Learning to Coordinate in Social Networks*

Pooya Molavi† Ceyhun Eksin† Alejandro Ribeiro† Ali Jadbabaie†

This version: July 2013

Abstract

We study a repeated game in which a group of players attempt to coordinate on a desired, but only partially known, outcome. The desired outcome is represented by an unknown state of the world. Agents' stage payoffs are represented by a quadratic utility function that captures the kind of trade-off exemplified by the Keynesian beauty contest: each agent's stage payoff is decreasing in the distance between her action and the unknown state; it is also decreasing in the distance between her action and the average action taken by other agents. The agents thus have the incentive to correctly estimate the state while trying to coordinate with and learn from others. We show that myopic but Bayesian agents who repeatedly play this game and observe the actions of their neighbors over a network (that satisfies some weak connectivity condition) eventually succeed in coordinating on a single action. The agents also asymptotically receive similar payoffs in spite of differences in the quality of their information. Finally, we show that if the agents' private observations are not functions of the history of the game, then the private observations are optimally aggregated in the limit. Therefore, agents asymptotically coordinate on choosing the best estimate of the state given the aggregate information available throughout the network.

Keywords: Learning, coordination games, social networks.
JEL Classification: C73, D83, D85.

*We are grateful to George Mailath, Ali Shourideh, and Alireza Tahbaz-Salehi for useful feedback and suggestions. The authors acknowledge financial support from the Air Force Office of Scientific Research (MURI Award No. FA9550-10-1-0567) and the Army Research Office (MURI Award No. W911NF-12-1-0509).

†Department of Electrical and Systems Engineering, University of Pennsylvania.

1 Introduction

Social networks spread information, encourage imitation, and thereby, facilitate coordination by speeding the resolution of uncertainties. They help individuals make consumption and investment decisions or perform tasks in organizations, all by providing them with information on the social, economic, or political fundamentals and on the other individuals' actions and beliefs.

An example demonstrating the role of social networks in coordination and learning is the recent wave of popular protests in Egypt, Iran, Tunisia, and more recently in Brazil and Turkey. These large scale events are widely believed to have been instigated by social media. Social networks act as conduits of information about the time and location of protests and the accounts of the events, and help individuals decide on their level of participation by providing them with information about the unknown "states" (such as the forcefulness of the police at an event) or the participation decisions of others individuals.[1]

In light of these observations, this paper examines the coordination problem faced by a group of agents when the relevant information is dispersed throughout a social network. Consider a group of agents that wish to coordinate on a desired outcome that is not fully known to any one of them. Agents choose actions which are close to what they consider to be the desired outcome. Yet, they also need to coordinate with other agents by choosing actions that are similar to what they expect others to choose. There is a trade-off between acting according to one's best estimate of the desired outcome and trying to coordinate with other agents. Such trade-offs are important in—besides the example previously mentioned—trade decision in financial markets (Morris and Shin (2002)), consumption decisions (Bramoullé, Kranton, and D'Amours (2009)), and problems in cooperative robotics (Marden, Arslan, and Shamma (2009)) and organizational coordination (Calvó-Armengol and Beltran (2009)). The decisions of traders in 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 examples, agents make decisions by attempting to second-guess the decisions of others while also guessing the value of an unknown. Moreover, oftentimes agents can only communicate with a handful of other agents, while at the same time, trying to coordinate with and learn from everybody else.

We use the framework of repeated 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

[1] According to Zeynep Tufekci of Princeton University, who interviewed scores of Turkish protesters, most cited social media as a spur. ("The digital demo", *The Economist*, June 29, 2013)

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 re-evaluate their beliefs about the underlying parameter and their predictions of others' future behavior. These re-evaluations 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. In this paper, 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.[2] 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.

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 the results presented by Jadbabaie,

[2]A series of results in game theory, all of them known by the name "folk theorem", establishes that in games played by sufficiently patient forward-looking agents, any individually rational payoff can be obtained as an equilibrium payoff. We are not aware of any folk theorem that directly applies to our model. However, based on the results proved in the literature, a unique equilibrium is unlikely to obtain in our setting if the agents are forward-looking. For two examples of a folk theorem, a classic result and a more recent result proved for games played on networks, see the works of Abreu, Pearce, and Stacchetti (1990) and Laclau (2012).

Molavi, Sandroni, and Tahbaz-Salehi (2012) and 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.

Related Literature The paper 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 paper, 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 paper, 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 paper, 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).

Organization of the Paper The rest of the paper is organized as follows. We present the baseline model in Section 2. In Section 3, we introduce our equilibrium notion, prove its existence and uniqueness, and provide a simple characterization of the equilibrium strategies. In Section 4, we study the asymptotic equilibrium behavior of the agents and argue that they reach consensus in their actions and payoffs. Section 5 provides conditions under which agents coordinate on the efficient action. Finally, Section 6 concludes the paper. Proofs are provided in the Appendix.

2 Baseline Model

2.1 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 μ. We make the following technical assumption on μ.

Assumption 1. The state is square integrable with respect to μ, that is,

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

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:

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

where $\lambda \in [0, 1)$ is a constant and $\bar{a}_{-i} = \frac{1}{n-1} \sum_{j \neq i} a_j$ 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).

2.2 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 $[g_t]_{ji} = 1$ if and only if agent j is a time t neighbor of agent i, that is, if $j \in N_{it}$.

Assumption 2. The network g_t is generated according to some probability distribution ν_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 cases where an agent's realized neighborhoods are informative about the state or other agents' signals. We maintain Assumption 2 throughout the paper.

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 $\nu = \times_{t=1}^{\infty}\nu_t$ 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 ν.

Assumption 3. For ν-almost all $\{g_t\}_{t\in\mathbb{N}}$, there exists a strongly connected network $\bar{g}$ such that if j is a neighbor of i given $\bar{g}$, 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.

2.3 Histories

Let $(\Omega, \mathcal{B})$ denote the measurable space of plays, where $\Omega = \Theta \times (S \times A \times G)^{\mathbb{N}}$ and $\mathcal{B}$ 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 $\mathcal{H}_t \subseteq \mathcal{B}$ 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.[3]

[3] Given measurable spaces $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$, a function $f : X \times \mathcal{Y} \to [0, 1]$ 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 $(Y, \mathcal{Y})$; and (ii) given any measurable set $B \in \mathcal{Y}$, the function $x \mapsto f(x)[B]$ is measurable.

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 $\mathcal{H}_{it} \subseteq \mathcal{H}_t$ and $\mathcal{H}_i \subseteq \mathcal{B}$ denote the σ-algebras of subsets of Ω generated by the Borel sets of H_{it} and H_i, respectively.

2.4 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 1. A *pure behavior strategy* for agent i is a 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 $\sigma = (\sigma_1, \ldots, \sigma_n)$, where σ_i is a strategy for agent i.

Any strategy profile σ—together with the agents' common prior μ, the stochastic process ν, and the signaling functions $\{\pi_t\}_{t \in \mathbb{N}}$—induces a probability distribution over the measurable space $(\Omega, \mathcal{B})$, 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 $(\Omega, \mathcal{F}, \mathbb{P}^{\sigma})$. The realized time t private history of agent i is in turn a measurable function 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 2. A *belief system* for agent i is a transition probability $q_i : H_i \times \mathcal{B} \to [0, 1]$.

A belief is a probably distribution over the space of plays $(\Omega, \mathcal{B})$, 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 $\tilde{q}_{it}(\cdot)[\cdot] = q_i(\tilde{h}_{it}(\cdot))[\cdot] : \Omega \times \mathcal{B} \to [0, 1]$ denote the transition probability that determines agent i's time t belief as a function of ω.

3 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 $(\Omega, \mathcal{B})$ 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 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.[4]

(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 $\mathcal{H}_{it}$.[5]

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 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 $\mathcal{H}_{it}$ is a consistent time t belief for agent i.

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

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

Definition 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 1. *σ^* is an equilibrium strategy profile if and only if for all i and t,*

$$\mathbb{E}^* \left[u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta) | \mathcal{H}_{it} \right] \geq \mathbb{E}^* \left[u_i(\tilde{\sigma}_{it}, \tilde{\sigma}^*_{-it}, \theta) | \mathcal{H}_{it} \right], \tag{2}$$

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

In the rest of the paper, 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 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 1. *The square integrable strategy profile σ^* is an equilibrium strategy profile if and only if for all i and t,*

$$\tilde{\sigma}^*_{it} = (1-\lambda)\mathbb{E}^* \left[\theta | \mathcal{H}_{it} \right] + \lambda \frac{1}{n-1} \sum_{j \neq i} \mathbb{E}^* \left[\tilde{\sigma}^*_{jt} | \mathcal{H}_{it} \right], \tag{3}$$

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 (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 1. *Suppose that Assumption 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 1 to analyze the asymptotic behavior of the agents' equilibrium actions.

4 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 2. *Suppose that Assumption 1 is satisfied. Let σ^* be a square integrable equilibrium strategy profile. Then, $\tilde{\sigma}^*_{it}$ converges to some $\mathcal{H}_i$-measurable random variable $\tilde{\varsigma}^*_i$ in the L^2 sense, that is,*

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

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.[6] We use this line of reasoning to prove the following theorem on consensus in the agents' payoffs and actions.

Theorem 1. *Suppose that Assumptions 1–3 are satisfied. Let σ^* be a square integrable equilibrium strategy profile. Then, as t goes to infinity, for all $i, j \in N$,*

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

[6] The Imitation Principle was first introduced by Gale and Kariv (2003) to study rational social learning with purely informational externalities. For another application of the Imitation Principle, see Rosenberg, Solan, and Vieille (2009).

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 1. Consider two agents who observe each others' actions at all time periods. The common prior is the uniform distribution over the set $\{-2,-1,1,2\}$. Agents 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 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 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

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

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 2—the agents generically coordinate on the action that is efficient given their aggregate information.

5 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. Repeated games of incomplete information of the type discussed in this paper 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). Yet, Theorem 1 of Section 4 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 3. Consider a single agent who repeatedly plays a game with payoffs as in (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 6. The private signals are *exogenously generated* if for any t, there exists some transition probability $\hat{\pi}_t$ 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 $\pi_t(h_t) = \hat{\pi}_t(\theta; s_1; s_2; \ldots; s_{t-1})$.

To simplify the analysis, we also replace Assumption 1 with Assumption 1′ below and replace Assumptions 2 and 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 $\bar{g}$ such that $\bar{g} = g_t$ for all t and with ν-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 $(\Xi, \mathcal{Z})$ be the measurable space with $\Xi = \Theta \times S^{\mathbb{N}}$ and $\mathcal{Z}$ the corresponding Borel σ-algebra. Any prior μ and signaling functions $\{\pi_t\}_{t\in\mathbb{N}}$ induce a probability distribution P over $(\Xi, \mathcal{Z})$ 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_i-valued random variable $\tilde{h}_{it} : \Xi \to A_i$. We define $\mathcal{I}_{it}^{\sigma}$ to be the σ-algebra generated by $\tilde{h}_{it}$ when agents follow the strategy profile σ, and define $\mathcal{I}_i^{\sigma}$ to be the σ-algebra generated by the union of $\mathcal{I}_{it}^{\sigma}$ 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 $(\Xi, \mathcal{Z})$. For instance, we have the following counterpart of Theorem 1, the proof of which is the same as the proof of Theorem 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) $E\Big[\big|u_i\left(\tilde{\sigma}_{it}^*, \tilde{\sigma}_{-it}^*, \theta\right) - u_j\left(\tilde{\sigma}_{jt}^*, \tilde{\sigma}_{-jt}^*, \theta\right)\big|\Big] \longrightarrow 0,$

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

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

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 μ 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 the paper we assume that $\mathbf{P}$ is endowed with the topology of uniform convergence (metrized by d).[7] We have the following result on the generic optimality of asymptotic actions.

Theorem 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,*

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

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

[7] 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*.

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 $\mathcal{I}^{\sigma^P}$. Note that $\mathcal{I}^{\sigma^P}$ captures the aggregate information that is collectively available to the agents at the end of the game. Therefore, $E\left[\theta|\mathcal{I}^{\sigma^{\mathcal{P}}}\right]$ 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 2 shows that the agents' equilibrium actions converge to this asymptotically optimal action in the L^1 sense.

An important special case of Theorem 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 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 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 of the game is infinite dimensional. It also extends Proposition 4 of Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012) to the case where the agents communicate their conditional estimates of the state, rather than their entire beliefs.

Finally, we remark that the assumption that the signal space is finite is somehow crucial. If the agents' signals belong to spaces that are higher dimensional than the action space, the efficient aggregation of information seems impossible as a large set of signals map to the same action, thus making it hard for the agents to infer the other agents' signals by observing their actions.

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. There are $n = 6$ agents over a fixed strongly connected social network playing the game with payoffs given by (1) with $\lambda = 2/3$. We consider two network topologies: a directed ring network depicted in Figure 1(a) and a star network depicted in Figure 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 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.[8] In this example, although the agents' signal spaces are not finite, convergence to the efficient action is achieved. [9]

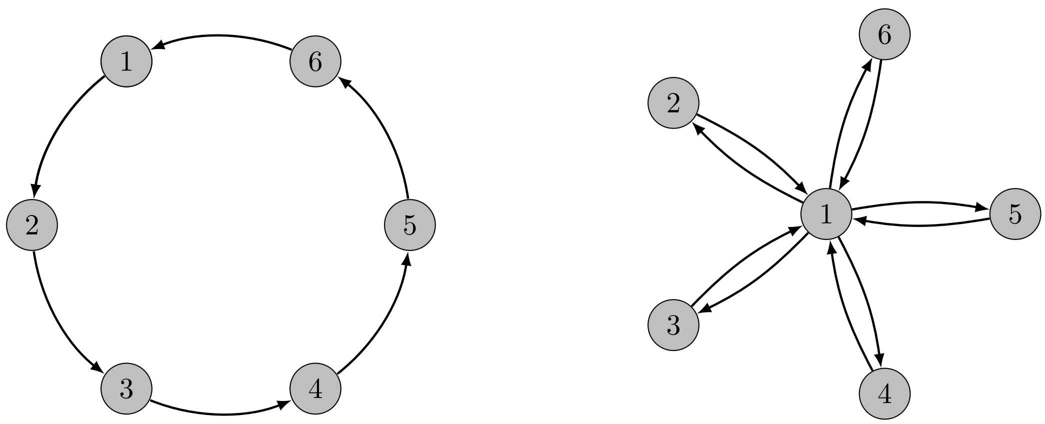

(a) The ring social network

(b) The star social network

Figure 1. The ring and star social networks

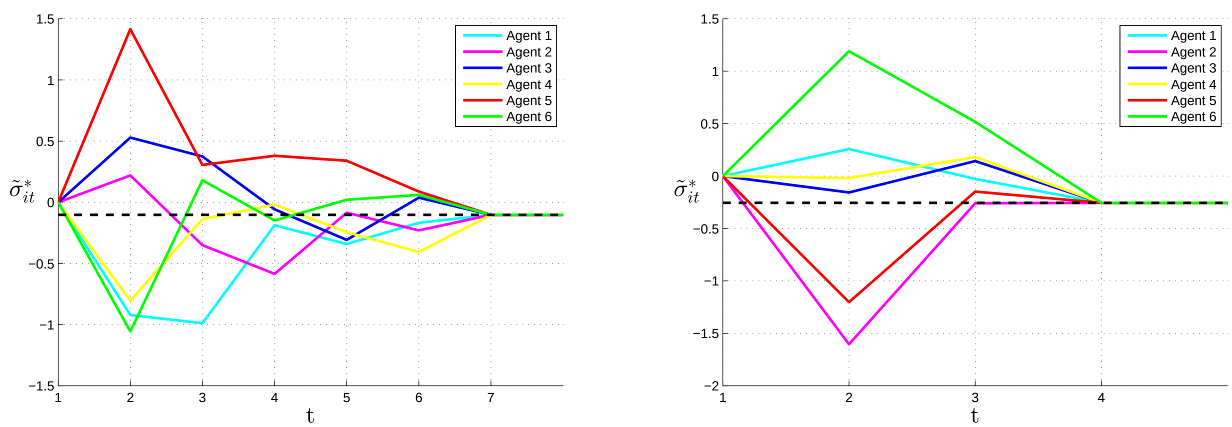

(a) Agents' actions given the ring network

(b) Agents' actions given the star network

Figure 2. Evolution of the agents' actions over time in Example 4

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

[9]For a recursive characterization of the agents' equilibrium actions in the Bayesian quadratic network games similar to the one studied in Example 4, see the complementary paper by the authors (Eksin et al. (2013)).

6 Conclusions

This paper studies a repeated 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 prove formal results regarding the asymptotic outcomes obtained when myopic agents play the actions prescribed by the weak perfect Bayesian equilibrium. In particular, we show 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.

Finally, we remark that although we proved the results assuming that the agents' preferences are represented by a quadratic utility function, the insights of our analysis do not hinge on the particular utility function used. In fact, similar results can be proved for more general coordination games with payoffs that satisfy some symmetry, concavity, and supermodularity conditions.

Appendix: Proofs

Proof of Lemma 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 $\mathcal{H}_{it}$, 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 $\mathcal{H}_{it}$, then (2) is equivalent to condition (a) of Definition 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

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

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 $\mathcal{H}_{it}$, then $\tilde{v}_{it}(\sigma_i, \sigma_{-i}; q^*_i)$ is a version of $\mathbb{E}^*\left[u_i(\tilde{\sigma}_t, \theta) | \mathcal{H}_{it}\right]$. □

Proof of Proposition 1

Before presenting the proof, we first introduce some notation and prove a technical lemma. Let $1 \leq p < \infty$, and let $(X, \mathcal{X}, P)$ 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 $\mathcal{L}^p(X, P)$. 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 $\|\cdot\|_p$:

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

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

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

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 2. *Let $(X, \mathcal{X}, P)$ 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 $\mathcal{X}_i \subseteq \mathcal{X}$ be a sub σ-algebra for any $i \in N$. Then, there exists a unique $f \in \mathbf{L}^2(X, P)$ such that*

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

for all $i \in N$.

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

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

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

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

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 $f \in \mathbf{L}^2(X, P)$. □

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

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

In particular, by Corollary 1, for any square integrable equilibrium strategy profile σ^*,

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

By Assumption 1 and Lemma 2, the above set of equations has a unique fixed-point in $\mathbf{L}^2(H_1, P_1)$. Consequently, (i) there exists a fixed-point $\tilde{\sigma}^*_1 = (\tilde{\sigma}^*_{11}, \ldots, \tilde{\sigma}^*_{n1})$ such that $\tilde{\sigma}^*_{i1} \in \mathcal{L}^2(H_1, P_1)$; 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 $\mathcal{H}_{i1}$-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 $(H_2, \mathcal{H}_2)$ induced by μ, ν, π_1 and π_2, and the time 1 profile $(\sigma^*_{11}, \ldots, \sigma^*_{n1})$ 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 $(H_2, \mathcal{H}_2)$. We can thus repeat the same argument to conclude that there exist functions $\sigma^*_{i2} : H_{i2} \to A_i$ such that $\sigma^*_{i2}(\tilde{h}_{i2}) = \tilde{\sigma}^*_{i2} \in \mathcal{L}^2(H_2, P_2)$ and

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

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 2

Consider the following system of equations:

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

By Lemma 2, the above set of equations has some solution $(\tilde{\varsigma}^*_i)_{i\in N}$, where $\tilde{\varsigma}^*_i \in \mathcal{L}^2(\Omega, \mathbb{P}^*)$. Moreover, by construction $\tilde{\varsigma}^*_i$ is $\mathcal{H}_i$-measurable. We prove the lemma by showing that $\tilde{\sigma}^*_{it} \to \tilde{\varsigma}^*_i$ in the L^2 sense as t goes to infinity. By Corollary 1,

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

Using the triangle inequality, we can conclude that

$$\begin{aligned}\|\tilde{\sigma}_{it}^* - \tilde{\varsigma}_i^*\|_2 \leq (1-\lambda)&\left\|\mathbb{E}^*[\theta|\mathcal{H}_{it}] - \mathbb{E}^*[\theta|\mathcal{H}_i]\right\|_2 \\ &+ \frac{\lambda}{n-1}\sum_{j\neq i}\left\|\mathbb{E}^*[\tilde{\sigma}_{jt}^* - \tilde{\varsigma}_j^*|\mathcal{H}_{it}]\right\|_2 \\ &+ \frac{\lambda}{n-1}\sum_{j\neq i}\left\|\mathbb{E}^*[\tilde{\varsigma}_j^*|\mathcal{H}_{it}] - \mathbb{E}^*[\tilde{\varsigma}_j^*|\mathcal{H}_i]\right\|_2. \end{aligned} \tag{4}$$

Since the conditional expectation is a contraction,

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

Summing (4) over i and using the above inequality imply

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

Hence,

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

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 1. Thus, by the L^p convergence theorem, $\mathbb{E}^*[\theta|\mathcal{H}_{it}]$ converges in the L^2 sense to $\mathbb{E}^*[\theta|\mathcal{H}_i]$.[10] 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,

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

□

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

Proof of Theorem 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 ν-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,

$$\left\|\tilde{\sigma}_{it}^\dagger - \tilde{\varsigma}_j^*\right\|_2 \leq \left\|\tilde{\sigma}_{it}^\dagger - \tilde{\sigma}_{j\tau}^*\right\|_2 + \left\|\tilde{\sigma}_{j\tau}^* - \tilde{\varsigma}_j^*\right\|_2 \tag{5}$$

We first use a truncation argument to bound the first term of (5). For $\tau \leq r < t$, let B_{rt} be the event defined as

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

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

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

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,[11]

$$\begin{aligned} \mathbb{E}^*\left[\left(\tilde{\sigma}_{it}^\dagger - \tilde{\sigma}_{j\tau}^*\right)^2\right] &= \sum_{r=\tau}^{t-1} \mathbb{E}^*\left[\left(\tilde{\sigma}_{it}^\dagger - \tilde{\sigma}_{j\tau}^*\right)^2 \middle| B_{rt}\right] \mathbb{P}^*(B_{rt}) + \mathbb{E}^*\left[\left(\tilde{\sigma}_{it}^\dagger - \tilde{\sigma}_{j\tau}^*\right)^2 \middle| D_{\tau t}\right] \mathbb{P}^*(D_{\tau t}) \\ &= \sum_{r=\tau}^{t-1} \mathbb{E}^*\left[\left(\tilde{\sigma}_{jr}^* - \tilde{\sigma}_{j\tau}^*\right)^2\right] \mathbb{P}^*(B_{rt}) + \mathbb{E}^*\left[\left(\tilde{\sigma}_{i\tau}^\dagger - \tilde{\sigma}_{j\tau}^*\right)^2\right] \mathbb{P}^*(D_{\tau t}), \end{aligned} \tag{6}$$

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

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

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

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 $\|\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau}\|_2 \leq \frac{\epsilon}{\sqrt{8}}$ for all $r \geq \tau$, implying that

$$\sum_{r=\tau}^{t-1} \mathbb{E}^* \left[\left(\tilde{\sigma}^*_{jr} - \tilde{\sigma}^*_{j\tau} \right)^2 \right] \mathbb{P}^*(B_{rt}) \leq \frac{\epsilon^2}{8}. \tag{7}$$

Next, we consider the second term of (6). By construction, $\tilde{\sigma}^\dagger_{i\tau} \in \{0\} \cup \{\tilde{\sigma}^*_{jr}\}_{r\in[1,\tau-1]}$. Thus,

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

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 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

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

Combining (7) and (8) with (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 (5). By Proposition 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 (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,

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

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

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

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

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^*_{it}, \tilde{\sigma}^*_{-it}, \theta)\right] \longrightarrow \mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta)\right] \qquad \text{as} \qquad t \to \infty. \tag{9}$$

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

$$\mathbb{E}^*\left[u_i(\tilde{\sigma}^\dagger_{it}, \tilde{\sigma}^*_{-it}, \theta)\right] \longrightarrow \mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_j, \tilde{\varsigma}^*_{-i}, \theta)\right] \qquad \text{as} \qquad t \to \infty. \tag{10}$$

On the other hand, since σ^* is an equilibrium strategy profile, by Lemma 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 (9) and (10),

$$\mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_i, \tilde{\varsigma}^*_{-i}, \theta)\right] \geq \mathbb{E}^*\left[u_i\left(\tilde{\varsigma}^*_j, \tilde{\varsigma}^*_{-i}, \theta\right)\right]. \tag{11}$$

Step three. By Assumption 3, there exists a sequence of agents $i_0, i_1, i_2, \ldots, i_n$ 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 ν-almost surely. For any k, thus by the result of step two,

$$\mathbb{E}^*\left[u_i(\tilde{\varsigma}^*_{i_k}, \tilde{\varsigma}^*_{-i_k}, \theta)\right] \geq \mathbb{E}^*\left[u_i\left(\tilde{\varsigma}^*_{i_{k+1}}, \tilde{\varsigma}^*_{-i_k}, \theta\right)\right]. \tag{12}$$

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

$$\sum_{k=1}^{n} \mathbb{E}^* \left[\tilde{\varsigma}^*_{i_k} \tilde{\varsigma}^*_{i_{k-1}} \right] \geq \sum_{k=1}^{n} \mathbb{E}^* \left[\left(\tilde{\varsigma}^*_{i_k} \right)^2 \right]. \tag{13}$$

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

$$\sum_{k=1}^{n} \mathbb{E}^* \left[\left(\tilde{\varsigma}^*_{i_k} \right)^2 \right] \geq \sum_{k=1}^{n} \mathbb{E}^* \left[\tilde{\varsigma}^*_{i_k} \tilde{\varsigma}^*_{i_{k-1}} \right], \tag{14}$$

with equality if and only if $\tilde{\varsigma}^*_{i_k} = \tilde{\varsigma}^*_{i_{k-1}}$ for all k, $\mathbb{P}^*$-almost surely; equation (13) implies that (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 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 (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.[12] 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 1, we can conclude that $\tilde{\sigma}^*_{it} - \mathbb{E}^*[\theta | \mathcal{H}_{it}]$ goes to zero in the L^2 sense. On the other hand, since $\mathcal{H}_{it} \uparrow \mathcal{H}_i$, we have that $\mathbb{E}^*[\theta | \mathcal{H}_{it}]$ converges to $\mathbb{E}^*[\theta | \mathcal{H}_i]$ in the L^1 sense. Therefore, $\tilde{\sigma}^*_{it}$ converges to $\mathbb{E}^*[\theta | \mathcal{H}_i]$ in the L^1 sense. Another use of the result of part (b) completes the proof. □

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

Proof of Theorem 2

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

Lemma 3. *Let $(X, \mathcal{B})$ be a measurable space, and let $(\mathbf{P}, d)$ be the metric space where $\mathbf{P}$ is the collection of all probability measures on $(X, \mathcal{B})$ and d is the total variation distance. Let $\mathcal{F}_1$ and $\mathcal{F}_2$ be two arbitrary sub σ-algebras of $\mathcal{B}$, let $\mathcal{F}$ be the σ-algebra generated by the union of $\mathcal{F}_1$ and $\mathcal{F}_2$, and let f be an arbitrary bounded random variable. The set*

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

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

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

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

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

We also have a similar equality for $\mathcal{F}_2$. 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,

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

We also have a similar equality for $\mathcal{F}_2$. 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

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

$\mathcal{F}_1$ and $\mathcal{F}_2$ 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 $\sigma(\Pi) = \sigma(\mathcal{F}_1 \cup \mathcal{F}_2) = \mathcal{F}$.

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

$$\mathbf{R} = \big\{P \in \mathbf{P} : E_P[f|\mathcal{F}_1] = E_P[f|\mathcal{F}_2]\big\}.$$

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 $\mathcal{U}$ 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 $\mathcal{U}$, 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 $A_1 \in \mathcal{F}_1$ and $A_2 \in \mathcal{F}_2$ be arbitrary sets with $C = A_1 \cap A_2$. Since $A_1 \in \mathcal{F}_1$, by the definition of conditional expectation, for all $Q \in b_\epsilon$,

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

Therefore,

$$\int_{A_1 \setminus C} f dQ + \int_C f dQ = \int_{A_1 \setminus C} E_Q[f|\mathcal{F}_1] dQ + \int_C E_Q[f|\mathcal{F}_1] dQ. \tag{15}$$

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, $E_{Q_k}[f|\mathcal{F}_1]$ converges in Q-probability to $E_Q[f|\mathcal{F}_1]$.[13] Therefore, since f is bounded and Q_k converges in total variation distance to Q,

$$\int_{A_2} E_{Q_k}[f|\mathcal{F}_1] dQ_k \longrightarrow \int_{A_2} E_Q[f|\mathcal{F}_1] dQ, \tag{16}$$

and

$$\int_{A_2} f dQ_k \longrightarrow \int_{A_2} f dQ. \tag{17}$$

Moreover, for all k,

$$\int_{A_2} f dQ_k = \int_{A_2} E_{Q_k}[f|\mathcal{F}_2] dQ_k = \int_{A_2} E_{Q_k}[f|\mathcal{F}_1] dQ_k, \tag{18}$$

where the first equality is by the definition of conditional expectation and the assumption that $A_2 \in \mathcal{F}_2$, and the second equality is a consequence of the fact that $Q_k \in \mathbf{R}$. Equations (16)–(18) imply that

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

And hence,

$$\int_{A_2 \setminus C} f dQ + \int_C f dQ = \int_{A_2 \setminus C} E_Q[f|\mathcal{F}_1] dQ + \int_C E_Q[f|\mathcal{F}_1] dQ. \tag{19}$$

We use (15) and (19) to conclude that $\int_C f dQ = \int_C E_Q[f|\mathcal{F}_1] dQ$ 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 $\hat{Q}_\delta$ over $(X, \mathcal{B})$ as follows: for any $B \in \mathcal{B}$,

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

It is easy to verify that $\hat{Q}_\delta$ is indeed a probability measure. We next show that $E_{\hat{Q}_\delta}[f|\mathcal{F}_1] = E_Q[f|\mathcal{F}_1]$.

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

Let $B \in \mathcal{F}_1$ be arbitrary.

$$\begin{aligned}
\int_B f d\hat{Q}_\delta &= \int_{B\cap A_1} f d\hat{Q}_\delta + \int_{B\cap A_1^c} f d\hat{Q}_\delta \\
&= (1+\delta Q(A_1^c)) \int_{B\cap A_1} f dQ + (1-\delta Q(A_1)) \int_{B\cap A_1^c} f dQ \\
&= (1+\delta Q(A_1^c)) \int_{B\cap A_1} E_Q[f|\mathcal{F}_1] dQ + (1-\delta Q(A_1)) \int_{B\cap A_1^c} E_Q[f|\mathcal{F}_1] dQ \\
&= \int_{B\cap A_1} E_Q[f|\mathcal{F}_1] d\hat{Q}_\delta + \int_{B\cap A_1^c} E_Q[f|\mathcal{F}_1] d\hat{Q}_\delta \\
&= \int_B E_Q[f|\mathcal{F}_1] d\hat{Q}_\delta, \qquad (20)
\end{aligned}$$

where the third equality follows from the assumption that $E_Q[f|\mathcal{F}_1]$ is a conditional expectation of f given $\mathcal{F}_1$ and the fact that $B \cap A_1 \in \mathcal{F}_1$ and $B \cap A_1^c \in \mathcal{F}_1$. Since $E_Q[f|\mathcal{F}_1]$ is $\mathcal{F}_1$-measurable, equation (20) proves that $E_Q[f|\mathcal{F}_1]$ is a version of $E_{\hat{Q}_\delta}[f|\mathcal{F}_1]$. Let $B_1 = A_1 \setminus C$ and $B_2 = A_2 \setminus C$. Equations (15) and (19) imply that

$$\int_{B_1} [f - E_Q[f|\mathcal{F}_1]]\, dQ = \int_{B_2} [f - E_Q[f|\mathcal{F}_1]]\, dQ. \qquad (21)$$

Since $B_1 \cap A_1 = B_1$,

$$\int_{B_1} \left[f - E_{\hat{Q}_\delta}[f|\mathcal{F}_1]\right] d\hat{Q}_\delta = (1+\delta Q(A_1^c)) \int_{B_1} [f - E_Q[f|\mathcal{F}_1]]\, dQ. \qquad (22)$$

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

$$\int_{B_2} \left[f - E_{\hat{Q}_\delta}[f|\mathcal{F}_1]\right] d\hat{Q}_\delta = (1-\delta Q(A_1)) \int_{B_2} [f - E_Q[f|\mathcal{F}_1]]\, dQ. \qquad (23)$$

On the other hand, if δ is sufficiently small, $\hat{Q}_\delta \in b_\epsilon$. Therefore, by (15) and (19),

$$\int_{B_1} \left[f - E_{\hat{Q}_\delta}[f|\mathcal{F}_1]\right] d\hat{Q}_\delta = \int_{B_2} \left[f - E_{\hat{Q}_\delta}[f|\mathcal{F}_1]\right] d\hat{Q}_\delta. \qquad (24)$$

Equations (21)–(24) imply that

$$\int_{B_1} [f - E_Q[f|\mathcal{F}_1]]\, dQ = \int_{B_2} [f - E_Q[f|\mathcal{F}_1]]\, dQ = 0. \qquad (25)$$

Thus, by (15),

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

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

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

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, $\sigma(\Pi) = \mathcal{F} \subseteq \Lambda_Q$ for $Q \in b_\epsilon$; that is, for any $A \in \mathcal{F}$,

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

Together with the fact that $E_Q[f|\mathcal{F}_1]$ and $E_Q[f|\mathcal{F}_2]$ are both measurable with respect to $\mathcal{F}$, this shows that $E_Q[f|\mathcal{F}] = E_Q[f|\mathcal{F}_1] = E_Q[f|\mathcal{F}_2]$ for all $Q \in b_\epsilon$. Thus, b_ϵ and $\mathbf{Q}$ are disjoint. Recall that the set $u \in \mathcal{U}$ was arbitrary. Therefore, for any set u in $\mathcal{U}$, 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 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 $P \in \mathbf{R}$,

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

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

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

where $\mathcal{I}_{ij}^{\sigma^P}$ is the σ-algebra generated by the union of $\mathcal{I}_i^{\sigma^P}$ and $\mathcal{I}_j^{\sigma^P}$. 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.

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

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 $\mathcal{I}_i^{\sigma^P}$ is a sub σ-algebra of $\mathcal{S}_t$, the Borel σ-algebra generated by the signal profiles up to time t. Given two arbitrary sub σ-algebras $\mathcal{S}_{it}, \mathcal{S}_{jt} \subseteq \mathcal{S}_t$ and the σ-algebra generated by their union $\mathcal{S}_{ijt}$, define

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

$\mathbf{Q}_{ijt}$ is a subset of the union of the above sets over all σ-algebra pairs $\mathcal{S}_{it}, \mathcal{S}_{jt}$. Since S is finite and for $P \in \mathbf{D}_t$ the signals are constant after time t, the σ-algebra $\mathcal{S}_t$ 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, $\mathcal{S}_{it}, \mathcal{S}_{jt}$ are two arbitrary fixed sub σ-algebras of the Borel σ-algebra of $\Theta \times S^t$. Therefore, we can directly use Lemma 3 to conclude that $\mathbf{S}_{ijt}(\mathcal{S}_{it}, \mathcal{S}_{jt})$ 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 $\mathbf{P}$.

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

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

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063.

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), "Bayesian learning in social networks." *The Review of Economic Studies*, 78, 1201–1236.

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.

Banerjee, Abhijit (1992), "A simple model of herd behavior." *The Quarterly Journal of Economics*, 107, 797–817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy*, 992–1026.

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*.

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.

Crimaldi, Irene and Luca Pratelli (2005), "Two inequalities for conditional expectations and convergence results for filters." *Statistics & probability letters*, 74, 151–162.

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 e-prints, URL http://arxiv.org/abs/1302.0249. arXiv:1302.0249.

Foster, Dean P. and H. Peyton Young (2003), "Learning, hypothesis testing, and Nash equilibrium." *Games and Economic Behavior*, 45, 73–96.

Gale, Douglas and Shachar Kariv (2003), "Bayesian learning in social networks." *Games and Economic Behavior*, 45, 329–346.

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.

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.

Keynes, John M. (1936), *The General Theory of Interest, Employment and Money*. London: Macmillan.

Laclau, Marie (2012), "A folk theorem for repeated games played on a network." *Games and Economic Behavior*, 76, 711–737.

Landers, Dieter and Lothar Rogge (1976), "Convergence of conditional p-means given a σ-lattice." *The Annals of Probability*, 4, 147–150.

Marden, Jason R., Gürdal Arslan, and Jeff S. Shamma (2009), "Cooperative control and potential games." *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39, 1393–1407.

Morris, Stephen and Hyun Song Shin (2002), "Social value of public information." *The American Economic Review*, 92, 1521–1534.

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.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2009), "Informational externalities and emergence of consensus." *Games and Economic Behavior*, 66, 979–994.

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.