moreover, since $\ell_i^{\hat{\theta}}$ has full support for all i and $\hat{\theta}$, $f_t^\theta(x)$ is continuous, implying that assumption (c) holds as well. In the remainder of the proof, we show that condition (d) of the theorem, which is a form of Lipschitz-continuity, is also satisfied.

For any given $x \in \mathbb{R}^{n(|\Theta|-1)}$, define $G_t^\theta(x)$ as the diagonal matrix with diagonal entries

$$\left[G_t^\theta(x)\right]_{i,\hat{\theta}} = \alpha \ell_i^{\hat{\theta}}(\omega_{it}) \left(\frac{1}{g_{it}^\theta(x)} - \frac{1}{\ell_i^\theta(\omega_{it})} \right).$$

Note that, by construction, $f_t^\theta(x) = G_t^\theta(x)x$. Furthermore,

$$\left[G_t^\theta(x)\right]_{i,\hat{\theta}} = \frac{\alpha \ell_i^{\hat{\theta}}(\omega_{it})}{g_{it}^\theta(x) \ell_i^\theta(\omega_{it})} \sum_{\tilde{\theta} \neq \theta} \left(\ell_i^\theta(\omega_{it}) - \ell_i^{\tilde{\theta}}(\omega_{it}) \right) x_{i,\tilde{\theta}},$$

where we are replacing $\ell_i^\theta(\omega_{it}) - g_{it}^\theta(x)$ from the definition of $g_{it}^\theta(x)$ in (3.28). On the other hand, given that $\ell_i^{\hat{\theta}}$ has full support for all i and $\hat{\theta}$,

$$\frac{\alpha \ell_i^{\hat{\theta}}(\omega_{it})}{g_{it}^\theta(x) \ell_i^\theta(\omega_{it})} \leq C_1 < \infty,$$

for some positive constant C_1 defined as

$$C_1 = \alpha \frac{\max_{\hat{\theta},i,s} \ell_i^{\hat{\theta}}(s)}{\min_{\hat{\theta},i,s} \left(\ell_i^{\hat{\theta}}(s) \right)^2}.$$

Hence, for any $x, y \in \mathbb{R}^{n(|\Theta|-1)}$,

$$\begin{aligned} \left| \left[G_t^\theta(x) - G_t^\theta(y) \right]_{i,\hat{\theta}} \right| &\leq C_1 \left| \sum_{\tilde{\theta} \neq \theta} \left(\ell_i^\theta(\omega_{it}) - \ell_i^{\tilde{\theta}}(\omega_{it}) \right) \left(x_{i,\tilde{\theta}} - y_{i,\tilde{\theta}} \right) \right| \\ &\leq C_1 \sum_{\tilde{\theta} \neq \theta} \left| \ell_i^\theta(\omega_{it}) - \ell_i^{\tilde{\theta}}(\omega_{it}) \right| \left| x_{i,\tilde{\theta}} - y_{i,\tilde{\theta}} \right|. \end{aligned}$$

Let

$$C_2 = C_1 \max_{\tilde{\theta} \neq \theta, i, s} \left| \ell_i^\theta(s) - \ell_i^{\tilde{\theta}}(s) \right|,$$

In the current setting, this condition is satisfied for a generic set of parameter values. Alternatively, one can use more general variants of the theorem that extend the results to the case of noninvertible matrices.

which by Assumption 3.1 is guaranteed to be strictly positive. Then,

$$\begin{aligned}\left|\left[G_t^\theta(x)-G_t^\theta(y)\right]_{i,\hat{\theta}}\right| &\le C_2\sum_{\tilde{\theta}\neq\theta}\left|x_{i,\tilde{\theta}}-y_{i,\tilde{\theta}}\right|\\ &\le C\max_{\tilde{\theta}\neq\theta}\left|x_{i,\tilde{\theta}}-y_{i,\tilde{\theta}}\right|,\end{aligned}$$

where $C=(|\Theta|-1)C_2>0$. Consequently,

$$\begin{aligned}\left\|G_t^\theta(x)-G_t^\theta(y)\right\|_\infty &= \max_i\max_{\hat{\theta}\neq\theta}\left|\left[G_t^\theta(x)-G_t^\theta(y)\right]_{i,\hat{\theta}}\right|\\ &\le C\max_i\max_{\tilde{\theta}\neq\theta}\left|x_{i,\tilde{\theta}}-y_{i,\tilde{\theta}}\right|\\ &\le C\left\|x-y\right\|_\infty.\end{aligned}\tag{3.29}$$

Therefore, noting that $f_t^\theta(x)=G_t^\theta(x)x$ implies

$$\begin{aligned}\|f_t(x)-f_t(y)\| &= \|G_t^\theta(x)x-G_t^\theta(y)y\|\\ &\le \left\|G_t^\theta(x)x-G_t^\theta(x)y\right\|+\left\|G_t^\theta(x)y-G_t^\theta(y)y\right\|\\ &\le \left\|G_t^\theta(x)\right\|\left\|x-y\right\|+\left\|G_t^\theta(x)-G_t^\theta(y)\right\|\left\|y\right\|,\end{aligned}$$

where the first inequality is due to the triangle inequality. Subsequently, using the fact that $G_t^\theta(0)=0$ for all t and ω leads to

$$\begin{aligned}\|f_t(x)-f_t(y)\| &\le \left\|G_t^\theta(x)-G_t^\theta(0)\right\|\left\|x-y\right\|+\left\|G_t^\theta(x)-G_t^\theta(y)\right\|\left\|y\right\|\\ &\le C\|x\|\left\|x-y\right\|+C\left\|x-y\right\|\left\|y\right\|\\ &= C\left\|x-y\right\|\left(\|x\|+\left\|y\right\|\right),\end{aligned}$$

where the second inequality is a consequence of (3.29). This establishes that condition (d) of the theorem of Barreira and Valls is satisfied for $q=1$, and hence, completes the proof. □

Chapter 4

Social Learning in a Coordination Game

This chapter examines the coordination problem faced by a group of agents when the relevant information is dispersed throughout a social network. We use the framework of dynamic games of incomplete information to model the agents' coordination problem. A number of agents play a game with payoffs that have two components: an estimation term and a coordination term. The estimation term serves to capture the agents' desire to make decisions that are optimal given their private information about an unknown parameter. The coordination term captures the payoffs agents receive by taking actions that are close to the average action taken by the rest of the population. The game is played over multiple stages. At each stage of the game, agents observe the previous choices made by a subset of other agents, called their neighbors. An agent's action may reveal some information to her neighbors that was previously unknown to them. The neighbors can use this information to reevaluate their beliefs about the underlying parameter and their predictions of others' future behavior. These reevaluations may, in turn, lead agents to revise their actions over time.

Given this dynamic environment, different behavioral assumptions lead to different outcomes. In particular, the way agents revise their views in face of new information and the actions they choose given these views determine the long run outcome of the game. We assume that agents are Bayesian and myopic. Bayesian agents use Bayes' rule to incorporate new observations in their beliefs. Myopic agents choose actions at each stage of the game which maximize their stage payoffs, without regard for the effect of these actions on their future payoffs. The assumption on myopic agent behavior simplifies the analysis significantly and results in an essentially unique equilibrium, which is unlikely with forward-looking agents. We use this behavioral assumption to define an equilibrium, and prove formal results regarding the agents' asymptotic equilibrium behavior, assuming a quadratic utility function.

Our analysis yields several important results. First, each agent's action asymptotically converges to some limit action. By making use of this result, we show that if an agent (she) observes the actions of some other agent (he) infinitely often, she will eventually be able to imitate his actions and achieve a payoff at least as high as his limit payoffs. We then use this argument to prove that if the social network is sufficiently connected over time, agents asymptotically receive similar payoffs. In our symmetric coordination game, this implies that different agents' actions also converge to the same value. In other words, agents eventually coordinate on the same action. These results extend some of the results in the social learning literatures to the setting where each agent's actions directly affect others' payoffs. To the best of our knowledge, this is the first such result on reaching consensus in social networks in presence of payoff externalities.

This chapter is based on Molavi, Eksin, Ribeiro, and Jadbabaie (2013).

Second, we show that if the agents' private observations are only functions of the unknown state (and not their own actions), then generically the agents eventually coordinate on the "efficient" action—the action on which the agents would have coordinated if each agent had access to the private observations of every other one. Thus, the dispersed information is asymptotically optimally aggregated through the agents' repeated interactions. This result is true because the agents play a coordination game wherein their incentives are aligned, and hence, they do not have an incentive to withhold their private information. This theorem extends Theorem 2.1 of this thesis, as well as the result of Mueller-Frank (2013) on optimal aggregation of information in Bayesian learning, to the cases where the state space is not finite and the agents face payoff externalities.

The work in this chapter is related to three main lines of research. The first is the literature on Bayesian learning over networks. The focus of the social learning literature is on modeling the way agents use their observations to update their beliefs and characterizing the outcomes of the learning process. Examples include, Bikhchandani, Hirshleifer, and Welch (1992), Banerjee (1992), and Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) that study sequential decision problems; and Borkar and Varaiya (1982), Gale and Kariv (2003), Rosenberg, Solan, and Vieille (2009), and Mueller-Frank (2013) that study repeated and simultaneous interactions. Due to the complexity of social learning, the focus in the latter family of models is on asymptotic outcomes. In this chapter of the thesis, we extend the repeated Bayesian social learning framework to an environment with payoff externalities, i.e., one where an agent's stage payoff is a function of other agents' actions.

The current work is also related to the literature on learning in games, such as the works by Jordan (1991, 1995), Kalai and Lehrer (1993), Jackson and Kalai (1997), Nachbar (1997), and Foster and Young (2003). The central question in this literature is whether agents learn to play a Nash (or Bayesian Nash) equilibrium. Whereas, in the current chapter of the thesis, the focus is on whether agents in a network asymptotically receive the same payoffs and whether they optimally aggregate the dispersed information.

Finally, our work is related to the literature in economic theory that studies the effect of public and private information on welfare, pioneered by the work of Morris and Shin (2002) who study the effect of public information on the equilibrium welfare when agents play a beauty contest game. In this chapter of the thesis, we borrow the payoff function introduced by Morris and Shin (2002) to model the agents' coordination problem. However, unlike the model of Morris and Shin, the focus of the current work is on coordination and aggregation of information dispersed in social networks. Among other related papers that study effect of public and private information on welfare are the works by Angeletos and Pavan (2007, 2009), Vives (2010), and Amador and Weill (2012).

4.1 Baseline Model

4.1.1 Agents and Payoffs

Consider n agents indexed by $i \in N = \{1, \ldots, n\}$ who repeatedly play a game with uncertain payoffs. The payoff-relevant uncertainty is captured by a common unknown parameter θ, called the state of the world, that takes values in $\Theta = \mathbb{R}$. Agents start with a common prior belief about θ denoted by ν. We make the following technical assumption on ν.

Assumption 4.1. The state is square integrable with respect to ν, that is,

$$\int_{\Theta} \theta^2 d\nu < \infty.$$

The game is played over a countable set of time periods that is indexed by the positive integers. At time t, each agent observes a private signal in addition to the time $t-1$ actions of a subset of agents, takes an action simultaneously with other agents, and receives a payoff. We use $s_{it} \in S_i$ to denote the private signal observed by agent i at time t, where S_i is a complete separable metric space, and use $s_t = (s_{1t}, \ldots, s_{nt}) \in S = \times_{i=1}^{n} S_i$ to denote the corresponding signal profile. Furthermore, we let $a_{it} \in A_i = \mathbb{R}$ denote the action taken by agent i at time t, and let $a_t = (a_{1t}, \ldots, a_{nt}) \in A = \mathbb{R}^n$ denote the corresponding action profile. Finally, $u_i(a, \theta)$ denotes the stage payoff received by agent i when agents play the action profile a and given that the realized state is θ. Agent i's stage payoff has the following representation:

$$u_i(a, \theta) = -(1-\lambda)(a_i - \theta)^2 - \lambda(a_i - \overline{a}_{-i})^2, \tag{4.1}$$

where $\lambda \in [0, 1)$ is a constant and $\overline{a}_{-i} = \frac{1}{n-1}\sum_{j \neq i} a_j$ denotes the average payoff across other agents. The first term is a quadratic loss in the distance between the realized state and agent i's action, capturing the agent's preference for actions which are close to the unknown state. The second term is the "beauty contest" term representing the agent's preference for acting in conformity with the rest of the population. This utility function was introduced by Morris and Shin (2002) to represent the preferences of the agents who engage in second-guessing others' actions as postulated by Keynes (1936).

4.1.2 Social Network

At time $t+1$, in addition to her private signal, each agent also observes the time t actions of a subset of other agents, denoted by $N_{it} \subseteq N$ and called her time t neighbors. We use the convention that agents are their own neighbors at all times, that is, $i \in N_{it}$ for all i and t. The time t interactions between agents can be summarized by a directed network $g_t \in G = \{0, 1\}^{n \times n}$ where $\left[g_t\right]_{ji} = 1$ if and only if agent j is a time t neighbor of agent i, that is, if $j \in N_{it}$.

Assumption 4.2. The network g_t is generated according to some probability distribution ψ_t independently of other random variables in the model.

This assumption is satisfied by many commonly used models of social networks such as fixed networks, i.i.d networks, and deterministically time-varying ones. But it excludes the case where an agent's realized neighborhood is informative about the state or other agents' signals. We maintain Assumption 4.2 throughout this chapter.

A directed path from i to j is a sequence of agents starting with i and ending with j such that each agent is a neighbor of the next one in the sequence. We say that a social network is strongly connected if there exists a directed path from each node to any other. Let $\psi = \times_{t=1}^{\infty} \psi_t$ denote the probability distribution over the sequences of networks $\{g_t\}_{t \in \mathbb{N}}$. We impose the following mild connectivity assumption on the networks generated by the stochastic process ψ.

Assumption 4.3. For ψ-almost all $\{g_t\}_{t \in \mathbb{N}}$, there exists a strongly connected network $\widehat{g}$ such that if j is a neighbor of i given $\widehat{g}$, then j is also a neighbor of i given g_t for infinitely many t.

The above assumption guarantees that information obtained by an agent at any given time period can eventually flow to any other agent in the network. That said, we have to remark that an agent's private information may never become available to other agents. Whether this is indeed the case depends on the actions chosen by the agents in the equilibrium of the game.

4.1.3 Histories

Let $(\Omega, \mathscr{B})$ denote the measurable space of plays, where $\Omega = \Theta \times (S \times A \times G)^{\mathbb{N}}$ and $\mathscr{B}$ is the corresponding Borel σ-algebra. A generic element of the set Ω is denoted by ω and is called a *path of play*. This is an infinite history of the game, consisting of the state and a list of all the private signals, actions, and realized networks at all time periods. Similarly, let h_t denote the time t *history* of the game defined recursively as

$$h_t = (h_{t-1}; s_{t-1}, a_{t-1}, g_{t-1}),$$

with $h_1 = \theta$. This is a complete description of the game up to time period t that belongs to the measurable space $H_t = \Theta \times (S \times A \times G)^{t-1}$. We let $\mathscr{H}_t \subseteq \mathscr{B}$ denote the σ-algebra of subsets of Ω generated by the Borel sets of H_t.

Agents' private signals are endogenously generated according to some probability distribution which is a function of the history of the game. Given $h_t \in H_t$, the time t signal profile is generated according to the probability distribution $\pi_t(h_t)[\cdot]$, where π_t is a transition probability from H_t to S.[30]

The time t *private history* of agent i is a list of all of her observations, denoted by h_{it} and defined recursively as

$$h_{it} = \left(h_{it-1}; s_{it-1}, (a_{jt-1})_{j \in N_{it-1}}\right),$$

with $h_{i1} = \emptyset$. We let H_{it} denotes the set of agent i's time t private histories, let $H_i = \cup_{t=1}^{\infty} H_{it}$ denote the set of agent i's private histories of any length, and let $\mathscr{H}_{it} \subseteq \mathscr{H}_t$ and $\mathscr{H}_i \subseteq \mathscr{B}$ denote the σ-algebras of subsets of Ω generated by the Borel sets of H_{it} and H_i, respectively.

4.1.4 Strategies and Belief Systems

A strategy is a function that maps an agent's private histories to her actions, whereas a belief system is mapping from private histories to probability distributions over the space of plays.

Definition 4.1. A *pure behavior strategy* for agent i is a measurable function $\sigma_i : H_i \to A_i$.

Agent i's strategy is a complete contingency plan determining the action to be taken by her at all time periods and given any private history. More generally, the joint behavior of the agents is fully described by the strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$, where σ_i is a strategy for agent i.

Any strategy profile σ—together with the agents' common prior ν, the stochastic process ψ, and the signaling functions $\{\pi_t\}_{t \in \mathbb{N}}$—induces a probability distribution over the measurable space $(\Omega, \mathscr{B})$, denoted by $\mathbb{P}^\sigma$. We let $\mathbb{E}^\sigma$ denote the expectation operator corresponding to $\mathbb{P}^\sigma$. Given that agents follow the strategy profile σ, the path of play ω is simply a point in the probability space $(\Omega, \mathscr{F}, \mathbb{P}^\sigma)$. The realized time t private history of agent i is in turn a measurable function

[30] Given measurable spaces $(X, \mathscr{X})$ and $(Y, \mathscr{Y})$, a function $f : X \times \mathscr{Y} \to [0,1]$ is called a transition probability from X to Y if (i) for any given $x \in X$, $f(x)[\cdot]$ is a probability distribution over $(Y, \mathscr{Y})$; and (ii) given any measurable set $B \in \mathscr{Y}$, the function $x \mapsto f(x)[B]$ is measurable.

of the realized path of play, denoted by $\tilde{h}_{it}(\cdot) : \Omega \to H_{it}$. We let $\tilde{\sigma}_{it}(\cdot) = \sigma_i(\tilde{h}_{it}(\cdot)) : \Omega \to A_i$ denote the random variable that determines the time t action of agent i as a function of the realized path of play ω.

Definition 4.2. A *belief system* for agent i is a transition probability $q_i : H_i \times \mathscr{B} \to [0,1]$.

A belief is a probably distribution over the space of plays $(\Omega, \mathscr{B})$, whereas a belief system is a collection of beliefs—one for every possible private history—that describes the agent's belief after observing any private history. More generally, the beliefs of the agents are fully described by $q = (q_1, \ldots, q_n)$, where q_i is a belief system for agent i. Finally, given a belief system q_i, we let $\tilde{q}_{it}(\cdot)[\cdot] = q_i(\tilde{h}_{it}(\cdot))[\cdot] : \Omega \times \mathscr{B} \to [0,1]$ denote the transition probability that determines agent i's time t belief as a function of ω.

4.2 Equilibrium

In this section, we introduce our equilibrium notion and provide a characterization of the equilibrium behavior. Our notion is a variant of the weak perfect Bayesian Equilibrium according to which (i) agents' strategies maximize their expected stage payoffs given their beliefs; and (ii) agents' equilibrium beliefs are consistent with their strategies. Before formally presenting our equilibrium notion, we introduce some notation.

Agent i's expected utility of taking an action is dependent on her belief about the path of play as well as what she expects other agents to do. However, if we fix a strategy profile, the other agents' actions are only functions of the realized path of play. Thus, given a strategy profile σ, the expected time t payoff to agent i of taking action a_i is uniquely determined as a function of her belief p_i over $(\Omega, \mathscr{B})$ as

$$v_{it}(a_i, \sigma_{-i}; p_i) = \int_{\Omega} u_i\left(a_i, \tilde{\sigma}_{-it}, \theta\right) dp_i,$$

where $\tilde{\sigma}_{-it} = \left(\tilde{\sigma}_{jt}\right)_{j \neq i}$.

Definition 4.3. A *weak perfect Bayesian equilibrium* consists of a strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ and a collection of belief systems $q^* = (q_1^*, \ldots, q_n^*)$ that satisfy the following conditions for all i and t.[31]

(a) For $\mathbb{P}^*$-almost all $h_{it} \in H_{it}$ and all $a_i \in A_i$,

$$v_{it}\left(\sigma_i^*(h_{it}), \sigma_{-i}^*; q_i^*(h_{it})\right) \geq v_{it}\left(a_i, \sigma_{-i}^*; q_i^*(h_{it})\right).$$

(b) $\tilde{q}_{it}^*$ is a regular conditional probability of $\mathbb{P}^*$ given $\mathscr{H}_{it}$.[32]

According to the first condition, in equilibrium agents do not have access to profitable unilateral deviations given all, but possibly a set of measure zero, of private histories. Bayesian Nash equilibrium is typically defined by requiring the agents to maximize their expected utilities given all information sets, including the ones that are reached with zero probability. Our equilibrium

[31] We use $\mathbb{P}^*$ and $\mathbb{E}^*$ to denote the probability distribution and expectation operator, respectively, induced by σ^*.

[32] Given a probability space $(X, \mathscr{X}, \mathbb{P})$ and a sub σ-algebra $\mathscr{Y} \subseteq \mathscr{X}$, the transition probability $f : X \times \mathscr{X} \to [0,1]$ is a regular conditional probability of $\mathbb{P}$ given $\mathscr{Y}$ if for each $B \in \mathscr{X}$, $x \mapsto f(x)[B]$ is a version of $\mathbb{P}(B|\mathscr{Y})$.

notion is thus weaker than the standard Bayesian Nash equilibrium. However, the requirement is sufficiently strong to ensure the existence of an equilibrium that is unique up to sets of measure zero.

The assumption that the agents maximize their stage payoffs corresponds to myopia on agents' behalf. An alternative equilibrium notion is obtained by assuming that the agents choose actions that maximize the average (or discounted sum) of their payoffs over their lifetime. However, using this alternative equilibrium notion significantly complicates the analysis and more importantly results in multiplicity of equilibria.

The second equilibrium condition is the consistency requirement according to which the agents' beliefs are obtained using Bayes' rule given their prior and the equilibrium strategy profile σ^*. We remark that, as typically is the case with weak perfect Bayesian equilibria, agents' beliefs are not uniquely determined given the equilibrium strategy profile. Rather, any regular probability distribution distribution of $\mathbb{P}^*$ given $\mathscr{H}_{it}$ is a consistent time t belief for agent i.

Definition 4.4. σ^* is an *equilibrium strategy profile* if there exists some q^* such that the pair (σ^*, q^*) constitutes an equilibrium.

The following lemma provides a characterization of the equilibrium strategy profiles.

Lemma 4.1. *σ^* is an equilibrium strategy profile if and only if for all i and t,*

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta)|\mathscr{H}_{it}\right] \geq \mathbb{E}^*\left[u_i(\tilde{\sigma}_{it}, \tilde{\sigma}^*_{-it}, \theta)|\mathscr{H}_{it}\right], \tag{4.2}$$

for any strategy σ_i and with $\mathbb{P}^*$-probability one.

In the rest of the current chapter, we restrict our attention to square integrable strategies in order to rule out the uninteresting equilibria wherein each agent's expected payoff is equal to minus infinity, regardless of her own strategy.

Definition 4.5. A strategy profile σ is *square integrable* if

$$\mathbb{E}^{\sigma}\left[\tilde{\sigma}^2_{it}\right] < \infty,$$

for all i and t.

If agents follow square integrable strategies, their expected payoffs of taking any action given any private history is finite. Moreover, agents' expected stage payoffs are quadratic, and concave in their own actions. Thus, the equilibria of the game can be characterized by a set of necessary and sufficient first-order conditions that result in the following simple characterization of the square integrable strategy profiles.

Corollary 4.1. *The square integrable strategy profile σ^* is an equilibrium strategy profile if and only if for all i and t,*

$$\tilde{\sigma}^*_{it} = (1-\lambda)\mathbb{E}^*\left[\theta|\mathscr{H}_{it}\right] + \lambda \frac{1}{n-1}\sum_{j\neq i}\mathbb{E}^*\left[\tilde{\sigma}^*_{jt}|\mathscr{H}_{it}\right], \tag{4.3}$$

with $\mathbb{P}^*$-probability one.

Agents' equilibrium strategies are linear in their expectation of the state and others' actions. This feature of the equilibrium keeps the analysis tractable. Moreover, equation (4.3) can be used to show that square integrable strategy profiles are the fixed points of a contraction mapping in the L^p space. We use this property to show that square integrable equilibrium strategies always exist and result in equilibrium actions which are almost always unique.

Proposition 4.1. *Suppose that Assumption 4.1 is satisfied. Then, a square integrable equilibrium strategy profile σ^* exists. Furthermore, for any other square integrable equilibrium strategy profile $\sigma^\dagger$ and all i and t,*

$$\tilde{\sigma}^*_{it} = \tilde{\sigma}^\dagger_{it},$$

$\mathbb{P}^*$-almost surely and $\mathbb{P}^\dagger$-almost surely.

Thus, the agents' equilibrium actions are uniquely determined after a set of full measure of histories. In the next section, we use this result and the characterization of the equilibrium actions in Corollary 4.1 to analyze the asymptotic behavior of the agents' equilibrium actions.

4.3 Reaching Consensus

In this section, we show that the agents eventually reach consensus in their actions and that their realized payoffs are asymptotically the same. To prove these results, we first show that agents' actions converge to some limit action.

Proposition 4.2. *Suppose that Assumption 4.1 is satisfied. Let σ^* be a square integrable equilibrium strategy profile. Then, $\tilde{\sigma}^*_{it}$ converges to some $\mathscr{H}_i$-measurable random variable $\tilde{\varsigma}^*_i$ in the L^2 sense, that is,*

$$\mathbb{E}^*\left[\left(\tilde{\sigma}^*_{it} - \tilde{\varsigma}^*_i\right)^2\right] \longrightarrow 0 \qquad \text{as} \qquad t \to \infty,$$

for all $i \in N$.

Agent j's action converges in L^2 to some limit action that is a function of the realized path of play. If agent i can observe the actions of j infinitely often, she can asymptotically imitate the actions of agent j. In a strongly connected network, agent j can in turn imitate the actions of some other agent k, and so on, with some agent being able to imitate the actions of agent i. All agents in such a chain must, therefore, asymptotically believe that their actions are better than the ones taken by the others. However, since the agents' payoffs are symmetric and their actions are strategic complements, this is only possible if any two agents asymptotically choose the same action, regardless of the realization of the state of the world. This is an instance of argument by the so-called Imitation Principle, according to which each agent's asymptotic payoff is always at least as high as the asymptotic payoff of any agent she can imitate, and hence, in a connected network, agents' asymptotic payoffs are the same.[33] We use this line of reasoning to prove the following theorem on consensus in the agents' payoffs and actions.

Theorem 4.1. *Suppose that Assumptions 4.1–4.3 are satisfied. Let σ^* be a square integrable equilibrium strategy profile. Then, as t goes to infinity, for all $i, j \in N$,*

[33] The Imitation Principle was first introduced by Bala and Goyal (1998) to study boundedly rational social learning with purely informational externalities. For other applications of the Imitation Principle, see Gale and Kariv (2003) and Rosenberg, Solan, and Vieille (2009).

(a) $\mathbb{E}^*\left[\left|u_i\left(\tilde{\sigma}^*_{it},\tilde{\sigma}^*_{-it},\theta\right)-u_j\left(\tilde{\sigma}^*_{jt},\tilde{\sigma}^*_{-jt},\theta\right)\right|\right]\longrightarrow 0,$

(b) $\mathbb{E}^*\left[\left(\tilde{\sigma}^*_{it}-\tilde{\sigma}^*_{jt}\right)^2\right]\longrightarrow 0,$

(c) $\tilde{\sigma}^*_{it}\xrightarrow{L^1}\mathbb{E}^*\left[\theta|\mathscr{H}_i\right]=\mathbb{E}^*\left[\theta|\mathscr{H}_j\right].$

According to part (a) of the theorem, the differences between the agents' payoffs asymptotically vanish in the L^1 sense. Thus, in spite of the differences in their location in the network and the quality of their private signals, agents asymptotically receive similar payoffs. This is due to the structure of the game wherein agents' incentives are aligned, and thus, each agent would benefit from making her private information available to the rest of the population. From the point of view of the agents, however, the asymptotic payoffs are not necessarily the same. That is, the conditional expectations of the agents' limit payoffs given their information at the end of the game could be dissimilar. The following example illustrates this possibility.

Example 4.1. Consider two agents who observe each others' actions at all time periods. The common prior is the uniform distribution over the set $\{-2,-1,1,2\}$. Agent 2 receives no signal ($S_2=\emptyset$), whereas Agent 1's private signals belong to the set $S_1=\{1,2\}$ and her signaling functions π_t are given by

$$\pi_t(h_t)=\begin{cases}\delta_1 & \text{if} \quad |\theta|<2,\\ \delta_2 & \text{if} \quad |\theta|\geq 2,\end{cases}$$

where $\delta_{s_{1t}}$ is the degenerate probability distribution with unit mass on the signal $s_{1t}\in S_1$. Thus, agent 1 is informed of the absolute value of θ. Observe that in any equilibrium of the game $\tilde{\sigma}^*_{it}=0$ at all times and for both agents, Agent 1 learns the absolute value of θ, whereas Agent 2 never makes any informative observations. At the end of the game, Agent 1's expected payoff conditional on her information is equal to $-(1-\lambda)|\theta|^2$, while the corresponding payoff for Agent 2 is given by $-(1-\lambda)\frac{5}{2}$. Although these conditional expected payoffs are unequal for any realization of the state, the unconditional expected payoffs and the realized payoffs are the same for both agents—as also implied by Theorem 4.1.

Part (b) of the theorem proves that the agents asymptotically coordinate their actions without ever communicating their private signals, whereas part (c) shows that agents asymptotically reach an agreement in their conditional expectations of the state. Nevertheless, it is not immediately obvious whether the agents coordinate on the "optimal" action—on which they would have coordinated, had they been able to fully communicate their private signals—or whether their consensus estimate of the state is the best possible. The following example shows that this may indeed not be the case.

Example 4.2. Consider two agents who observe each others' actions at all time periods. The common prior $\mathbb{P}$ is the uniform distribution over the set $\{-1,1\}$. Agents' private signals belong to the sets $S_1=S_2=\{\text{H},\text{T}\}$ and the signaling functions π_t are given by

$$\pi_t(h_t)=\begin{cases}\dfrac{1}{2}\delta_{(\text{H},\text{H})}+\dfrac{1}{2}\delta_{(\text{T},\text{T})} & \text{if} \quad \theta\geq 0,\\ \dfrac{1}{2}\delta_{(\text{H},\text{T})}+\dfrac{1}{2}\delta_{(\text{T},\text{H})} & \text{if} \quad \theta<0,\end{cases}$$

where δ_{s_t} is the degenerate probability distribution with unit mass on the signal profile $s_t \in S$. We first show that, in the unique equilibrium of the game, both agents choose $\tilde{\sigma}^*_{it} = 0$ at all times. Given the prior, in any equilibrium of the game, agents choose $\tilde{\sigma}^*_{i1} = 0$ at $t = 1$. Agents then each receive a signal that is H (T) with probability one half, regardless of the realization of θ. Agents' private signals are thus completely uninformative about the realized state. As a result, agents also choose $\tilde{\sigma}^*_{i2} = 0$ at $t = 2$, regardless of the realized state. These actions reveal no information; moreover, the time 2 private signals are uninformative. Therefore, agents continue to choose the zero action in all subsequent stages of the game.

Next, consider the alternative setting in which both agents observe the signal profile $s_t = (s_{1t}, s_{2t})$ at time t. (This setup is equivalent to one in which each agent communicates her private signals to the other.) In this modified game, both of the agents learn the realized state at $t = 2$. Therefore, in any equilibrium $\sigma^\dagger$ of the modified game both agents choose $\tilde{\sigma}^\dagger_{it} = \theta$ for all $t \geq 2$ and given any realization of θ. This shows that in the original game the agents did not coordinate on the optimal action—which they would have chosen if they had observed each others' private signals.

In the above example, the information content of the private signals is not successfully aggregated through the agents' repeated interactions. The reason for this failure is that the agents' equilibrium actions reveal no information about their private signals, although the signals contain useful information about the realized state. This example is however nongeneric in the sense that the transition probabilities π_t are "fine-tuned" to make all the states equally likely after the observation of any private signal. In the next section, we argue that when the signals are exogenously generated—as is in fact the case in Example 4.2—the agents generically coordinate on the action that is efficient given their aggregate information.

4.4 Exogenous Signals and Asymptotic Efficiency

In this section, we provide conditions under which agents aggregate the dispersed information and asymptotically coordinate on the efficient action. Dynamic games of incomplete information of the type discussed in this chapter generally exhibit two distinct inefficiencies. The first inefficiency is the result of the payoff externality whereby agents try to second-guess the actions of others by choosing actions that are close to their estimates of the average action across the population. A social planner that wants to maximize the sum of agents' payoffs, in contrast, would make them take actions which are simply close to their estimates of the state. This inefficiency is present even in static variants of the game, such as the model studied by Morris and Shin (2002). Theorem 4.1 of Section 4.3 shows that this inefficiency asymptotically disappears as each agent learns to correctly predict the actions of other agents.

The second inefficiency is due to the informational externalities present in a dynamic setting, wherein agents do not internalize the effect of their actions on the informativeness of the future observations. This inefficiency is also present in models of social learning, such as the model proposed by Vives (1997), in which each agent's payoff is independent of the actions taken by the rest of the population. This learning inefficiency could especially be severe if the distribution of the agents' private signals is a function of their previous actions. The following example illustrates some of the complications that can arise with endogenously generated signals.

Example 4.3. Consider a single agent who repeatedly plays a game with payoffs as in (4.1) with $\lambda = 0$. The agent's prior is given by the standard normal $\mathcal{N}(0,1)$. The signaling functions are given by $\pi_t(h_t) = \mathcal{N}(\theta, 1)$ for $t \leq 2$, and

$$\pi_t(h_t) = \begin{cases} \mathcal{N}(\theta, 1) & \text{if} \quad |\theta - a_2| > 1, \\ \mathcal{N}(0, 1) & \text{if} \quad |\theta - a_2| \leq 1. \end{cases}$$

for $t > 2$. The agent observes informative signals and chooses actions in the first two periods. If her time 2 action is not within unit distance of the realized state, she continues to observe informative private signals and asymptotically learns the state with arbitrary precision. However, if the agent's time 2 action is sufficiently close to the realized state, she does not observe any informative signals after the second time period and thus never learns the state. In this example, there is an externality associated with the effect of the agent's time 2 action on the distribution of the private signals observed by her future incarnations. If the agent is myopic (or sufficiently impatient), this informational externality is not internalized in the equilibrium.

This example illustrates the path dependence that learning with endogenously generated signals can exhibit: The total amount of information available to the agents is not fixed; it rather is a function of the realized path of play. Consequently, no well-defined notion of the efficient aggregation of information is readily available when learning is endogenous. We thus restrict our attention in this section to a setting where the signals are exogenously generated in the following sense.

Definition 4.6. The private signals are *exogenously generated* if for any t, there exists some transition probability $\widehat{\pi}_t$ from $\Theta \times S^{t-1}$ to S such that for all $h_t = (\theta; s_1, a_1, g_1; \ldots; s_{t-1}, a_{t-1}, g_{t-1}) \in H_t$ one has $\pi_t(h_t) = \widehat{\pi}_t(\theta; s_1; s_2; \ldots; s_{t-1})$.

To simplify the analysis, we also replace Assumption 4.1 with Assumption 1′ below and replace Assumptions 4.2 and 4.3 with Assumption 2′ below.

Assumption 1′. The set Θ is a bounded and measurable subset of $\mathbb{R}$.

Assumption 2′. There exists a strongly connected network $\widehat{g}$ such that $\widehat{g} = g_t$ for all t and with ψ-probability one.

When the private signals are exogenously generated and Assumptions 1′ and 2′ are satisfied, we can express our results more simply by using an alternative representation of the space of plays. Let $(\Xi, \mathscr{Z})$ be the measurable space with $\Xi = \Theta \times S^{\mathbb{N}}$ and $\mathscr{Z}$ the corresponding Borel σ-algebra. Any prior ν and signaling functions $\{\pi_t\}_{t\in\mathbb{N}}$ induce a probability distribution P over $(\Xi, \mathscr{Z})$ which is independent of the strategy profile followed by the agents—unlike in the case of endogenously generated private signals. On the other hand, given a strategy profile σ, the private history of agent i at time t is an H_{it}-valued random variable $\tilde{h}_{it} : \Xi \to H_{it}$. We define $\mathscr{I}^{\sigma}_{it}$ to be the σ-algebra generated by $\tilde{h}_{it}$ when agents follow the strategy profile σ and define $\mathscr{I}^{\sigma}_{i}$ to be the σ-algebra generated by the union of $\mathscr{I}^{\sigma}_{it}$ over all $t \in \mathbb{N}$. If the signals are exogenously generated, the results of the previous sections can alternatively be expressed in terms of the probability distribution P over the measurable space $(\Xi, \mathscr{Z})$. For instance, we have the following counterpart of Theorem 4.1, the proof of which is the same as the proof of Theorem 4.1 and is thus omitted.

Theorem 1′. *Suppose that the private signals are exogenously generated and Assumptions 1′ and 2′ are satisfied. Let σ^* be an equilibrium strategy profile. Then, as t goes to infinity, for all $i, j \in N$,*

(a) $E\left[\left|u_i\left(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta\right) - u_j\left(\tilde{\sigma}^*_{jt}, \tilde{\sigma}^*_{-jt}, \theta\right)\right|\right] \longrightarrow 0,$

(b) $E\left[\left(\tilde{\sigma}^*_{it} - \tilde{\sigma}^*_{jt}\right)^2\right] \longrightarrow 0,$

(c) $\tilde{\sigma}^*_{it} \xrightarrow{L^1} E\left[\theta | \mathscr{I}^*_i\right] = E\left[\theta | \mathscr{I}^*_j\right].$

Let $(\mathbf{P}, d)$ denote the metric space of all probability distributions over $(\Xi, \mathcal{Z})$, where d is the total variation distance. When agents' private signals are exogenously generated, any prior distribution ν and signaling functions $\{\pi_t\}_{t\in\mathbb{N}}$ induce some probability measure over $(\Xi, \mathcal{Z})$, and any probability distribution over $(\Xi, \mathcal{Z})$ is induced uniquely (up to sets of measure zero) by some prior and signaling functions. We can make use of this correspondence to define a *generic* set of priors and signaling functions as a set of priors and signaling functions such that their corresponding set of induced probability measures over $(\Xi, \mathcal{Z})$ is a residual subset of $\mathbf{P}$, where here and in the rest of this chapter of the thesis we assume that $\mathbf{P}$ is endowed with the topology of uniform convergence (metrized by d).[34] We have the following result on the generic optimality of asymptotic actions.

Theorem 4.2. *Suppose that the private signals are exogenously generated and Assumptions 1′ and 2′ are satisfied. If S is a finite set, then for P in a residual subset of $\mathbf{P}$ and all i,*

$$E^P\left[\left|\tilde{\sigma}^P_{it} - E^P\left[\theta \middle| \mathscr{I}^{\sigma^P}\right]\right|\right] \longrightarrow 0,$$

where E^P and σ^P denote the expectation operator and equilibrium strategy profile given P, respectively, and $\mathscr{I}^{\sigma^P}$ is the σ-algebra generated by the union of $\mathscr{I}^{\sigma^P}_i$ over all $i \in N$.

The theorem states that for a generic set of probability distributions P, the agents asymptotically play as if they all had the information captured by the σ-algebra $\mathscr{I}^{\sigma^P}$. Note that $\mathscr{I}^{\sigma^P}$ captures the aggregate information that is collectively available to the agents at the end of the game. Therefore, $E\left[\theta | \mathscr{I}^{\sigma^{\mathscr{P}}}\right]$ is the optimal action given all the signals that the agents receive through the course of the game. To formalize this idea of optimality, one could consider an alternative setting in which a coordinator asks the agents to play according to a strategy profile that maximizes the sum of their expected payoffs. The asymptotically optimal action profile is then the agents' limit action profile when they follow the coordinator's prescription. Theorem 4.2 shows that the agents' equilibrium actions converge to this asymptotically optimal action in the L^1 sense.

An important special case of Theorem 4.2 is obtained by letting $\lambda = 0$. In this case, the agents only attempt to form the best possible estimate of the state given the information available to them. Their equilibrium actions are in turn simply their estimate of θ conditional on their information. The agents' problem then becomes an instance of social learning. Theorem 4.2 states that the agents asymptotically learn to estimate the state as if they had access to all the available information. In this sense, Theorem 4.2 extends and complements some of the earlier optimality results in the Bayesian social learning literature. In particular, it extends Theorem 4 of Mueller-Frank (2013) to the case where the join of the agents' partitions of the state space at the end

[34] Given a topological space X, a subset A of X is a *meager set* if it can be expressed as the union of countably many nowhere dense subsets of X. The complement of a meager set is called a *residual set.*

of the game is infinite dimensional. It also extends Theorem 2.1 to the case where the agents communicate their conditional estimates of the state, rather than their entire beliefs.

The following numerical example illustrates the evolution of the agents' actions over time and their convergence to the optimal action given a setting where the state and the private signals are normally distributed.

Example 4.4. There are $n = 6$ agents over a fixed strongly connected social network playing the game with payoffs given by (4.1) with $\lambda = 2/3$. We consider two network topologies: a directed ring network depicted in Figure 4.1(a) and a star network depicted in Figure 4.1(b). The common prior over θ is given by the standard normal distribution $\mathcal{N}(0,1)$. The signal spaces are given by $S_1 = S_2 = \mathbb{R}$, and the signaling functions are given by $\pi_1(\theta) = \mathcal{N}(\theta, 1)$ and $\pi_t(h_t) = \delta_0$ for $t \geq 2$, where δ_0 is the degenerate probability distribution with unit mass over zero. That is, the agents receive only one informative signal. The evolution of the agents' actions over time is depicted in Figure 4.2 for two realizations of the path of play with $\theta = 0$. The dashed line represents the efficient action given the agents' private signals—which in the context of this example is equal to the average of the private signals. Since agents' prior is $\mathcal{N}(0,1)$, they start by choosing the zero action. At time $t = 1$, agents each receive a private signal and choose the action that is equal to her time 1 private signal. Yet, as time passes, the agents' actions converge to the efficient action. Moreover, over both of the networks, convergence is complete after a number of time periods equal to the diameter of the graph.[35] In this example, although the agents' signal spaces are not finite, convergence to the efficient action is achieved. [36]

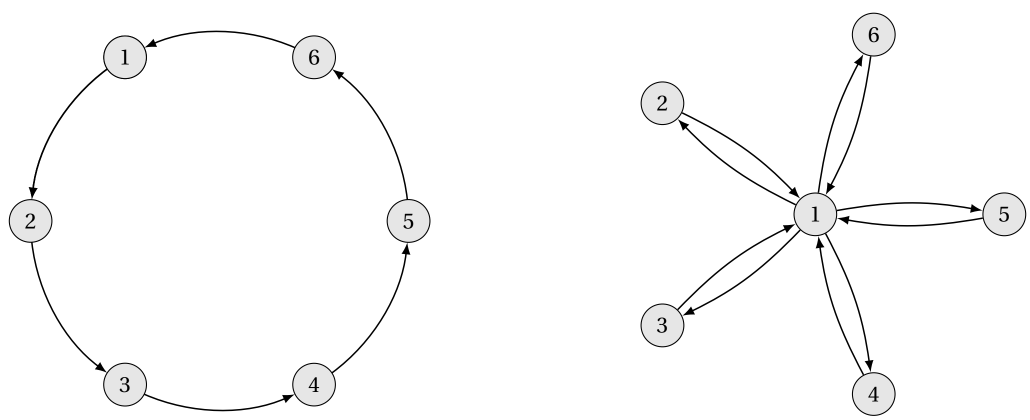

(a) The ring social network (b) The star social network

Figure 4.1. The ring and star social networks of Example 4.4

[35]The diameter of a directed network is defined as $\max_{i,j} l(i,j)$, where $l(i,j)$ is the length of the shortest directed path starting from i and ending at j.

[36]For a recursive characterization of the agents' equilibrium actions in the Bayesian quadratic network games similar to the one studied in Example 4.4, see the paper by Eksin, Molavi, Ribeiro, and Jadbabaie (2013).

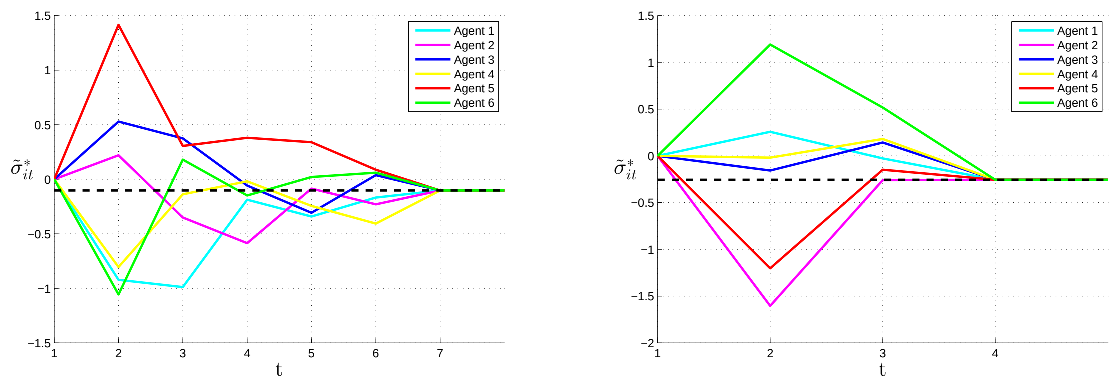

(a) Agents' actions given the ring network

(b) Agents' actions given the star network

Figure 4.2. Evolution of the agents' actions over time in Example 4.4

4.5 Proofs

Proof of Lemma 4.1

First, note that given any strategy profile σ, it is possible to construct a collection of consistent belief systems q by defining $q_i(h_{it})$ to be a regular conditional probability of $\mathbb{P}^\sigma$ given $\mathscr{H}_{it}$, evaluated at some $\omega \in \tilde{h}_{it}^{-1}(h_{it})$. Therefore, we only need to prove that if $\tilde{q}^*_{it}$ is a regular conditional probability of $\mathbb{P}^*$ given $\mathscr{H}_{it}$, then (4.2) is equivalent to condition (a) of Definition 4.3. Given the strategy profile σ and collection of belief systems q, let $\tilde{v}_{it}(\sigma_i, \sigma_{-i}; q_i)$ be the real-valued random variable defined as

$$\tilde{v}_{it}(\sigma_i, \sigma_{-i}; q_i) = v_{it}\left(\sigma_i\left(\tilde{h}_{it}\right), \sigma_{-i}; q_i\left(\tilde{h}_{it}\right)\right).$$

Condition (a) of the equilibrium definition can be expressed in terms of $\tilde{v}_{it}$ as follows: for any strategy σ_i and with $\mathbb{P}^*$-probability one,

$$\tilde{v}_{it}(\sigma^*_i, \sigma^*_{-i}; q^*_i) \geq \tilde{v}_{it}(\sigma_i, \sigma^*_{-i}; q^*_i).$$

On the other hand, it is easy to verify that if $\tilde{q}^*_{it}$ is a regular conditional probability of $\mathbb{P}^*$ given $\mathscr{H}_{it}$, then $\tilde{v}_{it}(\sigma_i, \sigma_{-i}; q^*_i)$ is a version of $\mathbb{E}^*\left[u_i(\tilde{\sigma}_t, \theta) | \mathscr{H}_{it}\right]$. □

Proof of Proposition 4.1

Before presenting the proof, we first introduce some notation and prove a technical lemma. Let $1 \leq p < \infty$, and let $(X, \mathscr{X}, P)$ be a measure space. Consider the set of all L^p-integrable random variables, that is, the set of all measurable functions $f : X \to \mathbb{R}$ such that

$$\|f\|_p = \left(\int_X |f|^p \, dP\right)^{\frac{1}{p}} < \infty.$$

The set of such functions, together with the function $\|\cdot\|_p$, defines a seminormed vector space, denoted by $\mathscr{L}^p(X, P)$. This can be made into a normed vector space in a standard way; one simply takes the quotient space with respect to the kernel of $\|\cdot\|_p$:

$$\ker(\|\cdot\|_p) = \left\{f : f = 0 \quad P\text{-almost everywhere}\right\}.$$

In the quotient space, two functions f and g are identified if $f = g$ almost everywhere. The resulting normed vector space is, by definition,

$$L^p(X,P) = \mathscr{L}^p(X,P)/\ker(\|\cdot\|_p).$$

Further let $\mathbf{L}^p(X,P) = (L^p(X,P))^n$ denote the Banach space with the norm $\|\cdot\|_p$ defined as

$$\|f\|_p = \sum_{i=1}^{n} \|f_i\|_p.$$

By Riesz-Fischer theorem, $L^p(X,P)$ and $\mathbf{L}^p(X,P)$, together with the corresponding $\|\cdot\|_p$, are Banach spaces. In our notation, we have suppressed the dependence of $\|\cdot\|_p$ on the underlying probability measure. Whenever we $\|\cdot\|_p$ use without reference to a specific probability measure, the correct measure will be obvious from the context.

Lemma 4.2. *Let $(X,\mathscr{X},P)$ be a measure space, and let E be the expectation operator corresponding to P. Also let θ be a square integrable random variable, and let $\mathscr{X}_i \subseteq \mathscr{X}$ be a sub σ-algebra for any $i \in N$. Then, there exists a unique $f \in \mathbf{L}^2(X,P)$ such that*

$$f_i = (1-\lambda)E[\theta|\mathscr{X}_i] + \frac{\lambda}{n-1}\sum_{j\neq i} E[f_j|\mathscr{X}_i],$$

for all $i \in N$.

Proof. Let $T : \mathbf{L}^2(X,P) \to \mathbf{L}^2(X,P)$ be the mapping defined as

$$T_i(f) = (1-\lambda)E[\theta|\mathscr{X}_i] + \frac{\lambda}{n-1}\sum_{j\neq i} E[f_j|\mathscr{X}_i],$$

where we are using the fact that θ is square integrable. Note that

$$\begin{aligned}
\|T_i(f) - T_i(g)\|_2 &= \frac{\lambda}{n-1}\left\|\sum_{j\neq i} E[f_j - g_j|\mathscr{X}_i]\right\|_2 \\
&\leq \frac{\lambda}{n-1}\sum_{j\neq i}\|E[f_j - g_j|\mathscr{X}_i]\|_2 \\
&\leq \frac{\lambda}{n-1}\sum_{j\neq i}\|f_j - g_j\|_2,
\end{aligned}$$

where the first inequality is the triangle inequality and the second one is a consequence of the fact that conditional expectation is a contraction with respect to the norm $\|\cdot\|_2$. Therefore,

$$\begin{aligned}
\|T(f)-T(g)\|_2 &= \sum_{i=1}^{n} \|T_i(f)-T_i(g)\|_2 \\
&\leq \frac{\lambda}{n-1}\sum_{i=1}^{n}\sum_{j\neq i}\|f_j-g_j\|_2 \\
&= \lambda\|f-g\|_2.
\end{aligned}$$

Thus, T is a contraction mapping with the Lipschitz constant $\lambda<1$. Hence, by the Banach fixed-point theorem, T has a unique fixed point $f\in\mathbf{L}^2(X,P)$. $\square$

Proof of Propositions 4.1 The proof is constructive. We start at $t=1$ and inductively construct the functions $\sigma^*_{it}: H_{it}\to A_i$. The equilibrium strategies are then defined as $\sigma^*_i(H_{it})=\sigma^*_{it}(H_{it})$ for all i and t.

Let P_1 be the probability distribution over $(H_1,\mathscr{H}_1)$ induced by ν, ψ, and π_1, and let E_1 be the corresponding expectation operator. Consider some strategy profile σ. The marginal of $\mathbb{P}^\sigma$ over $(H_1,\mathscr{H}_1)$ is equal to P_1. Furthermore, θ is measurable with respect to $\mathscr{H}_1$ and $\tilde{\sigma}_{i1}$ is measurable with respect to $\mathscr{H}_{i1}\subseteq\mathscr{H}_1$ for all i. Therefore, for any σ,

$$(1-\lambda)\mathbb{E}^\sigma[\theta|\mathscr{H}_{i1}]+\frac{\lambda}{n-1}\sum_{j\neq i}\mathbb{E}^\sigma[\tilde{\sigma}_{i1}|\mathscr{H}_{i1}]=(1-\lambda)E_1[\theta|\mathscr{H}_{i1}]+\frac{\lambda}{n-1}\sum_{j\neq i}E_1[\tilde{\sigma}_{i1}|\mathscr{H}_{i1}].$$

In particular, by Corollary 4.1, for any square integrable equilibrium strategy profile σ^*,

$$\begin{aligned}
\tilde{\sigma}^*_{i1} &= (1-\lambda)\mathbb{E}^*[\theta|\mathscr{H}_{i1}]+\frac{\lambda}{n-1}\sum_{j\neq i}\mathbb{E}^*\left[\tilde{\sigma}^*_{i1}|\mathscr{H}_{i1}\right] \\
&= (1-\lambda)E_1[\theta|\mathscr{H}_{i1}]+\frac{\lambda}{n-1}\sum_{j\neq i}E_1\left[\tilde{\sigma}^*_{j1}|\mathscr{H}_{i1}\right].
\end{aligned}$$

By Assumption 4.1 and Lemma 4.2, the above system of equations has a unique fixed point in $\mathbf{L}^2(H_1,P_1)$. Consequently, (i) there exists a fixed point $\tilde{\sigma}^*_1=(\tilde{\sigma}^*_{11},\dots,\tilde{\sigma}^*_{n1})$ such that $\tilde{\sigma}^*_{i1}\in\mathscr{L}^2(H_1,P_1)$; and (ii) for any other square integrable equilibrium strategy profile $\sigma^\dagger$, we have that $\tilde{\sigma}^*_{i1}=\tilde{\sigma}^\dagger_{i1}$ with P_1-probability one. Furthermore, by construction $\tilde{\sigma}^*_{i1}$ is $\mathscr{H}_{i1}$-measurable for all i. This implies that there exists some function $\sigma^*_{i1}: H_{i1}\to A_i$ such that $\sigma^*_{i1}(\tilde{h}_{i1})=\tilde{\sigma}^*_{i1}$.

Next, let P_2 be the probability distribution over $(H_2,\mathscr{H}_2)$ induced by ν, ψ, π_1 and π_2, and the time 1 profile $(\sigma^*_{11},\dots,\sigma^*_{n1})$ constructed earlier. Recall that, for any two square integrable equilibrium strategy profiles, $\tilde{\sigma}^*_{i1}=\tilde{\sigma}^\dagger_{i1}$ with P_1-probability one. Thus, all such strategy profiles induce the same probability distribution over $(H_2,\mathscr{H}_2)$. We can thus repeat the same argument to conclude that there exist functions $\sigma^*_{i2}: H_{i2}\to A_i$ such that $\sigma^*_{i2}(\tilde{h}_{i2})=\tilde{\sigma}^*_{i2}\in\mathscr{L}^2(H_2,P_2)$ and

$$\begin{aligned}
\tilde{\sigma}^*_{i2} &= (1-\lambda)\mathbb{E}^*[\theta|\mathscr{H}_{i2}]+\frac{\lambda}{n-1}\sum_{j\neq i}\mathbb{E}^*\left[\tilde{\sigma}^*_{i2}|\mathscr{H}_{i2}\right] \\
&= (1-\lambda)E_2[\theta|\mathscr{H}_{i2}]+\frac{\lambda}{n-1}\sum_{j\neq i}E_2\left[\tilde{\sigma}^*_{i2}|\mathscr{H}_{i2}\right],
\end{aligned}$$

for all i. Moreover, for any other square integrable strategy profile $\sigma^\dagger$, we have that $\tilde{\sigma}^*_{i2} = \tilde{\sigma}^\dagger_{i2}$ with P_2-probability one. We can proceed inductively to complete the proof. □

Proof of Proposition 4.2

Consider the following system of equations:

$$\tilde{\zeta}_i = (1-\lambda)\mathbb{E}^*[\theta|\mathscr{H}_i] + \frac{\lambda}{n-1}\sum_{j\neq i}\mathbb{E}^*[\tilde{\zeta}_j|\mathscr{H}_i].$$

By Lemma 4.2, the above set of equations has some solution $(\tilde{\zeta}^*_i)_{i\in N}$, where $\tilde{\zeta}^*_i \in \mathscr{L}^2(\Omega, \mathbb{P}^*)$. Moreover, by construction $\tilde{\zeta}^*_i$ is $\mathscr{H}_i$-measurable. We prove the lemma by showing that $\tilde{\sigma}^*_{it} \to \tilde{\zeta}^*_i$ in the L^2 sense as t goes to infinity. By Corollary 4.1,

$$\tilde{\sigma}^*_{it} - \tilde{\zeta}^*_i = (1-\lambda)\left(\mathbb{E}^*[\theta|\mathscr{H}_{it}] - \mathbb{E}^*[\theta|\mathscr{H}_i]\right) + \frac{\lambda}{n-1}\sum_{j\neq i}\left(\mathbb{E}^*[\tilde{\sigma}^*_{jt}|\mathscr{H}_{it}] - \mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_i]\right).$$

Using the triangle inequality, we can conclude that

$$\begin{aligned}\|\tilde{\sigma}^*_{it} - \tilde{\zeta}^*_i\|_2 \leq{}& (1-\lambda)\left\|\mathbb{E}^*[\theta|\mathscr{H}_{it}] - \mathbb{E}^*[\theta|\mathscr{H}_i]\right\|_2\\ &+ \frac{\lambda}{n-1}\sum_{j\neq i}\left\|\mathbb{E}^*[\tilde{\sigma}^*_{jt} - \tilde{\zeta}^*_j|\mathscr{H}_{it}]\right\|_2\\ &+ \frac{\lambda}{n-1}\sum_{j\neq i}\left\|\mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_{it}] - \mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_i]\right\|_2.\end{aligned} \tag{4.4}$$

Since the conditional expectation is a contraction,

$$\left\|\mathbb{E}^*[\tilde{\sigma}^*_{jt} - \tilde{\zeta}^*_j|\mathscr{H}_{it}]\right\|_2 \leq \left\|\tilde{\sigma}^*_{jt} - \tilde{\zeta}^*_j\right\|_2.$$

Summing (4.4) over i and using the above inequality imply

$$\begin{aligned}\sum_{i=1}^n \|\tilde{\sigma}^*_{it} - \tilde{\zeta}^*_i\|_2 \leq{}& (1-\lambda)\sum_{i=1}^n\left\|\mathbb{E}^*[\theta|\mathscr{H}_{it}] - \mathbb{E}^*[\theta|\mathscr{H}_i]\right\|_2\\ &+ \lambda\sum_{i=1}^n\|\tilde{\sigma}^*_{it} - \tilde{\zeta}^*_i\|_2\\ &+ \frac{\lambda}{n-1}\sum_{i=1}^n\sum_{j\neq i}\left\|\mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_{it}] - \mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_i]\right\|_2.\end{aligned}$$

Hence,

$$\begin{aligned}\sum_{i=1}^n \|\tilde{\sigma}^*_{it} - \tilde{\zeta}^*_i\|_2 \leq{}& \sum_{i=1}^n\left\|\mathbb{E}^*[\theta|\mathscr{H}_{it}] - \mathbb{E}^*[\theta|\mathscr{H}_i]\right\|_2\\ &+ \frac{\lambda}{1-\lambda}\frac{1}{n-1}\sum_{i=1}^n\sum_{j\neq i}\left\|\mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_{it}] - \mathbb{E}^*[\tilde{\zeta}^*_j|\mathscr{H}_i]\right\|_2.\end{aligned}$$

It is easy to verify that $\mathbb{E}^*[\theta|\mathcal{H}_{it}]$ is a martingale with respect to the filtration $\mathcal{H}_{it} \uparrow \mathcal{H}_i$. Furthermore,

$$\sup_t \|\mathbb{E}^*[\theta|\mathcal{H}_{it}]\|_2 \leq \|\theta\|_2 < \infty,$$

where the first inequality is a consequence of the fact that conditional expectation is a contraction and the second one is due to Assumption 4.1. Thus, by the L^p convergence theorem, $\mathbb{E}^*[\theta|\mathcal{H}_{it}]$ converges in the L^2 sense to $\mathbb{E}^*[\theta|\mathcal{H}_i]$.[37] That is,

$$\lim_{t\to\infty} \left\|\mathbb{E}^*[\theta|\mathcal{H}_{it}] - \mathbb{E}^*[\theta|\mathcal{H}_i]\right\|_2 = 0.$$

By a similar argument, relying on the fact that $\tilde{\varsigma}_j^*$ is square integrable, for all j,

$$\lim_{t\to\infty} \left\|\mathbb{E}^*[\tilde{\varsigma}_j^*|\mathcal{H}_{it}] - \mathbb{E}^*[\tilde{\varsigma}_j^*|\mathcal{H}_i]\right\|_2 = 0.$$

Therefore,

$$\lim_{t\to\infty} \sum_{i=1}^n \|\tilde{\sigma}_{it}^* - \tilde{\varsigma}_i^*\|_2 = 0.$$

□

Proof of Theorem 4.1

We first prove part (b) of the theorem. Let i, j be a pair of agents such that i observes the actions of j infinitely often ψ-almost surely. Consider the strategy $\sigma_i^\dagger : H_i \to A_i$ defined as follows.

$$\sigma_i^\dagger(h_{it}) = \begin{cases} 0 & \text{if} \quad t = 1, \\ a_{jt-1} & \text{if} \quad [g_{t-1}]_{ji} = 1, \\ \sigma_i(h_{it-1}) & \text{otherwise.} \end{cases}$$

The strategy $\sigma_i^\dagger$ describes the following plan of action: Agent i starts by choosing zero; she plays the same action until observing some action taken by agent j, in which case agent i switches to the observed action and continues choosing it until agent j's action is observed again. We use this strategy to prove the result in three steps. In step one, we show that $\tilde{\sigma}_i^\dagger$ converges in the L^2 sense to $\tilde{\varsigma}_j^*$. That is, if agent i follows strategy $\sigma_i^\dagger$, her actions will asymptotically coincide with those of agent j. In step two, we use this result to show that the limit of agent i's expected payoff from following σ_i^* is not lower than what it would be, had she followed $\sigma_i^\dagger$. In step three, we show that in a strongly connected network the limits of all agents' expected payoffs, and hence, the limits of their actions, are the same.

Step one. Let $\tau < t$. By the triangle inequality,

$$\left\|\tilde{\sigma}_{it}^\dagger - \tilde{\varsigma}_j^*\right\|_2 \leq \left\|\tilde{\sigma}_{it}^\dagger - \tilde{\sigma}_{j\tau}^*\right\|_2 + \left\|\tilde{\sigma}_{j\tau}^* - \tilde{\varsigma}_j^*\right\|_2 \tag{4.5}$$

[37] For a statement and proof of the L^p convergence theorem, see, for instance, Durrett (2010, p. 215).

We first use a truncation argument to bound the first term of (4.5). For $\tau \le r < t$, let B_{rt} be the event defined as

$$B_{rt} = \left\{\omega : [g_r]_{ji} = 1, \quad \text{and} \quad [g_s]_{ji} = 0 \quad \text{for all } r < s < t\right\},$$

and let $D_{\tau t}$ be the event defined as

$$D_{\tau t} = \left\{\omega : [g_r]_{ji} = 0 \quad \text{for all } \tau \le r < t\right\}.$$

B_{rt} is the event that after observing the time r action of agent j, agent i does not observe agent j's action again until after time t. $D_{\tau t}$ is the event that agent i does not observe agent j's actions between time periods τ and t. By definition,

$$B_{\tau t} \cup B_{\tau+1 t} \cdots \cup B_{t-1 t} \cup D_{\tau t} = \Omega.$$

Therefore,[38]

$$\begin{aligned}
\mathbb{E}^*\left[\left(\tilde{\sigma}^{\dagger}_{it} - \tilde{\sigma}^*_{j\tau}\right)^2\right] &= \sum_{r=\tau}^{t-1} \mathbb{E}^*\left[\left(\tilde{\sigma}^{\dagger}_{it} - \tilde{\sigma}^*_{j\tau}\right)^2 \middle| B_{rt}\right] \mathbb{P}^*(B_{rt}) + \mathbb{E}^*\left[\left(\tilde{\sigma}^{\dagger}_{it} - \tilde{\sigma}^*_{j\tau}\right)^2 \middle| D_{\tau t}\right] \mathbb{P}^*(D_{\tau t}) \\
&= \sum_{r=\tau}^{t-1} \mathbb{E}^*\left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right)^2\right] \mathbb{P}^*(B_{rt}) + \mathbb{E}^*\left[\left(\tilde{\sigma}^{\dagger}_{i\tau} - \tilde{\sigma}^*_{j\tau}\right)^2\right] \mathbb{P}^*(D_{\tau t}), \qquad (4.6)
\end{aligned}$$

where in the second equality we are using the fact that $\tilde{\sigma}^*_{j\tau}$ is independent of $D_{\tau t}$ and of any of the events $\{B_{rt}\}_{r\in[\tau,t-1]}$. We have

$$\begin{aligned}
\sum_{r=\tau}^{t-1} \mathbb{E}^*\left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right)^2\right] \mathbb{P}^*(B_{rt}) &\le \max_{r\in[\tau,t-1]} \mathbb{E}^*\left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right)^2\right] \sum_{r=\tau}^{t-1} \mathbb{P}^*(B_{rt}) \\
&\le \max_{r\in[\tau,t-1]} \mathbb{E}^*\left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right)^2\right] \\
&= \max_{r\in[\tau,t-1]} \left\|\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right\|_2^2.
\end{aligned}$$

Since $\{\tilde{\sigma}^*_{jt}\}_{t\in\mathbb{N}}$ is a convergent sequence in the $(L^2, \|\cdot\|_2)$ space, it is also a Cauchy sequence. Therefore, for any $\epsilon > 0$, if τ is sufficiently large, then $\|\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\|_2 \le \frac{\epsilon}{\sqrt{8}}$ for all $r \ge \tau$, implying that

$$\sum_{r=\tau}^{t-1} \mathbb{E}^*\left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right)^2\right] \mathbb{P}^*(B_{rt}) \le \frac{\epsilon^2}{8}. \qquad (4.7)$$

Next, we consider the second term of (4.6). By construction, $\tilde{\sigma}^{\dagger}_{i\tau} \in \{0\} \cup \{\tilde{\sigma}^*_{jr}\}_{r\in[1,\tau-1]}$. Thus,

$$\begin{aligned}
\mathbb{E}^*\left[\left(\tilde{\sigma}^{\dagger}_{i\tau} - \tilde{\sigma}^*_{j\tau}\right)^2\right] &\le \max\left\{\mathbb{E}^*\left[\left(\tilde{\sigma}^*_{j\tau}\right)^2\right], \max_{r\in[1,\tau-1]} \mathbb{E}^*\left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right)^2\right]\right\} \\
&= \max\left\{\left\|\tilde{\sigma}^*_{j\tau}\right\|_2^2, \max_{r\in[1,\tau-1]} \left\|\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right\|_2^2\right\}.
\end{aligned}$$

[38]We provide the proof for the case that $\mathbb{P}^*(B_{rt}) > 0$ for all $\tau \le r < t$ and $\mathbb{P}^*(D_{\tau t}) > 0$ for all t. The proof can be extended to other cases through obvious modifications.

Since $\{\tilde{\sigma}^*_{jt}\}_{t\in\mathbb{N}}$ is a convergent sequence in the $(L^2, \|\cdot\|_2)$ space, it is a also a bounded Cauchy sequence. Therefore, there exists some $M > 0$ such that for all τ,

$$\max\left\{\left\|\tilde{\sigma}^*_{j\tau}\right\|_2^2, \max_{r\in[1,\tau-1]}\left\|\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\right\|_2^2\right\} \leq M.$$

Finally, by Assumption 4.3, $\mathbb{P}^*(D_{\tau t}) \to 0$ as t goes to infinity. Therefore, for any $\epsilon > 0$, if t is sufficiently large, then $\mathbb{P}^*(D_{\tau t}) \leq \frac{\epsilon^2}{8M}$, implying that

$$\mathbb{E}^*\left[\left(\tilde{\sigma}^\dagger_{i\tau} - \tilde{\sigma}^*_{j\tau}\right)^2\right]\mathbb{P}^*(D_{\tau t}) \leq \frac{\epsilon^2}{8}. \tag{4.8}$$

Combining (4.7) and (4.8) with (4.6), we get that, for sufficiently large values of τ and $t > \tau$,

$$\left\|\tilde{\sigma}^\dagger_{it} - \tilde{\sigma}^*_{j\tau}\right\|_2 = \left(\mathbb{E}^*\left[\left(\tilde{\sigma}^\dagger_{it} - \tilde{\sigma}^*_{j\tau}\right)^2\right]\right)^{\frac{1}{2}} \leq \frac{\epsilon}{2}.$$

We next bound the second term of (4.5). By Proposition 4.2, for any arbitrary $\epsilon > 0$, if τ is sufficiently large,

$$\left\|\tilde{\sigma}^*_{j\tau} - \tilde{\varsigma}^*_j\right\|_2 \leq \frac{\epsilon}{2}.$$

Together with (4.5), the last two inequities show that if t is sufficiently large, then

$$\left\|\tilde{\sigma}^\dagger_{it} - \tilde{\varsigma}^*_j\right\|_2 \leq \epsilon.$$

Since $\epsilon > 0$ was arbitrary, we can conclude that $\tilde{\sigma}^\dagger_{it}$ converges to $\tilde{\varsigma}^*_j$ in the L^2 sense as t goes to infinity.

Step two. We first prove that $\mathbb{E}^*\left[u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta)\right]$ converges to $\mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta)\right]$. By the reverse triangle inequality,

$$\begin{aligned}\left|\left(\mathbb{E}^*\left[\left(\tilde{\sigma}^*_{it} - \theta\right)^2\right]\right)^{\frac{1}{2}} - \left(\mathbb{E}^*\left[\left(\tilde{\varsigma}^*_i - \theta\right)^2\right]\right)^{\frac{1}{2}}\right| &= \left|\left\|\tilde{\sigma}^*_{it} - \theta\right\|_2 - \left\|\tilde{\varsigma}^*_i - \theta\right\|_2\right| \\ &\leq \left\|\left(\tilde{\sigma}^*_{it} - \theta\right) - \left(\tilde{\varsigma}^*_i - \theta\right)\right\|_2.\end{aligned}$$

By Proposition 4.2, as t goes to infinity, $\tilde{\sigma}^*_{it}$ converges to $\tilde{\varsigma}^*_i$ in the L^2 sense. Therefore,

$$\mathbb{E}^*\left[\left(\tilde{\sigma}^*_{it} - \theta\right)^2\right] \longrightarrow \mathbb{E}^*\left[\left(\tilde{\varsigma}^*_i - \theta\right)^2\right] \qquad \text{as} \qquad t \to \infty.$$

A similar argument shows that

$$\mathbb{E}^*\left[\left(\tilde{\sigma}^*_{it} - \frac{1}{n-1}\sum_{j\neq i}\tilde{\sigma}^*_{jt}\right)^2\right] \longrightarrow \mathbb{E}^*\left[\left(\tilde{\varsigma}^*_i - \frac{1}{n-1}\sum_{j\neq i}\tilde{\varsigma}^*_j\right)^2\right] \qquad \text{as} \qquad t \to \infty,$$

thus implying that

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta)\right] \longrightarrow \mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta)\right] \qquad \text{as} \qquad t \to \infty. \tag{4.9}$$

We can use the result of the step one to show, in a similar manner, that

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^{\dagger}_{it},\tilde{\sigma}^*_{-it},\theta)\right] \longrightarrow \mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_j,\tilde{\varsigma}^*_{-i},\theta)\right] \qquad \text{as} \qquad t\to\infty. \tag{4.10}$$

On the other hand, since σ^* is an equilibrium strategy profile, by Lemma 4.1, for all t,

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^*_{it},\tilde{\sigma}^*_{-it},\theta)\middle|\mathcal{H}_{it}\right] \geq \mathbb{E}^*\left[u_i\left(\tilde{\sigma}^{\dagger}_{it},\tilde{\sigma}^*_{-it},\theta\right)\middle|\mathcal{H}_{it}\right],$$

with $\mathbb{P}^*$-probability one. Therefore,

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^*_{it},\tilde{\sigma}^*_{-it},\theta)\right] \geq \mathbb{E}^*\left[u_i\left(\tilde{\sigma}^{\dagger}_{it},\tilde{\sigma}^*_{-it},\theta\right)\right].$$

Thus, taking limits of both sides of the above inequality as $t\to\infty$ and using (4.9) and (4.10),

$$\mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_i,\tilde{\varsigma}^*_{-i},\theta)\right] \geq \mathbb{E}^*\left[u_i\left(\tilde{\varsigma}^*_j,\tilde{\varsigma}^*_{-i},\theta\right)\right]. \tag{4.11}$$

Step three. By Assumption 4.3, there exists a sequence of agents $i_0, i_1, i_2, \dots, i_n$ starting and ending with the same agent that includes each agent other than agent i_0 exactly once, and such that, for all k, agent i_k observes i_{k+1} infinitely often ψ-almost surely. For any k, thus by the result of step two,

$$\mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_{i_k},\tilde{\varsigma}^*_{-i_k},\theta)\right] \geq \mathbb{E}^*\left[u_i\left(\tilde{\varsigma}^*_{i_{k+1}},\tilde{\varsigma}^*_{-i_k},\theta\right)\right]. \tag{4.12}$$

Summing over k and reindexing the right-hand side sum imply

$$\sum_{k=0}^{n-1}\mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_{i_k},\tilde{\varsigma}^*_{-i_k},\theta)\right] \geq \sum_{k=1}^{n}\mathbb{E}^*\left[u_i\left(\tilde{\varsigma}^*_{i_k},\tilde{\varsigma}^*_{-i_{k-1}},\theta\right)\right].$$

Expanding both sides of the inequality, all terms except for one cancel, resulting in

$$\sum_{k=0}^{n-1}\mathbb{E}^*\left[\tilde{\varsigma}^*_{i_k}\sum_{j\neq k}\tilde{\varsigma}^*_{i_j}\right] \geq \sum_{k=1}^{n}\mathbb{E}^*\left[\tilde{\varsigma}^*_{i_k}\sum_{j\neq k-1}\tilde{\varsigma}^*_{i_j}\right].$$

Further simplification implies that

$$\sum_{k=1}^{n}\mathbb{E}^*\left[\tilde{\varsigma}^*_{i_k}\tilde{\varsigma}^*_{i_{k-1}}\right] \geq \sum_{k=1}^{n}\mathbb{E}^*\left[\left(\tilde{\varsigma}^*_{i_k}\right)^2\right]. \tag{4.13}$$

On the other hand, $\sum_{k=1}^n \mathbb{E}^*[(\tilde{\varsigma}^*_{i_k}-\tilde{\varsigma}^*_{i_{k-1}})^2] \geq 0$ with equality if and only if $\tilde{\varsigma}^*_{i_k}=\tilde{\varsigma}^*_{i_{k-1}}$ for all k with $\mathbb{P}^*$-probability one. Thus, using the fact that $\sum_{k=1}^n \mathbb{E}^*[(\tilde{\varsigma}^*_{i_k})^2] = \sum_{k=1}^n \mathbb{E}^*[(\tilde{\varsigma}^*_{i_{k-1}})^2]$, we can conclude that

$$\sum_{k=1}^{n}\mathbb{E}^*\left[\left(\tilde{\varsigma}^*_{i_k}\right)^2\right] \geq \sum_{k=1}^{n}\mathbb{E}^*\left[\tilde{\varsigma}^*_{i_k}\tilde{\varsigma}^*_{i_{k-1}}\right], \tag{4.14}$$

with equality if and only if $\tilde{\varsigma}^*_{i_k}=\tilde{\varsigma}^*_{i_{k-1}}$ for all k, $\mathbb{P}^*$-almost surely; equation (4.13) implies that (4.14) indeed holds with equality. Thus, for all i and j and with $\mathbb{P}^*$-probability one,

$$\tilde{\varsigma}^*_i = \tilde{\varsigma}^*_j.$$

Together with Proposition 4.2, this completes the proof of part (b) of the theorem.

We now prove part (a). Since $\tilde{\sigma}^*_{it}$ converges to $\tilde{\varsigma}^*_i$ in the L^2 sense, it also converges in probability. Therefore, by the continuous mapping theorem, $u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta)$ converges to $u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta)$ in probability. Together with (4.9), this implies that $u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta)$ converges to $u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta)$ in the L^1 sense.[39] This is true for any two agents. Moreover, by part (b) of the theorem, $u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta) = u_j(\tilde{\varsigma}^*_j, \tilde{\varsigma}^*_{-j}, \theta)$ for all $i, j \in N$. This proves that, as t goes to infinity,

$$\left\| u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta) - u_j(\tilde{\sigma}^*_{jt}, \tilde{\sigma}^*_{-jt}, \theta) \right\|_1 \longrightarrow 0,$$

for any i, j.

We next prove part (c). By part (b) of the theorem, $\sum_{j \neq i} \tilde{\sigma}^*_{it} - \tilde{\sigma}^*_{jt}$ goes to zero for all i in the L^2 sense. Therefore, by Corollary 4.1, we can conclude that $\tilde{\sigma}^*_{it} - \mathbb{E}^*[\theta|\mathscr{H}_{it}]$ goes to zero in the L^2 sense. On the other hand, since $\mathscr{H}_{it} \uparrow \mathscr{H}_i$, we have that $\mathbb{E}^*[\theta|\mathscr{H}_{it}]$ converges to $\mathbb{E}^*[\theta|\mathscr{H}_i]$ in the L^1 sense. Therefore, $\tilde{\sigma}^*_{it}$ converges to $\mathbb{E}^*[\theta|\mathscr{H}_i]$ in the L^1 sense. Another use of the result of part (b) completes the proof. □

Proof of Theorem 4.2

Before proving the theorem, we first prove a technical lemma.

Lemma 4.3. *Let $(X, \mathscr{B})$ be a measurable space, and let $(\mathbf{P}, d)$ be the metric space where $\mathbf{P}$ is the collection of all probability measures on $(X, \mathscr{B})$ and d is the total variation distance. Let $\mathscr{F}_1$ and $\mathscr{F}_2$ be two arbitrary sub σ-algebras of $\mathscr{B}$, let $\mathscr{F}$ be the σ-algebra generated by the union of $\mathscr{F}_1$ and $\mathscr{F}_2$, and let f be an arbitrary bounded random variable. The set*

$$\mathbf{Q} = \left\{ P \in \mathbf{P} : E_P[f|\mathscr{F}_1] = E_P[f|\mathscr{F}_2] \neq E_P[f|\mathscr{F}] \right\},$$

is nowhere dense in the metric space $(\mathbf{P}, d)$.

Proof. To prove the lemma, we use Dynkin's π-λ theorem. Let us first construct the appropriate λ and π-systems. For any $P \in \mathbf{P}$, define

$$\Lambda_P = \left\{ B \in \mathscr{B} : \int_B f\, dP = \int_B E_P[f|\mathscr{F}_1]\, dP = \int_B E_P[f|\mathscr{F}_2]\, dP \right\}.$$

We first verify that for any $P \in \mathbf{P}$, the set Λ_P is a λ-system of subsets of X. (i) By the law of total expectation $X \in \Lambda_P$. (ii) Let B^c denote the complement of B in X. If $B \in \Lambda_P$, then

$$\int_{B^c} f\, dP = \int_X f\, dP - \int_B f\, dP = \int_X E_P[f|\mathscr{F}_1]\, dP - \int_B E_P[f|\mathscr{F}_1]\, dP = \int_{B^c} E_P[f|\mathscr{F}_1]\, dP.$$

We also have a similar equality for $\mathscr{F}_2$. Therefore, $B^c \in \Lambda_P$. (iii) If $B_1, B_2, \ldots$ is a sequence of subsets of X in Λ_P such that $B_i \cap B_j = \emptyset$ for all $i \neq j$, then by the countable additivity of the integral,

$$\int_{\cup_{i=1}^\infty B} f\, dP = \sum_{i=1}^\infty \int_{B_i} f\, dP = \sum_{i=1}^\infty \int_{B_i} E_P[f|\mathscr{F}_1]\, dP = \int_{\cup_{i=1}^\infty B} E_P[f|\mathscr{F}_1]\, dP.$$

[39] This is due to a variant of the dominated convergence theorem that can be found, among other places, as Theorem 5.5.2. in the book of Durrett (2010).

We also have a similar equality for $\mathscr{F}_2$. Therefore, $\cup_{i=1}^{\infty} B_i \in \Lambda_P$. This proves that Λ_P is a λ-system. Consider next the set Π defined as

$$\Pi = \{A_1 \cap A_2 : A_1 \in \mathscr{F}_1, A_2 \in \mathscr{F}_2\}.$$

$\mathscr{F}_1$ and $\mathscr{F}_2$ are σ-algebras; thus, Π is nonempty and closed under intersections. This proves that Π is indeed a π-system of subsets of X. It is also easy to verify that $\sigma(\Pi) = \sigma(\mathscr{F}_1 \cup \mathscr{F}_2) = \mathscr{F}$.

Define the set $\mathbf{R} \supseteq \mathbf{Q}$ as

$$\mathbf{R} = \left\{P \in \mathbf{P} : E_P[f|\mathscr{F}_1] = E_P[f|\mathscr{F}_2]\right\}.$$

We consider the following two cases: If $\mathbf{R}$ is nowhere dense in $\mathbf{P}$, then $\mathbf{Q}$ is nowhere dense in $\mathbf{P}$, and we have the desired result. If, on the other hand, $\mathbf{R}$ is not nowhere dense in $\mathbf{P}$, then it must be somewhere dense in it. Let $\mathscr{U}$ be the collection of all open subsets u of $\mathbf{P}$, such that there exists no nonempty open set v contained in u such that v and $\mathbf{R}$ are disjoint. We prove that $\mathbf{Q}$ is nowhere dense in $\mathbf{R}$ by showing that any such u contains an open subset that is disjoint from $\mathbf{Q}$. Let u be an arbitrary set in $\mathscr{U}$, and let b_ϵ be an open ball of radius ϵ in the interior of u. In what follows, we first show that for every $Q \in b_\epsilon$, we have $\Pi \subseteq \Lambda_Q$. Let $A_1 \in \mathscr{F}_1$ and $A_2 \in \mathscr{F}_2$ be arbitrary sets with $C = A_1 \cap A_2$. Since $A_1 \in \mathscr{F}_1$, by the definition of conditional expectation, for all $Q \in b_\epsilon$,

$$\int_{A_1} f\,dQ = \int_{A_1} E_Q[f|\mathscr{F}_1]\,dQ.$$

Therefore,

$$\int_{A_1\setminus C} f\,dQ + \int_C f\,dQ = \int_{A_1\setminus C} E_Q[f|\mathscr{F}_1]\,dQ + \int_C E_Q[f|\mathscr{F}_1]\,dQ. \tag{4.15}$$

On the other hand, since $\mathbf{R}$ is dense in b_ϵ, for any $Q \in b_\epsilon$, there exists a sequence $\{Q_k\}_{k\in\mathbb{N}}$ such that $Q_k \in b_\epsilon \cap \mathbf{R}$ for all k, and Q_k converges in the total variation distance to Q. Therefore, $E_{Q_k}[f|\mathscr{F}_1]$ converges in Q-probability to $E_Q[f|\mathscr{F}_1]$.[40] Therefore, since f is bounded and Q_k converges in total variation distance to Q,

$$\int_{A_2} E_{Q_k}[f|\mathscr{F}_1]\,dQ_k \longrightarrow \int_{A_2} E_Q[f|\mathscr{F}_1]\,dQ, \tag{4.16}$$

and

$$\int_{A_2} f\,dQ_k \longrightarrow \int_{A_2} f\,dQ. \tag{4.17}$$

Moreover, for all k,

$$\int_{A_2} f\,dQ_k = \int_{A_2} E_{Q_k}[f|\mathscr{F}_2]\,dQ_k = \int_{A_2} E_{Q_k}[f|\mathscr{F}_1]\,dQ_k, \tag{4.18}$$

where the first equality is by the definition of conditional expectation and the assumption that $A_2 \in \mathscr{F}_2$, and the second equality is a consequence of the fact that $Q_k \in \mathbf{R}$. Equations (4.16)–(4.18) imply that

$$\int_{A_2} f\,dQ = \int_{A_2} E_Q[f|\mathscr{F}_1]\,dQ.$$

[40]This follows a result of Landers and Rogge (1976) (cf. Theorem 3.3. of Crimaldi and Pratelli (2005)).

And hence,

$$\int_{A_2\setminus C} f\,dQ + \int_C f\,dQ = \int_{A_2\setminus C} E_Q[f|\mathscr{F}_1]\,dQ + \int_C E_Q[f|\mathscr{F}_1]\,dQ. \tag{4.19}$$

We use (4.15) and (4.19) to conclude that $\int_C f\,dQ = \int_C E_Q[f|\mathscr{F}_1]\,dQ$ for all $Q \in b_\epsilon$. Pick some arbitrary $Q \in b_\epsilon$. If $Q(A_1) = 0$ or $Q(A_1) = 1$, by boundedness of f we are done. If $0 < Q(A_1) < 1$, for any $\delta \in (0,1)$ construct the measure $\widehat{Q}_\delta$ over $(X, \mathscr{B})$ as follows: for any $B \in \mathscr{B}$,

$$\widehat{Q}_\delta(B) = (1 + \delta Q(A_1^c))Q(B \cap A_1) + (1 - \delta Q(A_1))Q(B \cap A_1^c).$$

It is easy to verify that $\widehat{Q}_\delta$ is indeed a probability measure. We next show that $E_{\widehat{Q}_\delta}[f|\mathscr{F}_1] = E_Q[f|\mathscr{F}_1]$. Let $B \in \mathscr{F}_1$ be arbitrary.

$$\begin{aligned}
\int_B f\,d\widehat{Q}_\delta &= \int_{B\cap A_1} f\,d\widehat{Q}_\delta + \int_{B\cap A_1^c} f\,d\widehat{Q}_\delta \\
&= (1 + \delta Q(A_1^c)) \int_{B\cap A_1} f\,dQ + (1 - \delta Q(A_1)) \int_{B\cap A_1^c} f\,dQ \\
&= (1 + \delta Q(A_1^c)) \int_{B\cap A_1} E_Q[f|\mathscr{F}_1]\,dQ + (1 - \delta Q(A_1)) \int_{B\cap A_1^c} E_Q[f|\mathscr{F}_1]\,dQ \\
&= \int_{B\cap A_1} E_Q[f|\mathscr{F}_1]\,d\widehat{Q}_\delta + \int_{B\cap A_1^c} E_Q[f|\mathscr{F}_1]\,d\widehat{Q}_\delta \\
&= \int_B E_Q[f|\mathscr{F}_1]\,d\widehat{Q}_\delta,
\end{aligned} \tag{4.20}$$

where the third equality follows from the assumption that $E_Q[f|\mathscr{F}_1]$ is a conditional expectation of f given $\mathscr{F}_1$ and the fact that $B \cap A_1 \in \mathscr{F}_1$ and $B \cap A_1^c \in \mathscr{F}_1$. Since $E_Q[f|\mathscr{F}_1]$ is $\mathscr{F}_1$-measurable, equation (4.20) proves that $E_Q[f|\mathscr{F}_1]$ is a version of $E_{\widehat{Q}_\delta}[f|\mathscr{F}_1]$. Let $B_1 = A_1 \setminus C$ and $B_2 = A_2 \setminus C$. Equations (4.15) and (4.19) imply that

$$\int_{B_1} \left[f - E_Q[f|\mathscr{F}_1]\right] dQ = \int_{B_2} \left[f - E_Q[f|\mathscr{F}_1]\right] dQ. \tag{4.21}$$

Since $B_1 \cap A_1 = B_1$,

$$\int_{B_1} \left[f - E_{\widehat{Q}_\delta}[f|\mathscr{F}_1]\right] d\widehat{Q}_\delta = (1 + \delta Q(A_1^c)) \int_{B_1} \left[f - E_Q[f|\mathscr{F}_1]\right] dQ. \tag{4.22}$$

Likewise, since $B_2 \cap A_1^c = B_2$,

$$\int_{B_2} \left[f - E_{\widehat{Q}_\delta}[f|\mathscr{F}_1]\right] d\widehat{Q}_\delta = (1 - \delta Q(A_1)) \int_{B_2} \left[f - E_Q[f|\mathscr{F}_1]\right] dQ. \tag{4.23}$$

On the other hand, if δ is sufficiently small, $\widehat{Q}_\delta \in b_\epsilon$. Therefore, by (4.15) and (4.19),

$$\int_{B_1} \left[f - E_{\widehat{Q}_\delta}[f|\mathscr{F}_1]\right] d\widehat{Q}_\delta = \int_{B_2} \left[f - E_{\widehat{Q}_\delta}[f|\mathscr{F}_1]\right] d\widehat{Q}_\delta. \tag{4.24}$$

Equations (4.21)–(4.24) imply that

$$\int_{B_1} [f - E_Q[f|\mathscr{F}_1]]\, dQ = \int_{B_2} [f - E_Q[f|\mathscr{F}_1]]\, dQ = 0. \tag{4.25}$$

Thus, by (4.15),

$$\int_C f\, dQ = \int_C E_Q[f|\mathscr{F}_1]\, dQ.$$

A similar argument shows that for all $Q \in b_\epsilon$,

$$\int_C f\, dQ = \int_C E_Q[f|\mathscr{F}_2]\, dQ,$$

Therefore, $A_1 \cap A_2 \in \Lambda_Q$ for every $Q \in b_\epsilon$. Since A_1 and A_2 were arbitrary, this shows that $\Pi \in \Lambda_Q$ for all $Q \in b_\epsilon$. Therefore, by the Dynkin's π-λ theorem, $\sigma(\Pi) = \mathscr{F} \subseteq \Lambda_Q$ for $Q \in b_\epsilon$; that is, for any $A \in \mathscr{F}$,

$$\int_A f\, dQ = \int_A E_P[f|\mathscr{F}_1]\, dQ = \int_A E_P[f|\mathscr{F}_2]\, dQ.$$

Together with the fact that $E_Q[f|\mathscr{F}_1]$ and $E_Q[f|\mathscr{F}_2]$ are both measurable with respect to $\mathscr{F}$, this shows that $E_Q[f|\mathscr{F}] = E_Q[f|\mathscr{F}_1] = E_Q[f|\mathscr{F}_2]$ for all $Q \in b_\epsilon$. Thus, b_ϵ and $\mathbf{Q}$ are disjoint. Recall that the set $u \in \mathscr{U}$ was arbitrary. Therefore, for any set u in $\mathscr{U}$, there exists some v contained in u such that v and $\mathbf{Q}$ are disjoint. This shows that $\mathbf{Q}$ is nowhere dense in $\mathbf{P}$. $\square$

Proof of Theorem 4.2 In light of part (c) of Theorem 1′, in order to prove the theorem it is sufficient to show that there exists a residual set $\mathbf{R} \subseteq \mathbf{P}$ such that for all $P \in \mathbf{R}$,

$$E_P\left[\theta \middle| \mathscr{I}_i^{\sigma^P}\right] = E_P\left[\theta \middle| \mathscr{I}^{\sigma^P}\right] \quad \text{for all } i \in N.$$

For any pair of agents $i, j \in N$, define $\mathbf{M}_{ij} \subseteq \mathbf{P}$ as

$$\mathbf{M}_{ij} = \left\{P \in \mathbf{P} : E_P\left[\theta \middle| \mathscr{I}_i^{\sigma^P}\right] = E_P\left[\theta \middle| \mathscr{I}_j^{\sigma^P}\right] \neq E_P\left[\theta \middle| \mathscr{I}_{ij}^{\sigma^P}\right]\right\},$$

where $\mathscr{I}_{ij}^{\sigma^P}$ is the σ-algebra generated by the union of $\mathscr{I}_i^{\sigma^P}$ and $\mathscr{I}_j^{\sigma^P}$. We first prove that $\mathbf{M}_{ij}$ is a meager set. Let $\mathbf{D}_t \subset \mathbf{P}$ be the set of all probability measures P such that $P(s_\tau) = \mathbb{I}_{\{s_\tau = s\}}$ for some $s \in S$ and all $\tau > t$, where $\mathbb{I}$ is the indicator function. When the state and the private signals are realized according to some P belonging to $\mathbf{D}_t$, then the signal profiles generated after time t are constant and thus completely uninformative. Trivially, it is true that

$$\mathbf{P} = \bigcup_{t=1}^{\infty} \mathbf{D}_t.$$

Let $\mathbf{M}_{ijt} = \mathbf{M}_{ij} \cap \mathbf{D}_t$. Then, by the above equality,

$$\mathbf{M}_{ij} = \bigcup_{t=1}^{\infty} \mathbf{M}_{ijt}.$$

Therefore, for $\mathbf{M}_{ij}$ to be a meager set, it is sufficient that $\mathbf{M}_{ijt}$ is a meager subset of $\mathbf{D}_t$ for any t.

We prove this by proving that the set $\mathbf{Q}_{ijt} \supseteq \mathbf{M}_{ijt}$ defined below is meager.

$$\mathbf{Q}_{ijt} = \left\{P \in \mathbf{D}_t : E_P\left[\theta \middle| \mathscr{I}_i^{\sigma^{P'}}\right] = E_P\left[\theta \middle| \mathscr{I}_j^{\sigma^{P'}}\right] \neq E_P\left[\theta \middle| \mathscr{I}_{ij}^{\sigma^{P'}}\right] \quad \text{for some} \quad P' \in \mathbf{D}_t\right\}$$

Note that, for all $P \in \mathbf{D}_t$, the signal profiles generated by P after time t are constant. Therefore, for all i, any strategy profile σ, and any $P \in \mathbf{D}_t$, we have that $\mathscr{I}_i^{\sigma^P}$ is a sub σ-algebra of $\mathscr{S}_t$, the Borel σ-algebra generated by the signal profiles up to time t. Given two arbitrary sub σ-algebras $\mathscr{S}_{it}, \mathscr{S}_{jt} \subseteq \mathscr{S}_t$ and the σ-algebra generated by their union $\mathscr{S}_{ijt}$, define

$$\mathbf{S}_{ijt}(\mathscr{S}_{it}, \mathscr{S}_{jt}) = \left\{P \in \mathbf{D}_t : E_P[\theta|\mathscr{S}_{it}] = E_P\left[\theta|\mathscr{S}_{jt}\right] \neq E_P\left[\theta|\mathscr{S}_{ijt}\right]\right\},$$

$\mathbf{Q}_{ijt}$ is a subset of the union of the above sets over all σ-algebra pairs $\mathscr{S}_{it}, \mathscr{S}_{jt}$. Since S is finite and for $P \in \mathbf{D}_t$ the signals are constant after time t, the σ-algebra $\mathscr{S}_t$ is finite. Therefore, there are finitely many such $\mathbf{S}_{ijt}$ sets. Consequently, it is sufficient to show that any $\mathbf{S}_{ijt}$ is meager in $\mathbf{D}_t$ in order to conclude that $\mathbf{Q}_{ijt}$, and hence $\mathbf{M}_{ijt}$, are meager in $\mathbf{D}_t$. Note that the set $\mathbf{D}_t$ is the set of probability measures over $\Theta \times S^t$ and θ is a bounded random variable over this set. Moreover, $\mathscr{S}_{it}, \mathscr{S}_{jt}$ are two arbitrary fixed sub σ-algebras of the Borel σ-algebra of $\Theta \times S^t$. Therefore, we can directly use Lemma 4.3 to conclude that $\mathbf{S}_{ijt}(\mathscr{S}_{it}, \mathscr{S}_{jt})$ is nowhere dense in $\mathbf{D}_t$; therefore, $\mathbf{M}_{ij}$ is a meager subset of $\mathbf{P}$.

The above argument shows that for any pair of agents $i, j \in N$, the set $\mathbf{M}_{ij}$ is a meager subset of $\mathbf{P}$. We can use this result to argue similarly that for any $i, j, k \in N$, the set $\mathbf{M}_{ijk}$ defined below is a meager subset of $\mathbf{P}$.

$$\mathbf{M}_{ijk} = \left\{P \in \mathbf{P} : E_P\left[\theta \middle| \mathscr{I}_i^{\sigma^P}\right] = E_P\left[\theta \middle| \mathscr{I}_j^{\sigma^P}\right] = E_P\left[\theta \middle| \mathscr{I}_k^{\sigma^P}\right] \neq E_P\left[\theta \middle| \mathscr{I}_{ijk}^{\sigma^P}\right]\right\},$$

where $\mathscr{I}_{ijk}^{\sigma^P}$ is the σ-algebra generated by the union of $\mathscr{I}_i^{\sigma^P}$, $\mathscr{I}_j^{\sigma^P}$, and $\mathscr{I}_k^{\sigma^P}$. Proceeding inductively we can prove the lemma. $\square$

CHAPTER 5

CONCLUSIONS

5.1 Thesis Summary

In this thesis, we studied the problem of information aggregation in social networks under three distinct assumptions on agent behavior. In the second chapter, we studied the inference problem faced by a group of fully rational agents who interact over an unknown social network. In the third chapter, we studied a similar inference problem when agents are instead boundedly rational and showed that boundedly rational agents can asymptotically aggregate the dispersed information as well as their fully rational counterparts. We further provided an explicit characterization of the rate of learning and used the result to perform comparative analysis. In the fourth chapter of the thesis, we introduced a novel model of opinion formation in which agents attempt to act in conformity with each other while also matching the unknown state. For each of the models, the focus was on characterizing the conditions on network and information structures that lead to consensus and information aggregation. We summarize the presented results in the following.

In Chapter 2, we studied the behavior of a group of individuals who are fully rational and are only concerned with learning some unknown state of the world. We showed that communications between rational individuals with access to complementary pieces of information eventually direct everyone to discover the truth. This result holds so long as the state is collectively identifiable, information can flow from any individual to any other one, and the agents are endowed with a common prior—even if no individual is able to identify the state on her own or if the network structure is not commonly known by the agents. Repeated interactions between rational agents eventually lead to efficient aggregation of information. This result can be viewed as a positive result on the possibility of full information aggregation by Bayesian agents. Yet, fully rational agent behavior may not be a realistic assumption when dealing with large societies and complex networks due to the extreme computational complexity of Bayesian inference.

Motivated by this observation, in Chapter 3, we explored the implications of bounded rationality by introducing biases in the way agents interpret the opinions of others while at the same time maintaining the assumption that agents interpret their private observations rationally. Our analysis yields the result that when faced with overwhelming evidence in favor of the truth even biased agents will eventually learn to discover the truth.

We moreover characterized the rate of learning in terms of the relative entropy of different agents' signal structures and their eigenvector centralities. We showed that if the agents' signal structures are comparable in the sense of uniform informativeness, then the rate of learning is maximized when the most central agents receive signals of the highest quality. We also showed that in the presence of experts—i.e., agents with access to information crucial for learning—the role played by the social network structure is inverted. In particular, learning is slower if the high quality signals are assigned to the more central agents rather than the ones at the periphery of the network. This result is a consequence of the fact that learning is slowed down whenever the information and identification bottleneck effects reinforce one another. More specifically, if the information required for distinguishing the pair of states that are hardest to tell apart is only available to an agent who receives very little effective attention from others, then it would take a long time for (i) that agent to collect enough information to distinguish between the two states; and (ii) for this information to be diffused throughout the network. On the other hand, a negative assortative matching of signal qualities and eigenvector centralities guarantees that these two events happen in parallel, leading to a faster convergence rate.

Furthermore, by defining the novel notion of regularity as a measure of the extent of asymmetry in the network structure, we provided a comparative analysis of the role of the structural properties of the social network on the long run dynamics of the beliefs. We showed that even though the speed of learning is smaller in more regular networks when the signal structures are comparable in the uniform sense, the exact opposite is true when each agent possesses some information crucial for learning.

Finally, in Chapter 4, we studied a dynamic game in which a number of agents attempt to coordinate on an outcome about which they have incomplete and asymmetric information. Any agent's actions reveal information which is used by other agents to revise their beliefs, and hence, their actions. We proved formal results regarding the asymptotic outcomes obtained when myopic agents play the actions prescribed by the weak perfect Bayesian equilibrium. In particular, we showed that the agents reach consensus in their actions if the observation network is connected, and the consensus action is generically optimal if the agents' private observations are exogenously generated and the signal space is finite.

5.2 Future Directions

In Chapter 3, we studied a variant of the DeGroot (1974) learning model wherein agents repeatedly make new private observations. DeMarzo, Vayanos, and Zwiebel (2003) argue that the agents in the DeGroot model update their beliefs as Bayesian agents who make noisy observations of a normally distributed unknown parameter, except that the agents in DeGroot's model fail to account for repetitions in their observations. In other words, the agents in the DeGroot model are Bayesian except that they lack perfect recall of their past observations. As such, a promising research direction is to investigate the non-Bayesian updating rules that can be obtained from Bayes rule in different environments by imposing the additional restriction that the agents lack perfect recall. An advantage of such a framework over the seemingly *ad hoc* array of non-Bayesian models proposed in the literature is that it can be used even in environments where no obvious non-Bayesian update rule is readily available, e.g., when agents observe the actions of their neighbors instead of their beliefs.

Another interesting future direction is to test the predictions of the model studied in Chapter 3 using lab experiments or field data. The theoretical model yields refined predictions on the dependence of the speed of information aggregation in a network on the allocation of information across agents that can be empirically identified. Empirical evidence could further elucidate the usefulness and limitations of the theoretical results and inform the design of public policy regarding information dissemination when information sources are scarce or costly.[41]

In Section 4.5, we proved Lemma 4.3 as an intermediate step in the proof of Theorem 4.2. In simple terms it states the following: given a generic set of priors, consensus implies learning. The proof of this lemma makes no use of the structure of the model being studied in Chapter 4. Indeed, the result continues to hold in a much more general framework. Consider a group of agents who are endowed with a common prior over some state space and who reach an agreement in their beliefs about an event of interest based on some information. Lemma 4.3 states that such agents would generically not further benefit from direct communication of their information as their beliefs already reflect all the information relevant to the event of interest. This is the case irrespective of the opinion dynamic that lead the agents to consensus. In future research, we intend to generalize this result by showing that our topological notion of genericity can be substituted with a measure-theoretic one, and to study its implications for problems such as information aggregation in markets.

Finally, we remark that although the results of Chapter 4 were proved under the assumption that the agents' preferences are represented by a quadratic utility function, the insights of our analysis do not hinge on the particular functional form of the utility function. In fact, similar results can be proved for more general symmetric coordination games with payoffs that can be approximated well with a quadratic function. Similarly, the assumption on the myopia of the agents could be dispensed with if the agents were assumed to be atomistic. An interesting future research direction is to develop a formal model that relaxes the assumptions on quadratic utility function and myopic agent behavior.

[41] I would like to thank Dimitri Vayanos for suggesting this research direction.

Appendix A

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 3.3 provides a more complete order over the set of signal structures than either of the notions above.

Proposition A.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

$$\sum_{s\in S} \ell^{\theta}(s)\phi\left(\frac{\ell^{\hat{\theta}}(s)}{\ell^{\theta}(s)}\right) \geq \sum_{s\in S} \ell'^{\theta}(s)\phi\left(\frac{\ell'^{\hat{\theta}}(s)}{\ell'^{\theta}(s)}\right), \tag{A.1}$$

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 A.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 (A.1) should hold for all convex functions ϕ, whereas for uniform informativeness, it is sufficient that (A.1) is satisfied for $\phi(x) = \log(x)$. Thus, unlike most other information orders (such as Lehmann (1988)'s), uniform informativeness does not coincide with Blackwell's order on dichotomies.

APPENDIX B

REGULARITY AND NETWORK SYMMETRY

Acemoglu et al. (2012) define a measure of symmetry of network structures according to which network A is more symmetric than network A' if $\|v\|_2 \leq \|v'\|_2$, where v and v' are the centrality vectors corresponding to A and A', respectively. The next result relates this notion to the notion of regularity defined in Section 3.5.

Proposition B.1. *If social network A is more regular than A', then $\|v\|_2 \leq \|v'\|_2$.*

Before presenting the proof, we state and prove a simple lemma.

Lemma B.1. *Suppose that vector y' majorizes vector y as defined in* (3.7)*. Then, for any nonnegative vector x,*

$$\sum_{i=1}^{n} x_{[i]} y_{[i]} \leq \sum_{i=1}^{n} x_{[i]} y'_{[i]}.$$

Proof. By assumption,

$$\sum_{i=1}^{k} y_{[i]} \leq \sum_{i=1}^{k} y'_{[i]}$$

for all $k \leq n$. Multiplying both sides of the above inequality by the nonnegative number $x_{[k]} - x_{[k+1]}$ and summing over all k imply

$$\sum_{k=1}^{n} \left(x_{[k]} - x_{[k+1]}\right) \sum_{i=1}^{k} y_{[i]} \leq \sum_{k=1}^{n} \left(x_{[k]} - x_{[k+1]}\right) \sum_{i=1}^{k} y'_{[i]},$$

with the convention that $x_{[n+1]} = 0$. Finally, noticing that

$$\sum_{k=1}^{n} \left(x_{[k]} - x_{[k+1]}\right) \sum_{i=1}^{k} y_{[i]} = \sum_{i=1}^{n} x_{[i]} y_{[i]}$$

completes the proof. □

Proof of Proposition B.1 Given that A is more regular than A', the eigenvector centrality v is majorized by v'. Thus, by Lemma B.1,

$$\|v\|_2^2 = \sum_{i=1}^{n} v_{[i]} v_{[i]} \leq \sum_{i=1}^{n} v_{[i]} v'_{[i]}.$$

Applying Lemma B.1 once again implies

$$\|v'\|_2^2 = \sum_{i=1}^{n} v'_{[i]} v'_{[i]} \geq \sum_{i=1}^{n} v_{[i]} v'_{[i]}.$$

Combining the above two inequalities completes the proof. □

APPENDIX C

STABILITY OF RANDOM DYNAMICAL SYSTEMS

Dynamical Systems

The definitions regarding dynamical systems are standard and can be found, among many other places, in the book of Sastry (1999).

A discrete-time *dynamical system* on the Euclidean space (henceforth simply called dynamical system) is a mapping $\varphi : \mathbb{N} \times \mathbb{R}^k \to \mathbb{R}^k$ that gives the time t consequent $\varphi_t(x)$ to a state x. The trajectory of dynamical system φ with initial state x_0 is the time-ordered collection of states $x_0, x_1, \ldots$ such that $x_t = \varphi_t(x_{t-1})$ for all $t \in \mathbb{N}$.

A dynamical system is called linear if $\varphi_t(x) = M_t x$ for some $M_t \in \mathbb{R}^{k \times k}$ and all $t \in \mathbb{N}$. The *trajectory* of a linear dynamical system with initial state x_0 is given by $x_t = \Phi(t) x_0$, where $\Phi(t)$ is the *evolution operator* (or the state transition matrix) given by

$$\Phi(t) = M_t \cdots M_1 \tag{C.1}$$

for all $t \in \mathbb{N}$.

A *fixed point* (or equilibrium) of φ is a state x_{eq} such that $x_{\text{eq}} = \varphi_t(x_{\text{eq}})$ for all $t \in \mathbb{N}$.

Lyapunov Exponents

The concepts of Lyapunov exponents, regularity, and Lyapunov stability date back to Lyapunov. A modern treatment for the case of ergodic ODEs can be found in the book by Barreira and Pesin (2002).

Let $\varphi_t(x) = M_t x$ be a linear dynamical system with bounded M_t. The *Lyapunov exponent* of the trajectory $\Phi(t)x$ of the dynamical system starting with initial state x is the mapping $\lambda : \mathbb{R}^k \to \mathbb{R} \cup \{-\infty\}$ defined as

$$\lambda(x) = \limsup_{t \to \infty} \frac{1}{t} \log \|\Phi(t) x\|. \tag{C.2}$$

The Lyapunov exponents do not depend on the matrix norm used.

The function λ takes at most a number $r \leq k$ of distinct values on $\mathbb{R}^k \setminus \{0\}$. Write

$$-\infty \leq \lambda_r < \lambda_{r-1} < \cdots < \lambda_1 < \infty$$

for different values of $\lambda(x)$ and call them the Lyapunov exponents of the dynamical system. λ_1 is called the top Lyapunov exponent. The sets $V_i = \left\{x \in \mathbb{R}^k : \lambda(x) \leq \lambda_i\right\}$ are linear subspaces of $\mathbb{R}^k$ for $i = 1, \ldots, r$. The integer $d_i = \dim V_i - \dim V_{i+1}$ is called the multiplicity of λ_i.

The sequence $\Phi(t)$ is said to be *Lyapunov regular* if

$$\sum_{i=1}^{r} d_i \lambda_i = \liminf_{n\to\infty} \frac{1}{t} \log|\det \Phi(t)|.$$

For a Lyapunov regular system all lim sup's in (C.2) are in fact limits.

Lyapunov Stability

Let x_t be the trajectory of dynamical system φ with initial state x_0. Also assume that

$$\varphi_t(x) = M_t x + f_t(x) \tag{C.3}$$

for all $t \in \mathbb{N}$ and $x \in \mathbb{R}^k$. Without loss of generality, we assume that the origin is a fixed point of φ.

The origin is *stable* if for every neighborhood U of the origin and all $t_0 \in \mathbb{N}$ there is a neighborhood $V \subseteq U$ of the origin such that if $x_{t_0} \in V$, then $x_t \in U$ for all $t \geq t_0$.

The origin is *uniformly stable* if in the preceding definition V can be chosen independent of t_0.

The origin is *asymptotically stable* if it is stable and additionally V can be chosen so that $x_t \to 0$ as $t \to \infty$ for all $x_{t_0} \in V$.

The origin is *uniformly asymptotically stable* if it is uniformly stable and additionally V can be chosen so that $x_t \to 0$, uniformly in t_0, as $t \to \infty$ for all $x_{t_0} \in V$.

The origin is *exponentially stable* if there is a neighborhood V of the origin and constants $c, \lambda > 0$ such that if $x_{t_0} \in V$, then $\|x_t\| \leq ce^{-\lambda(t-t_0)}\|x_{t_0}\|$ for all $t \geq t_0$.

We say that the above properties hold globally if the neighborhood V can be chosen to be the entire Euclidean space.

A central question in Lyapunov stability theory is whether stability of a fixed point of a linear dynamical system is preserved when the system is subject to small perturbations. The following result is a partial answer to this question for Lyapunov regular systems. The following is a corollary of Theorems 1 and 2 of Barreira and Valls (2007).

Theorem (Barreira and Valls)**.** *Consider the dynamical system $\varphi_t(x) = M_t x + f_t(x)$ with trajectory $x_t \in \mathbb{R}^k$, where $M_t \in \mathbb{R}^{k\times k}$ for all t. Also, suppose that the following hold:*

(a) The linear dynamical system M_t is Lyapunov regular.

(b) The Lyapunov exponents corresponding to M_t are all negative, that is,

$$\lambda_1 = \lim_{t\to\infty} \frac{1}{t} \log\|M_t M_{t-1} \dots M_1\| < 0.$$

(c) $f_t(x)$ is a continuous map with $f_t(0) = 0$ for all t.

(d) There are constants $C, q > 0$ such that $\|f_t(x) - f_t(y)\| \leq C\|x - y\|\big(\|x\|^q + \|y\|^q\big)$, for all t and all $x, y \in \mathbb{R}^k$.

For all t_0 and $\epsilon > 0$, there exist a neighborhood V of the origin and a constant K such that if $x_{t_0} \in V$, then for all $t \geq t_0$,

$$\|x_t\| \leq Ke^{(t-t_0)(\lambda_1+\epsilon)+\epsilon t_0}\|x_{t_0}\|.$$

Oseledets' Theorem

Oseledets' multiplicative ergodic theorem (MET) is the most fundamental theorem in the study of random dynamical systems. The following result, which can be found in Oseledets (2008), is part of one of various versions of the MET.

Let $\mathscr{T}$ be a measure preserving transformation of a probability space $(\Omega, \mathscr{F}, \mathbb{P})$ and

$$\Phi(\omega; t) = M(\mathscr{T}^{t-1}\omega) \ldots M(\omega), \tag{C.4}$$

where M is a measurable map to the space of invertible $k \times k$ real valued matrices satisfying

$$\int \max\left\{\log \|M(\omega)\|, 0\right\} d\mathbb{P} < \infty.$$

Theorem (Oseledets)**.** *The function $t \mapsto \Phi(\omega; t)$ is* Lyapunov regular *for $\mathbb{P}$-almost every ω.*

BIBLIOGRAPHY

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar (forthcoming), “Dynamics of information exchange in endogenous social networks.” *Theoretical Economics*.

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2012), “The network origins of aggregate fluctuations.” *Econometrica*, 80, 1977–2016.

Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), “Bayesian learning in social networks.” *The Review of Economic Studies*, 78, 1201–1236.

Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi (2010), “Spread of (mis) information in social networks.” *Games and Economic Behavior*, 70, 194–227.

Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken (2012), “Network structure and the aggregation of information: Theory and evidence from Indonesia.” NBER Working Paper 18531.

Amador, Manuel and Pierre-Olivier Weill (2012), “Learning from private and public observations of others’ actions.” *Journal of Economic Theory*, 147, 910–940.

Angeletos, George-Marios and Alessandro Pavan (2007), “Efficient use of information and social value of information.” *Econometrica*, 75, 1103–1142.

Angeletos, George-Marios and Alessandro Pavan (2009), “Policy with dispersed information.” *Journal of the European Economic Association*, 7, 11–60.

Aumann, Robert J. (1976), “Agreeing to disagree.” *The Annals of Statistics*, 1236–1239.

Bala, Venkatesh and Sanjeev Goyal (1998), “Learning from neighbours.” *The Review of Economic Studies*, 65, 595–621.

Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou (2006), “Who’s who in networks. Wanted: the key player.” *Econometrica*, 74, 1403–1417.

Banerjee, Abhijit (1992), “A simple model of herd behavior.” *The Quarterly Journal of Economics*, 107, 797–817.

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson (2012), “The diffusion of microfinance.” NBER Working paper No. 17743.

Barreira, Luis and Yakov B. Pesin (2002), *Lyapunov Exponents and Smooth Ergodic Theory*. American Mathematical Society.

Barreira, Luis and Claudia Valls (2007), "Stability theory and lyapunov regularity." *Journal of Differential Equations*, 232, 675–701.

Berman, Abraham and Robert J. Plemmons (1979), *Nonnegative Matrices in the Mathematical Sciences*. Academic Press.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy*, 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.

Blume, Lawrence and David Easley (2006), "If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets." *Econometrica*, 74, 929–966.

Borkar, Vivek and Pravin Varaiya (1982), "Asymptotic agreement in distributed estimation." *IEEE Transactions on Automatic Control*, 27, 650–655.

Bramoullé, Y., R. Kranton, and M. D'Amours (2009), "Strategic interaction and networks." Working Paper.

Cabrales, Antonio, Olivier Gossner, and Roberto Serrano (2012), "The appeal of information transactions." Universidad Carlos III de Madrid Working Paper 12-24.

Calvó-Armengol, Antoni and Joan de Martí Beltran (2009), "Information gathering in organizations: equilibrium, welfare, and optimal network structure." *Journal of the European Economic Association*, 7, 116–161.

Calvó-Armengol, Antoni, Joan de Martí, and Andrea Prat (2011), "Communication and influence." Working paper.

Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri (2011), "Testing models of social learning on networks: Evidence from a framed field experiment." Working paper.

Choi, Syngjoo, Douglas Gale, and Shachar Kariv (2005), "Behavioral aspects of learning in social networks: an experimental study." In *Advances in Applied Microeconomics* (Experimental and Behavioral Economics) (John Morgan, ed.), volume 13, 25–61, Emerald Group Publishing Limited.

Choi, Syngjoo, Douglas Gale, and Shachar Kariv (2012), "Social learning in networks: a quantal response equilibrium analysis of experimental data." *Review of Economic Design*, 16, 135–157.

Corazzini, Luca, Filippo Pavesi, Beatrice Petrovich, and Luca Stanca (2012), "Influential listeners: An experiment on persuasion bias in social networks." *European Economic Review*, 56, 1276–1288.

Cover, Thomas M. and Joy A. Thomas (1991), *Elements of Information Theory*. John Wiley & Sons.

Crimaldi, Irene and Luca Pratelli (2005), "Two inequalities for conditional expectations and convergence results for filters." *Statistics & Probability Letters*, 74, 151–162.

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.

Durrett, Rick (2010), *Probability: Theory and Examples*. Cambridge University Press.

Eksin, Ceyhun, Pooya Molavi, Alejandro Ribeiro, and Ali Jadbabaie (2013), "Bayesian quadratic network game filters." Arxiv Preprint arXiv:1302.0249.

Elliott, Matthew and Benjamin Golub (2013), "A network approach to public goods." Working paper.

Foster, Dean P. and H. Peyton Young (2003), "Learning, hypothesis testing, and Nash equilibrium." *Games and Economic Behavior*, 45, 73–96.

Furstenberg, Harry and Harry Kesten (1960), "Products of random matrices." *The Annals of Mathematical Statistics*, 31, 457–469.

Gale, Douglas and Shachar Kariv (2003), "Bayesian learning in social networks." *Games and Economic Behavior*, 45, 329–346.

Geanakoplos, John D. and Heraklis M. Polemarchakis (1982), "We can't disagree forever." *Journal of Economic Theory*, 28, 192–200.

Gharavi, Reza and V. Anantharam (2005), "An upper bound for the largest Lyapunov exponent of a Markovian product of nonnegative matrices." *Theoretical Computer Science*, 332, 543–557.

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 (2012a), "Does homophily predict consensus times? Testing a model of network structure via a dynamic process." *Review of Network Economics*, 11.

Golub, Benjamin and Matthew O. Jackson (2012b), "How homophily affects the speed of learning and best-response dynamics." *The Quarterly Journal of Economics*, 127, 1287–1338.

Hardy, Godfrey Harold (1992), *Divergent Series*. American Mathematical Society.

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

Jackson, Mathew O. (2008), *Social and Economic Networks*. Princeton University Press.

Jackson, Matthew O. and Ehud Kalai (1997), "Social learning in recurring games." *Games and Economic Behavior*, 21, 102–134.

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." Working Paper.

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

Jordan, James S. (1991), "Bayesian learning in normal form games." *Games and Economic Behavior*, 3, 60–81.

Jordan, James S. (1995), "Bayesian learning in repeated games." *Games and Economic Behavior*, 9, 8–20.

Kalai, Ehud and Ehud Lehrer (1993), "Rational learning leads to nash equilibrium." *Econometrica*, 61, 1019–1045.

Kesten, Harry and Frank Spitzer (1984), "Convergence in distribution of products of random matrices." *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 67, 363–386.

Keynes, John M. (1936), *The General Theory of Interest, Employment and Money*. London: Macmillan.

Kullback, Solomon and Richard A. Leibler (1951), "On information and sufficiency." *The Annals of Mathematical Statistics*, 22, 79–86.

Landers, Dieter and Lothar Rogge (1976), "Convergence of conditional p-means given a σ-lattice." *The Annals of Probability*, 4, 147–150.

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

Lehrer, Ehud and Rann Smorodinsky (2000), "Relative entropy in sequential decision problems." *Journal of Mathematical Economics*, 33, 425–439.

Lobel, Ilan and Evan Sadler (2012), "Social learning and aggregate network uncertainty." Working paper.

McNiven, Scott C. and Daniel O. Gilligan (2011), "Networks and constraints on the diffusion of a biofortified agricultural technology: Evidence from a partial population experiment." University of California - Davis Working Paper.

Miguel, Edward and Michael Kremer (2003), "Networks, social learning, and technology adoption: The case of deworming drugs in Kenya." Working Paper.

Molavi, Pooya, Ceyhun Eksin, Alejandro Ribeiro, and Ali Jadbabaie (2013), "Learning to coordinate in social networks." Working Paper.

Morris, Stephen and Hyun Song Shin (2002), "Social value of public information." *The American Economic Review*, 92, 1521–1534.

Mossel, Elchanan and Omer Tamuz (2010), “Efficient Bayesian learning in social networks with Gaussian estimators.” Arxiv Preprint arXiv:1002.0747.

Mueller-Frank, Manuel (2013), “A general framework for rational learning in social networks.” *Theoretical Economics*, 8, 1–40.

Nachbar, John H. (1997), “Prediction, optimization, and learning in repeated games.” *Econometrica*, 65, 275–309.

Oseledets, Valery I. (1968), “A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems.” *Transactions of the Moscow Mathematical Society*, 19, 197–231.

Oseledets, Valery. I (2008), “Oseledets theorem.” *Scholarpedia*, 3, 1846.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2009), “Informational externalities and emergence of consensus.” *Games and Economic Behavior*, 66, 979–994.

Sandroni, Alvaro (2000), “Do markets favor agents able to make accurate predictions?” *Econometrica*, 68, 1303–1341.

Sastry, S. Shankar (1999), *Nonlinear Systems: Analysis, Stability, and Control*. Springer Verlag.

Tahbaz-Salehi, Alireza (2009), *Complex Networks: New Models and Distributed Algorithms*. Ph.D. thesis, University of Pennsylvania.

Vives, Xavier (1997), “Learning from others: A welfare analysis.” *Games and Economic Behavior*, 20, 177–200.

Vives, Xavier (2010), *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press.