Operations Research
Vol. 64, No. 3, May–June 2016, pp. 605–621
ISSN 0030-364X (print) | ISSN 1526-5463 (online)
http://dx.doi.org/10.1287/opre.2015.1381

Learning to Coordinate in Social Networks

Pooya Molavi
Department of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, mpooya@mit.edu

Ceyhun Eksin
School of Electrical and Computer Engineering and School of Biology, Georgia Inst. of Technology, Atlanta, Georgia 30332, ceyhuneksin@gatech.edu

Alejandro Ribeiro, Ali Jadbabaie
Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania 19104
{aribeiro@seas.upenn.edu, jadbabai@seas.upenn.edu}

We study a dynamic game in which short-run players repeatedly play a symmetric, strictly supermodular game whose payoffs depend on a fixed unknown state of nature. Each short-run player inherits the beliefs of his immediate predecessor in addition to observing the actions of the players in his social neighborhood in the previous stage. Because of the strategic complementary between their actions, players have the incentive to coordinate with others and learn from them. We show that in any Markov Bayesian equilibrium of the game, players eventually reach consensus in their actions. They also asymptotically receive similar payoffs despite initial differences in their access to information. We further show that, if the players' payoffs can be represented by a quadratic function, then the private observations are optimally aggregated in the limit for generic specifications of the game. Therefore, players asymptotically coordinate on choosing the best action given the aggregate information available throughout the network. We provide extensions of our results to the case of changing networks and endogenous private signals.

Keywords: consensus; information aggregation; supermodular games; networks.
Subject classifications: games/group decisions: noncooperative; networks/graphs: theory; technology.
Area of review: Special Issue on Information and Decisions in Social and Economic Networks.
History: Received July 2013; revisions received November 2014, February 2015; accepted February 2015. Published online in *Articles in Advance* November 30, 2015.

1. Introduction

Rational agents with aligned interests best served by coordinating their actions may nevertheless fail to do so if they disagree on the best course. A player who is sufficiently convinced of an action being the one resulting in the highest payoff will be willing to take that action in isolation, even at the expense of forgoing the payoff from acting in harmony with other players. This nonconforming action will serve as a strong signal for the players who play the same game in the future that the payoff-maximizing action could be different than what was chosen by the majority of players. But can future generations of players use this information to improve their payoffs? Is the unique insight of the nonconforming player lost if only a handful of other players observe his action, or will it eventually reach the entire population? How will the conclusions be different if each player of each generation has a unique insight about the game being played? This paper introduces a dynamic model of coordination games with asymmetric information to provide answers to these questions.

The scenario described above could represent a society wherein an informed leader's actions have the potential to change the prevailing social norm, or the market for a new technology in which adoption by an informed user can serve as signal of his belief in the future of the technology. Coordination models similar to ours have been used to study a wide-ranging set of phenomena including conventions (Shin and Williamson 1996), social norms and the rule of law (Acemoglu and Jackson 2015, 2014), currency runs (Obstfeld 1996), regime change (Angeletos et al. 2007), markets with externalities (Morris and Shin 2002; Angeletos and Pavan 2007, 2009), and Keynsian coordination failures (Cooper and John 1988), among others. Cooper (1999) contains additional examples of the applications of coordination games with asymmetric information in macroeconomics.

To answer the question posed above, we study a supermodular game of incomplete information played by generations of short-run players.[1] In every period, a new generation of players are born who live for one period. The players of each generation play a supermodular game with payoffs that depend on a fixed unknown state of the world. Each player inherits the belief of a player from the previous generation—the player previously occupying his *role*—and observes the actions of some of the players of the last generation—his *neighbors*. The players then simultaneously choose actions to maximize their payoffs given the information available to them.

Myopic behavior by short-run players is a good approximation to individuals' rational behavior in the examples mentioned above where we have a large number of small players, each having a negligible impact on the entire society. We have in mind a small customer deciding whether to purchase a product, a citizen deciding whether to follow a norm, or a protester deciding whether to join a protest. In all of these examples, each individual can ignore the effect of his current action on the actions of the individuals he encounters in the future. Alternatively, one can think of each role as a dynasty with each short-run player representing a member of the dynasty that has access to the entire history of the dynasty but who only makes a single decision.

We restrict attention to Markovian strategies in which the players' actions are not functions of the history of the game but only of their information. We use pure strategy *Markov Bayesian equilibrium* (MBE) as the solution concept.[2] Players choose actions that maximize their instantaneous payoffs given others' strategies and their beliefs, and the beliefs are consistent with the equilibrium stategies and the application of Bayes' rule.

We show that, if the social network is connected, players eventually reach consensus both in their actions and their payoffs despite occupying roles with asymmetric initial information about the fundamental state of the world. In other words, although players in initial generations might disagree on the best course of action, future generations cannot disagree in the long run. This result is similar in spirit to the argument presented by Aumann (1976) that Bayesian agents who share a prior cannot "agree to disagree." The key intuition for why this result holds is that the imitation principle applies to our setting.[3] According to the imitation principle, the mere fact that, in equilibrium, no player (he) wishes to deviate by imitating the action of a player (she) whose play he observes infinitely often is evidence that he believes that his equilibrium action results in a higher payoff. The imitation principle imposes restrictions on the equilibrium beliefs that can be leveraged to rule out strategies according to which two players in two roles that frequently observe each others' actions continue to miscoordinate.

In §4 we restrict our attention to quadratic symmetric supermodular games and develop sharper results. Because of their tractability, quadratic games are studied extensively in the recent literature.[4] We show that, if the information structure of each player can be represented by a finite partition of the state space, the asymptotic consensus action is generically the action players would have chosen if they had been able to directly pool their information at the beginning of the game. This result shows that, in a quadratic supermodular game of incomplete information, failure of information aggregation is nongeneric. We also prove that consensus in actions and payoffs continue to hold even if the network is random, directed, and time-varying, or if the players observe a stream of signals over time whose distribution depends on the previous actions of players.

Our assortment of results suggests that consensus is a ubiquitous phenomenon in games of strategic complementarity with a common prior. They can be interpreted as reinforcing the idea presented by Aumann (1976) that Bayesian agents cannot disagree forever. Aumann's argument was presented in a setting with no interaction among players other than the sharing of beliefs. Our results suggest that the conclusion that Bayesian players cannot agree to disagree is robust to the introduction of strategic interactions, as long as players' actions are strategic complements.

Our result on information aggregation shows that, when the utilities are quadratic, consensus generically implies optimal information aggregation. Similar results, with more stringent requirements, have already appeared in the literature; however, a general result has been lacking. In particular, Mossel et al. (2011) propose conditional independence of private signals as a sufficient condition for consensus to imply information aggregation. Ostrovsky (2012) presents a form of conditional independence called "separability" as a sufficient condition for information aggregation in a Kyle (1985)-style market. Our generic information aggregation result, on the other hand, does not rely on any conditional independence assumptions. Independently, in a recent paper, Arieli and Mueller-Frank (2015) prove generic information aggregation results in a more abstract setting.

In summary, we make three distinct sets of contributions: First, we contribute to the theory of symmetric supermodular games by showing that short-run players reach consensus on their payoffs and actions, even when the equilibrium is not unique. Our second contribution is to the social learning literature, where we show that presence of payoff externalities of the complementary nature will not hamper learning in societies that are sufficiently connected over time. Third, we show that in such a coordination setup information is generically aggregated. We present a set of examples from a variety of application domains, ranging from economics to engineering, in addition to a set of distinctive examples that are meant to illustrate the tightness of our results.

Related Literature. In addition to the papers already mentioned, our paper is related to three lines of research in game theory. The first is the literature on Bayesian learning over networks. The focus of this literature is on modeling the way agents use their observations to update their beliefs and characterizing the outcomes of the learning process. Examples include Borkar and Varaiya (1982), Banerjee (1992), Bikhchandani et al. (1992), Bala and Goyal (1998), Smith and Sørensen (2000), Gale and Kariv (2003), Çelen and Kariv (2004), Rosenberg et al. (2009), Mossel and Tamuz (2010), Mossel et al. (2011), Acemoglu et al. (2011), Mueller-Frank (2013), Lobel and Sadler (2015a, b), and Acemoglu et al. (2014). In this paper, we extend the Bayesian social learning framework to an environment with payoff externalities where each short-run player's stage payoff is a function of other players' 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. In the current paper in contrast, the focus is on whether agents in a network asymptotically reach consensus and whether they aggregate the dispersed information.

Finally, our paper is related to the literature on information aggregation in markets. The focus of this literature is on characterizing the conditions under which prices correctly aggregate the information dispersed throughout the market. Examples include Wolinsky (1990), Foster and Viswanathan (1996), Vives (2010), Lauermann and Wolinsky (2015), Amador and Weill (2012), Rostek and Weretka (2012), Ostrovsky (2012), and Bonatti et al. (2015). In our paper the players do not observe public signals, such as prices, but rather make local observations about the actions of the players in their social neighborhood.

Organization of the Paper. The rest of the paper is organized as follows. Section 2 presents the baseline model. Section 3 presents our main result on consensus in supermodular games. In §4 we specialize to a quadratic supermodular game and present our result on the genericity of information aggregation. We also present our results on consensus with endogenous signals and time-varying networks. Section 5 discusses the logic and implications of our results in more detail. In §6 we provide some applications of symmetric supermodular games in engineering and economics. All the proofs that are omitted from previous sections are provided in §7. Section 8 contains our concluding remarks.

2. Model

Throughout, we use the usual order and the standard topology on $\mathbb{R}$. Products of topological spaces are equipped with the product topology. All topological spaces are endowed with the Borel sigma-algebra. Two measurable mappings are said to be equal if they have the same domain and codomain and agree almost everywhere. Given a probability distribution P over a measurable space $X \times Y$, $\text{marg}_X P$ denotes the marginal distribution of P over X. Given sets $X_1, \ldots, X_n$, we use X to denote $\times_{i=1}^n X_i$ with generic element x and use X_{-i} to denote $\times_{j\neq i} X_i$ with generic element $x_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$.

2.1. The Game

Consider n roles indexed by $i \in N = \{1, \ldots, n\}$. Role i represents a sequence of short-run players, each of whom plays only once. We refer to the short-run player at role i playing in stage t as player it. We refer to the collection of all short-run players in role i as "big" player i or simply player i.

At the beginning of the game, nature chooses the payoff-relevant state of the world θ from a compact metric space Θ. The players in a given role all observe the same noisy signal of θ. We denote by s_i the signal observed by the players in role i. We assume that s_i belongs to a countable set S_i that is endowed with the discrete topology. The realized state otherwise remains unknown to the players.

The game is played over a countable set of stages indexed by the set of positive integers $\mathbb{N}$. In the beginning of stage $t \in \mathbb{N}$, player it observes the actions chosen in the previous stages by a subset of big players, called i's neighbors and denoted by N_i. We use the convention that each i is his own neighbor. We further assume that the neighborhood relationship is symmetric: i is a neighbor of j if and only if j is a neighbor of i.

At the end of period t, player it chooses an action $a_{it} \in A_i$ simultaneously with other short-run players and receives payoff $u_i(a_t, \theta)$. We assume that A_i is a compact subset of $\mathbb{R}$ and u_i is continuous in all its arguments. We further assume that the game is symmetric: for all $i, j \in N$, $A_i = A_j$ and $u_i(a_t, \theta) = u_j(a'_t, \theta)$ if $a_{it} = a'_{jt}$ and a_{-it} is a permutation of a'_{-jt}. Finally, we assume that $u_i(a_t, \theta)$ is strictly supermodular in a_t for all $\theta \in \Theta$. [A function $f\colon \mathbb{R}^n \to \mathbb{R}$ is *supermodular* if $f(\min\{x, y\}) + f(\max\{x, y\}) \geqslant f(x) + f(y)$ for all $x, y \in \mathbb{R}^n$, where $\min(\{x, y\})$ denotes the componentwise minimum and $\max(\{x, y\})$ denotes the componentwise maximum of x and y. The function is *strictly supermodular* if the inequality is strict for any incomparable pair of vectors x and y.[5]]

We summarize the players' uncertainty about the exogenous variables by some ω belonging to the measurable space $(\Omega, \mathcal{B})$, where $\Omega = \Theta \times S$ and $\mathcal{B}$ is the Borel sigma-algebra. Note that the canonical projection $s_i\colon \Omega \to S_i$ is continuous and therefore measurable. We assume that the payoff-relevant state θ and the private signals are jointly distributed according to some probability distribution $\mathbb{P}$ over $(\Omega, \mathcal{B})$ and that this is common knowledge. The expectation operator corresponding to $\mathbb{P}$ is denoted by $\mathbb{E}$.

We restrict our attention to Markovian strategies according to which the players' actions depend on the history of the game only to the extent that it is informative of the payoff-relevant state of the world. In particular, we define the players' strategies and information as follows. Let $\mathcal{H}_{i1}$ be the smallest sub-sigma algebra of $\mathcal{B}$ that makes s_i measurable. $\mathcal{H}_{i1}$ captures the information available to player $i1$. A Markovian strategy for player $i1$ is a mapping $\sigma_{i1}\colon \Omega \to A_i$ which is measurable with respect to $\mathcal{H}_{i1}$. For $t \geqslant 2$ define $\mathcal{H}_{it}^{\sigma^{t-1}}$ and σ_{it} recursively as follows: Denote by $\sigma^{t-1} = (\sigma_1, \sigma_2, \ldots, \sigma_{t-1})$, where $\sigma_\tau = (\sigma_{1\tau}, \ldots, \sigma_{n\tau})$, the Markovian strategy profile followed by the short-run players that are active before stage t. Given σ^{t-1}, the information available to player it is captured by $\mathcal{H}_{it}^{\sigma^{t-1}}$, the smallest sub-sigma algebra of $\mathcal{B}$ that makes s_i and $\{\sigma_{j1}, \ldots, \sigma_{j,t-1}\}_{j\in N_i}$ measurable. A Markovian strategy for player *it* is a mapping $\sigma_{it}\colon \Omega \to A_i$ that is measurable with respect to $\mathcal{H}_{it}^{\sigma^{t-1}}$. We let $\sigma = (\sigma_1, \sigma_2, \ldots)$ denote a Markovian strategy profile generated as above and let $\mathcal{H}_{i\infty}^{\sigma} = \bigvee_{t=1}^{\infty} \mathcal{H}_{it}^{\sigma^{t-1}}$ to be the information available to the players in role i "at the end of the game" given that players follow strategy profile σ. Note that, for any strategy profile σ and all i, $\mathcal{H}_{it}^{\sigma^{t-1}} \subseteq \mathcal{H}_{it'}^{\sigma^{t'-1}}$ if $t \leqslant t'$. Whenever there is no risk of confusion we use $\mathcal{H}_{it}^{\sigma}$ to mean $\mathcal{H}_{it}^{\sigma^{t-1}}$.

2.2. Equilibrium

Definition 1. A Markovian strategy profile σ is a *Markov Bayesian equilibrium* (MBE) if for all i, t, and $\mathcal{H}_{it}^{\sigma_{t-1}}$-measurable mappings $\sigma'_{it}\colon \Omega \to A_i$,

$$\mathbb{E}[u_i(\sigma_{it}, \sigma_{-it}, \theta) \mid \mathcal{H}_{it}^{\sigma_{t-1}}] \geqslant \mathbb{E}[u_i(\sigma'_{it}, \sigma_{-it}, \theta) \mid \mathcal{H}_{it}^{\sigma_{t-1}}].$$

According to our equilibrium notion, the short-run players who are active in stage t choose an interim pure-strategy Bayesian Nash equilibrium of a Bayesian game in which their information is induced by the equilibrium strategies of the short-run players that played before them.

Proposition 1. *A MBE σ exists.*

Proof. The proof involves repeated use of Theorem 23 of Van Zandt (2010). The game played by the short-run players in the first stage is a Bayesian supermodular game that satisfies the conditions of Theorem 23 of Van Zandt. Therefore, it has an interim pure-strategy Bayesian Nash equilibrium denoted by $\sigma_1 = (\sigma_{11}, \ldots, \sigma_{n1})$. Let $\mathcal{H}_{i2}^{\sigma^1}$ denote the smallest sub-sigma algebra of $\mathcal{B}$ that makes s_i and $\{\sigma_{j1}\}_{j\in N_i}$ measurable. The sigma-algebras $\mathcal{H}_{12}^{\sigma^1}, \ldots, \mathcal{H}_{n2}^{\sigma^1}$ define a Bayesian supermodular game in the second stage, which has an interim pure-strategy Bayesian Nash equilibrium σ_2. Repeating this argument inductively, we can construct an MBE $\sigma = (\sigma_1, \sigma_2, \ldots)$. □

2.3. Remarks on the Model

The model considers repeated interactions among rational short-run players in given roles. A role can represent a myopic individual with each short-run player representing the individual's one-time decision. Alternatively, a role can represent a dynasty with each short-run player representing a member of the dynasty that has access to the entire history of the dynasty but only makes a single decision.

Players' behavior is determined by an MBE strategy profile. The short-run players in the first generation are endowed with private information about the realized state of nature. The equilibrium strategy determines how players use their information to choose the actions that maximize their expected payoffs. The players that follow inherit the information of their predecessors. However, each member of successive generations also acquires information, through the observation of recent events in his social neighborhood, that was not available to his predecessors. The MBE imposes the usual requirements that the players' actions maximize their expected payoffs, and the information contained in new observations is incorporated in a fashion consistent with the equilibrium and the application of Bayes' rule.

MBE, in addition, restricts the players' strategies to be measurable with respect to the exogenous variables; hence the term Markovian. Players following Markovian strategies do not condition their actions on observations that are uninformative about the exogenous variables. Thus, the inference of player *it* about the actions of other players in his generation given his knowledge of the equilibrium strategies reduces to inference about the exogenous variables. In contrast, long-run players that follow non-Markovian strategies may experiment, try to build reputations, or punish other players based on past events.

We note that the players in our model only observe the actions of their neighbors and do not share their past experiences, signals, or beliefs with players in other roles. Furthermore, the players do not observe the realized payoffs of their predecessors. In §§4.2 and 5.3 we discuss an extension of the model to the case of observed payoffs.

Finally, strict supermodularity of the utility function captures the idea that it is in the players' interest to coordinate their actions: The sum of the utilities for players i and j when playing a_i and $a_j \neq a_i$, respectively, is less than the sum of the utility of i when player j moves to a_i (matching player i) and the utility of j when player i moves to a_j. This property of the utility function is key in proving our main result on consensus in the next section.

3. Main Result

Our main result states that short-run players asymptotically reach consensus when they act according to an MBE strategy profile. We discuss the implications of the results presented here in §5.1.

Theorem 1. *Let σ be an MBE. For all $i, j \in N$, $\sigma_{it} - \sigma_{jt} \to 0$, $\mathbb{P}$-almost surely, as t goes to infinity.*

Proof. We let $\mathcal{S}$ denote the smallest sub-sigma algebra of $\mathcal{B}$ that makes the mapping $\omega \mapsto s(\omega) = (s_1(\omega), \ldots, s_n(\omega))$ measurable, and let $\mathcal{H}_{\infty}^{\sigma} = \bigvee_{i=1}^{n} \mathcal{H}_{i\infty}^{\sigma}$. Since the information available to the players in any stage of the game is no more than the information jointly contained in their private signals, $\mathcal{H}_{\infty}^{\sigma} \subseteq \mathcal{S}$. Therefore, σ_{it} is measurable with respect to $\mathcal{S}$ for all i and t, so $\sigma_{it}(\omega) = \sigma_{it}(\omega')$ whenever $s(\omega) = s(\omega')$. We can thus define the mapping $\sigma_{it}\colon S \to A_i$, with some abuse of notation, by letting $\sigma_{it}(s) = \sigma_{it}(\omega(s))$, where $\omega(s)$ is a selection of $\Omega(s) = \{\omega \in \Omega\colon s(\omega) = s\}$. The statement of the theorem is therefore equivalent to the following: $\sigma_{it}(s) - \sigma_{jt}(s) \to 0$ for all $s \in S$ with $\mathbb{P}(s) = \mathbb{P}(\Theta \times \{s\}) > 0$.

Suppose, to the contrary, that there exists some neighboring $i, j \in N$, some $s_0 \in S$ with $\mathbb{P}(s_0) > 0$, and a divergent sequence $\{k_{0t}\}_{t\in\mathbb{N}}$ such that $|\sigma_{ik_{0t}}(s_0) - \sigma_{jk_{0t}}(s_0)|$ is uniformly bounded away from zero. Since S is countable, there exists an enumeration $s_1, s_2, \ldots$ of S. Since A is a compact metric space, there exists a further subsequence $\{k_{1t}\}_{t\in\mathbb{N}}$ of $\{k_{0t}\}_{t\in\mathbb{N}}$ such that the sequence $\{\sigma_{k_{1t}}(s_1)\}_{t\in\mathbb{N}}$ is convergent. Likewise, there exists a further subsequence $\{k_{2t}\}_{t\in\mathbb{N}}$ of $\{k_{1t}\}_{t\in\mathbb{N}}$ such that the sequence $\{\sigma_{k_{2t}}(s_2)\}_{t\in\mathbb{N}}$ is convergent, and by induction, for $m\in\mathbb{N}$, there exists a further subsequence $\{k_{m+1,t}\}_{t\in\mathbb{N}}$ of $\{k_{mt}\}_{t\in\mathbb{N}}$ such that the sequence $\{\sigma_{k_{m+1,t}}(s_{m+1})\}_{t\in\mathbb{N}}$ is convergent. Construct the sequence $\{l_t\}_{t\in\mathbb{N}}$ by letting $l_t = k_{tt}$. For all $s\in S$, as t goes to infinity $\sigma_{l_t}(s)$ converges to some $\sigma_\infty(s)\in A$ with $\sigma_{i\infty}(s_0)\neq\sigma_{j\infty}(s_0)$. With slight abuse of notation, define the measurable mapping $\sigma_\infty\colon \Omega\to A$ by letting $\sigma_\infty(\omega)=\sigma_\infty(s(\omega))$. Since u_i is continuous and A and Θ are compact, by the dominated convergence theorem,

$$\mathbb{E}[u_i(\sigma_{l_t},\theta)]\to\mathbb{E}[u_i(\sigma_\infty,\theta)].$$

Define the $\mathcal{H}^{\sigma}_{il_t}$-measurable mapping $\sigma'_{il_t}\colon \Omega\to A_i$ as $\sigma'_{il_t}=\sigma_{jl_{t-1}}$. This mapping constitutes a feasible strategy for player il_t according to which he imitates the action chosen by player jl_{t-1}. By construction, $(\sigma'_{il_t},\sigma_{-il_t})\to(\sigma_{j\infty},\sigma_{-i\infty})$ for all $\omega\in\Omega$. Thus,

$$\mathbb{E}[u_i(\sigma'_{il_t},\sigma_{-il_t},\theta)]\to\mathbb{E}[u_i(\sigma_{j\infty},\sigma_{-i\infty},\theta)].$$

Since σ is an equilibrium, $\mathbb{E}[u_i(\sigma_{l_t},\theta)]\geqslant\mathbb{E}[u_i(\sigma'_{il_t},\sigma_{-il_t},\theta)]$ for all $t\in\mathbb{N}$, so

$$\mathbb{E}[u_i(\sigma_{i\infty},\sigma_{-i\infty},\theta)]\geqslant\mathbb{E}[u_i(\sigma_{j\infty},\sigma_{-i\infty},\theta)]. \tag{1}$$

By a similar argument,

$$\mathbb{E}[u_j(\sigma_{j\infty},\sigma_{-j\infty},\theta)]\geqslant\mathbb{E}[u_j(\sigma_{i\infty},\sigma_{-j\infty},\theta)]. \tag{2}$$

Let $u(a_i;a_j,a_{-ij},\theta)$ denote the utility of a player in role i when he chooses a_i, player j chooses a_j, and other players choose a_{-ij}. By the symmetry assumption, the payoff of a player in role j when player j chooses a_i, player i chooses a_j, and others choose a_{-ij} is also equal to $u(a_i;a_j,a_{-ij},\theta)$. Equations (1) and (2) thus can be written as

$$\mathbb{E}[u(\sigma_{i\infty};\sigma_{j\infty},\sigma_{-ij\infty},\theta)]\geqslant\mathbb{E}[u(\sigma_{j\infty};\sigma_{j\infty},\sigma_{-ij\infty},\theta)],$$
$$\mathbb{E}[u(\sigma_{j\infty};\sigma_{i\infty},\sigma_{-ij\infty},\theta)]\geqslant\mathbb{E}[u(\sigma_{i\infty};\sigma_{i\infty},\sigma_{-ij\infty},\theta)].$$

Summing the above equations,

$$\mathbb{E}[u(\sigma_{i\infty};\sigma_{j\infty},\sigma_{-ij\infty},\theta)+u(\sigma_{j\infty};\sigma_{i\infty},\sigma_{-ij\infty},\theta)]$$
$$\geqslant\mathbb{E}[u(\sigma_{i\infty};\sigma_{i\infty},\sigma_{-ij\infty},\theta)+u(\sigma_{j\infty};\sigma_{j\infty},\sigma_{-ij\infty},\theta)]. \tag{3}$$

On the other hand, since u is strictly supermodular, for all $a_i\in A_i$ and $a_j\in A_j$,

$$u(a_i;a_j,a_{-ij},\theta)+u(a_j;a_i,a_{-ij},\theta)$$
$$\leqslant u(a_i;a_i,a_{-ij},\theta)+u(a_j;a_j,a_{-ij},\theta), \tag{4}$$

with equality if and only if $a_i=a_j$. Equations (3) and (4) imply that $\sigma_{i\infty}=\sigma_{j\infty}$ for $\mathbb{P}$-almost all ω, contradicting the assumption that $\sigma_{i\infty}(s_0)\neq\sigma_{j\infty}(s_0)$ and $\mathbb{P}(s_0)>0$. □

An immediate corollary of consensus in strategies is asymptotic consensus in payoffs.

Corollary 1. *Let σ be a MBE. For all $i,j\in N$, $u_i(\sigma_t,\theta)-u_j(\sigma_t,\theta)\to 0$, $\mathbb{P}$-almost surely, as t goes to infinity.*

Proof. Define $\sigma_{it}\colon S\to A_i$ as in the proof of Theorem 1. It is sufficient to show that $u_i(\sigma_t(s),\theta)-u_j(\sigma_t(s),\theta)\to 0$ for all $\theta\in\Theta$ and $s\in S$ with $\mathbb{P}(s)>0$. Suppose to the contrary that there exists some neighboring $i,j\in N$, some $\theta_0\in\Theta$ and $s_0\in S$ with $\mathbb{P}(s_0)>0$, and a divergent sequence $\{k_{0t}\}_{t\in\mathbb{N}}$ such that $|u_i(\sigma_{k_{0t}}(s_0),\theta_0)-u_j(\sigma_{k_{0t}}(s_0),\theta_0)|$ is uniformly bounded away from zero. As in the proof of Theorem 1, we can construct a further subsequence $\{l_t\}_{t\in\mathbb{N}}$ of $\{k_{0t}\}_{t\in\mathbb{N}}$ such that for all $s\in S$, as t goes to infinity, $\sigma_{l_t}(s)$ converges to some $\sigma_\infty(s)\in A$. Furthermore, by Theorem 1, $\sigma_{i\infty}(s)=\sigma_{j\infty}(s)$ for all $i,j\in N$ and $s\in S$. Therefore, since u_i is continuous and symmetric, $u_i(\sigma_{l_t}(s_0),\theta_0)-u_j(\sigma_{l_t}(s_0),\theta_0)\to 0$ for all $i,j\in N$, contradicting the assumption that $|u_i(\sigma_{k_{0t}}(s_0),\theta_0)-u_j(\sigma_{k_{0t}}(s_0),\theta_0)|$ is uniformly bounded away from zero for some i,j. □

The above result also implies ex ante consensus in the expectation of big players' asymptotic payoffs: prior to the start of the game, players in all roles expect their successors to asymptotically achieve similar payoffs. In §5.1 we show by means of an example that players might disagree in their conditional expected payoffs even when they are receiving the same payoffs.

4. Quadratic Games

In this section we study an important special case of the baseline model introduced in §2 in which θ is a real number and each player's utility function is quadratic in both θ and the average action of other players. Namely, we assume that

$$u_i(a,\theta)=-(1-\lambda)(a_i-\theta)^2-\lambda(a_i-\bar{a}_{-i})^2, \tag{5}$$

where $\lambda\in(0,1)$ is a constant and $\bar{a}_{-i}=\sum_{j\neq i}a_j/(n-1)$ denotes the average action of other players.[6] The game with players' payoffs given by (5) is a strictly supermodular game since the payoff's cross partial derivative is equal to $\lambda>0$. The first term is a quadratic loss in the distance between the realized state and player i's action, capturing the player's preference for actions that are close to the unknown state. The second term is the "beauty contest" term representing the player'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 players who engage in second-guessing others' actions as postulated by Keynes. Restricting the utility function allows us to show optimal information aggregation and extend our results on consensus to endogenous signals and time-varying networks. The following proposition states an important property of the quadratic game.

Proposition 2. *The quadratic game has a unique MBE σ.*

The result is based on expressing the MBE as a sequence of stage game Bayesian Nash equilibria. In each stage t, the corresponding short-run player it's equilibrium strategy must be a best response to other players' strategies. Given the payoff in (5) and since λ belongs to $(0, 1)$, the players' best-response function is a contraction mapping. This, together with the compactness of the action space, implies that the stage game has a unique Bayesian Nash equilibium.[7] Since a MBE is a sequence of stage game Bayesian Nash equilibria and all stages have unique equilibria, the MBE must be unique.

4.1. Information Aggregation

Consider the model with quadratic payoffs. Suppose that the signal space S_i is finite. Define information, strategies, and Markov Bayesian equilibrium as in the baseline model. Let $\mathscr{H}^\sigma_\infty = \bigvee_{i=1}^n \mathscr{H}^\sigma_{i\infty}$ denote the information collectively available to the players at the end of the game when they follow the strategy profile σ.

Instead of considering a single probability measure over $(\Omega, \mathscr{B})$, in this subsection we prove a result about the set of all probability measures over $(\Omega, \mathscr{B})$. Let $\mathbf{P}$ denote the space of all probability measures over $(\Omega, \mathscr{B})$ with elements denoted by P, and let d denote the total variation distance between measures. $(\mathbf{P}, d)$ is a metric space and therefore a topological space. We say that a property holds *generically* if it is true for all P belonging to a residual subset of $(\mathbf{P}, d)$.[8] We have the following result on information aggregation.

Theorem 2. *Let σ^P denote the unique MBE of the quadratic game given prior $P \in \mathbf{P}$. For generic $P \in \mathbf{P}$ and all i, $\sigma^P_{it} - E^P[\theta \mid \mathscr{H}^{\sigma^P}_\infty] \to 0$, P-almost surely.*

The theorem states that for priors P belonging to a generic set of probability distributions, the players asymptotically play as if they all had the information captured by the sigma-algebra $\mathscr{H}^{\sigma^P}_\infty$, which is the aggregate information collectively available to the players at the end of the game.

4.2. Endogenous Signals and Time-Varying Directed Networks

Suppose that each short-run player it privately observes a signal s_{it} belonging to a complete separable metric space S_i. The distribution of s_{it} depends on the history of the game up to stage t.

Additionally, suppose that each short-run player it has a different and random set of neighbors. Denote the set of neighbors of player it by N_{it}. We maintain the assumption that $i \in N_{it}$ for all i and t. However, the neighborhood relationship is no longer assumed to be symmetric. Let $\mathscr{N}_i$ denote the space of all possible neighborhoods for a player in role i, and let $\mathscr{N} = \prod_{i=1}^n \mathscr{N}_i$.

We summarize the players' uncertainty by some ω belonging to the measurable space $(\Omega, \mathscr{B})$, where $\Omega = \Theta \times S^{\mathbb{N}} \times \mathscr{N}^{\mathbb{N}}$ and $\mathscr{B}$ is the Borel sigma-algebra. The payoff-relevant state, the signals, and the neighborhoods are jointly distributed according to some *endogenous* probability distribution.

In what follows, we recursively define the players' strategies and information for the model with endogenous signals and time-varying networks. Let $\mathscr{H}_{i1}$ be the smallest sub-sigma algebra of $\mathscr{B}$ that makes s_{i1} and N_{i1} measurable. A Markovian strategy for player $i1$ is a mapping $\sigma_{i1}\colon \Omega \to A_i$ which is measurable with respect to $\mathscr{H}_{i1}$. For $t \geqslant 2$, define $\mathscr{H}^{\sigma^{t-1}}_{it}$ and σ_{it} recursively as follows. Denote by $\sigma^{t-1} = (\sigma_1, \sigma_2, \ldots, \sigma_{t-1})$ the Markovian strategy profile followed by the short-run players that are active before stage t. Given σ^{t-1}, the information available to player it is $\mathscr{H}^{\sigma^{t-1}}_{it}$, the smallest sub-sigma algebra of $\mathscr{B}$ that refines $\mathscr{H}^{\sigma^{t-2}}_{i,t-1}$ and makes s_{it}, N_{it}, and $\{\sigma_{j,t-1}\}_{j \in N_{it}}$ measurable. A Markovian strategy for player it is a mapping $\sigma_{it}\colon \Omega \to A_i$ that is measurable with respect to $\mathscr{H}^{\sigma^{t-1}}_{it}$. We let $\sigma = (\sigma_1, \sigma_2, \ldots)$ denote a Markovian strategy profile generated as above and let $\mathscr{H}^\sigma_{i\infty} = \bigvee_{t=1}^\infty \mathscr{H}^{\sigma^{t-1}}_{it}$. Note that, for any strategy profile σ and all i, $\mathscr{H}^{\sigma^{t-1}}_{it} \subseteq \mathscr{H}^{\sigma^{t'-1}}_{it'}$ if $t \leqslant t'$. Whenever there is no risk of confusion we use $\mathscr{H}^\sigma_{it}$ to mean $\mathscr{H}^{\sigma^{t-1}}_{it}$.

We next construct the endogenous probability distribution induced by a Markovian strategy profile σ over $(\Omega, \mathscr{B})$. The payoff-relevant state θ is distributed according to some exogenous probability distribution P_0. The history of the game at the end of stage t is defined recursively: $h_0 = \theta$ and $h_t = (h_{t-1}; s_t, a_t, N_t)$ for all $t \geqslant 1$, where s_t, a_t, and $N_t = (N_{1t}, \ldots, N_{nt})$ are the signal, action, and neighborhood profiles realized in stage t. Given history h_{t-1}, the private signals and neighborhoods in stage t are distributed according to $\pi_t(h_{t-1}) \in \Delta(S \times \mathscr{N})$. The mapping $h_{t-1} \mapsto \pi_t(h_{t-1})$ is a transition probability from the set of all histories to $S \times \mathscr{N}$.[9] We assume that the probability distribution P_0 and transition probabilities $\{\pi_t\}_{t \in \mathbb{N}}$ are common knowledge. P_0 and π_1 uniquely define a probability distribution P_1 over $\Theta \times S \times \mathscr{N}$. Likewise, for $t \geqslant 2$, the strategy profile $\sigma^{t-1} = (\sigma_1, \sigma_2, \ldots, \sigma_{t-1})$, probability distribution P_0, and transition probabilities $\pi_1, \pi_2, \ldots, \pi_t$ uniquely define a probability distribution $P^{\sigma^{t-1}}_t$ over $\Theta \times S^t \times \mathscr{N}^t$. The probability measures $P_0, P_1, P^{\sigma^1}_2, \ldots$ can be uniquely extended to $\Omega = \Theta \times S^{\mathbb{N}} \times \mathscr{N}^{\mathbb{N}}$ by Tulcea's extension theorem.[10] We denote this extension by P^σ and the corresponding expectation operator by E^σ.

Definition 2. A Markovian strategy profile σ for the game with endogenous signals is an *MBE* if for all i, t, and $\mathscr{H}^{\sigma_{t-1}}_{it}$-measurable mappings $\sigma'_{it}\colon \Omega \to A_i$,

$$E^\sigma[u_i(\sigma_{it}, \sigma_{-it}, \theta) \mid \mathscr{H}^{\sigma_{t-1}}_{it}] \geqslant E^\sigma[u_i(\sigma'_{it}, \sigma_{-it}, \theta) \mid \mathscr{H}^{\sigma_{t-1}}_{it}].$$

Our next result generalizes Proposition 2 on the uniqueness of equilibrium in the quadratic game to the setting with endogenous signals and time-varying networks.

Proposition 3. *The quadratic game with endogenous signals and time-varying networks has a unique MBE σ.*

We next present a series of results for the quadratic payoff in (5) with endogenous signals and time-varying networks that are counterparts of the results stated for symmetric strictly supermodular games in §3. We first state an intermediate result that shows that the actions of short-run players in a given role converge to some limit action.

Proposition 4. *Let σ be the MBE of the quadratic game with endogenous signals. For all i there exists some $\mathcal{H}^{\sigma}_{i\infty}$-measurable random variable $\sigma_{i\infty}: \Omega \to A_i$ such that $\sigma_{it} \to \sigma_{i\infty}$, P^{σ}-almost surely.*

The convergence follows by using the fact that equilibrium action of each player can be represented as a weighted sum of expectations with respect to beliefs of all orders of the payoff-relevant state, and that each term in this sum converges as $t \to \infty$.

Our next result shows consensus in actions of players when the network is sufficiently connected over time. A network process is a probability distribution over $\mathcal{N}^{\mathbb{N}}$.

Definition 3. A network process is called *infinitely often almost surely strongly connected* if for any two nodes i and j there exists a sequence of nodes $k_1, k_2, \ldots, k_m$ such that $k_1 = i$, $k_m = j$, and for all $1 \leqslant l < m$, $k_l \in N_{k_{l+1}, t}$ almost surely infinitely often.

Theorem 3. *Let σ be the MBE of the quadratic game with endogenous signals. If the network process* $\mathrm{marg}_{\mathcal{N}^{\mathbb{N}}} P^{\sigma}$ *is infinitely often almost surely connected, then for all $i, j \in N$, $\sigma_{it} - \sigma_{jt} \to 0$, P^{σ}-almost surely, as t goes to infinity.*

The following result is the counterpart of the result on consensus in payoffs presented in Theorem 1, presented here for the game with quadratic payoffs.

Corollary 2. *Let σ be the MBE of the quadratic game with endogenous signals. If the network process* $\mathrm{marg}_{\mathcal{N}^{\mathbb{N}}} P^{\sigma}$ *is infinitely often almost surely connected, then for all $i, j \in N$, $u_i(\sigma_t, \theta) - u_j(\sigma_t, \theta) \to 0$, P^{σ}-almost surely, as t goes to infinity.*

The above two results extend the results in §3 to endogenous signals and time-varying networks when the payoff function is quadratic. We discuss their implications further in §5.3.

5. Discussion

In the games considered in §3, players acquire exogenous private signals s_i at the beginning of the game that reveal information about the state of the world θ. They use this information to play the MBE action. The action played by each player becomes known to the players in neighboring roles. From the perspective of player i, the actions of neighbors $j \in N_i$ reveal information about their private signals, which can be used to improve the actions that they play in the subsequent stage. As time progresses, actions of neighbors reveal more information about their private signals as well as information about the private signals of their neighbors, and the signals of their neighbors' neighbors. If the network is connected, all players eventually observe actions that carry information about the private signals of all other players. The results in §§3 and 4 characterize the asymptotic behavior of the agents involved in this game. This section discusses the insights that these results provide.

5.1. Consensus

When players play this game with incomplete information over a network, how much do they learn of each other's private information? Perhaps not all, but Theorem 1 asserts that they achieve a steady state in which they have no reason to suspect they have not. Indeed, the claim in Theorem 1 is that given any pair of players i and j, their strategies σ_{it} and σ_{jt} approach each other as the number of plays grow, with probability one. Since the players use a common strategy in the limit, we say that they achieve consensus. In this consensus state players select identical actions, which they therefore must believe to be optimal given *all* their available information and the strategies of other players. Otherwise, deviations to strategies with better expected payoffs would be possible. To emphasize that players achieve this possibly misguided consensus, we show in Corollary 1 that the payoffs of all players eventually coincide.

That players achieve consensus is not unexpected given that the game being played is a symmetric supermodular game. If the state of the nature θ were known to the players, they would all play a common action in any equilibrium of the game. When the state of the world is not known but rather inferred from private signals and the observed actions of neighboring players, the incentive to coordinate is still present, but there is uncertainty on what exactly a coordinated action should be. Theorem 1 shows that such uncertainty is eventually resolved.

Expected as it may be, the result in Theorem 1 is not obvious because it is not clear that the uncertainty on what it means to have a coordinated action is resolved. The fundamental problem in resolving this uncertainty is that players have to guess the actions other players are about to take, yet they only know their strategies and observe only their realized actions. If other players' histories were observed, the incentive to coordinate, that is implicit in the supermodular assumption, would drive players to consensus. However, histories are not observed. The strategies of players other than i are, indeed, not necessarily measurable with respect to the information available to i. Lacking measurability, it is not possible for i to gauge the quality of his actions given the strategies of his neighbors.

The key step in the proof of Theorem 1 is to show that the strategies of neighbors become measurable in the limit. When strategies become measurable, it is possible for i to imitate j, if it so happens that the strategy of j is better. Since the player i acts with respect to MBE strategy, imitating j's strategy cannot be optimal. It follows that the strategy of j is not better than the strategy of i according to i. Yet, strategic complementarity implies that i cannot think that his strategy in the limit is better than j's limit strategy and vice versa, and at the same time their strategies be different.

According to Corollary 1, the differences between the players' payoffs asymptotically vanish. Thus, in spite of the differences in their location in the network and the quality of their private signals, players asymptotically receive similar payoffs. From the point of view of the players, however, the asymptotic payoffs are not necessarily the same. That is, conditional expectations of the players' limit payoffs given their information at the end of the game could be dissimilar. The following example illustrates this possibility.

Example 1. Consider two roles $i \in \{1, 2\}$ with payoffs given by (5) that observe each others' actions in all stages. The common prior is the uniform distribution over the set $\{-2, -1, 1, 2\}$. Player 2 receives no signal ($S_2 = \varnothing$), whereas Player 1's private signals belong to the set $S_1 = \{1, 2\}$, with $s_1 = |\theta|$. Thus, Player 1 is informed of the absolute value of θ. Observe that in any equilibrium of the game $\sigma_{it} = 0$ at all times and for both players, Player 1 learns the absolute value of θ, whereas Player 2 never makes any informative observations. At the end of the game, Player 1's expected payoff conditional on his information is equal to $-(1-\lambda)|\theta|^2$ whereas the corresponding payoff for Player 2 is given by $-(1-\lambda)\frac{5}{2}$.

In the above example, although the conditional expected payoffs are unequal for any realization of the state, the unconditional expected payoffs and the realized payoffs are the same for both players because Theorem 1 and Corollary 1 apply.

We remark that strategic complementarity is the main driver of the consensus results. In particular, in games with strategic substitutes, it is beneficial for the players to play different strategies. The games wherein players' actions are strategic substitutes might not even have any symmetric pure-strategy Nash equilibrium (e.g., the hawk-dove game). Hence, the consensus results cannot be generalized to games with strategic substitutability.

5.2. Information Aggregation

As we noted in §5.1, achieving consensus means that players have no reason to suspect there is more information to be learnt, but this does not necessarily mean that they have aggregated all the available information. To understand the difference, it is instructive to consider the following example.

Example 2. Consider two roles $i \in \{1, 2\}$ with players having the utility of the form in (5). The players in the two roles observe each others' actions in all stages. The state of the world θ belongs to the set $\Theta = \{-1, 1\}$. Initially, players have uniform common prior. They receive private signals belonging to the set $S_1 = S_2 = \{\mathrm{H}, \mathrm{T}\}$. As in the baseline model of §2, only the initial short-run players receive a signal, and the distribution of $s = (s_1, s_2)$ conditional on θ is given by

$$s(\theta) \sim \begin{cases} \frac{1}{2}\delta_{(\mathrm{H},\mathrm{H})} + \frac{1}{2}\delta_{(\mathrm{T},\mathrm{T})} & \text{if } \theta = 1, \\ \frac{1}{2}\delta_{(\mathrm{H},\mathrm{T})} + \frac{1}{2}\delta_{(\mathrm{T},\mathrm{H})} & \text{if } \theta = -1, \end{cases}$$

where δ_s is the degenerate probability distribution with unit mass on the signal profile $s \in S$. We first show that, in the unique equilibrium of the game, players in both roles choose $\sigma_{it} = 0$ for all t. Given the distribution of s_1 and s_2, each player in the first stage receives the signal H (T) with probability one half, regardless of the realization of θ. Players' private signals are thus completely uninformative about the realized state, and hence their expectation of θ is equal to zero: $\mathbb{E}[\theta \mid \mathscr{H}_{i1}] = 0$. Since the distribution of s is common knowledge, each player knows that the other player's expectation of θ is also zero. As a result, the equilibrium action of initial players is $\sigma_{i1} = 0$ for $i \in N$. These actions reveal no information to the players in the subsequent stages. Therefore, the short-run players in subsequent stages all continue to choose the zero action.

Next, consider the alternative setting in which both players observe the complete signal profile $s = (s_1, s_2)$. In this modified game, both of the short-run players playing in the first stage learn the realized state. Therefore, in equilibrium of the modified game players in both roles choose $\sigma_{it} = \theta$ for all $t \geqslant 1$ and given any realization of θ.

In the above example players achieve consensus on their strategy, as it should be because Theorem 1 applies, but the consensus strategy is not the one they would use if they had access to each others' private signals. The question then arises as to which conditions can prevent this failure to aggregate information through repeated network interactions. To answer this question, we restrict attention to utilities that have the quadratic form shown in (5).

The game in Example 2 has a quadratic utility, yet this restriction is not sufficient to guarantee information aggregation. Observe, however, that the parameters in Example 2 have been fine-tuned to make sure that private signals are uninformative about the state of the world when taken separately but informative when taken together. This causes players to take equilibrium actions that reveal no information about their private signals, although the signals contain useful information about the realized state. However, the delicate balance of parameters that creates this behavior is broken if we have infinitesimal variations in any parameter. In that sense, Example 2 is not generic in that it represents an isolated example that cannot be drawn if we consider games that are selected from a residual set.

Theorem 2 states that in network games with quadratic utilities, we can have examples where players fail to aggregate information, but these examples are nongeneric in the space of probability measures over $(\Omega, \mathscr{B})$ endowed with the total variation metric. Conversely, network games

Figure 1. The ring and star social networks of Example 3.

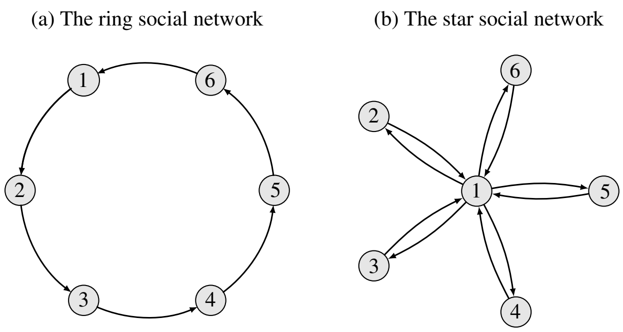

with quadratic utilities aggregate information in a set that is dense in $(\Omega, \mathcal{B})$. This means that failure to aggregate information is not a practical concern when utilities are quadratic. In practically all possible games of incomplete information, players not only achieve consensus on their strategies but also converge to the strategy defined by the expectation $E^P[\theta \mid \mathcal{H}_\infty^{\sigma^P}]$. This means that players end up playing the same action that would be selected by an omniscient planner that has access to all the private signals of all players. The question of whether this is also true for general supermodular utilities remains open.

An important special case of Theorem 2 is obtained by letting $\lambda = 0$. In this case, the players 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 players' problem then becomes an instance of social learning. Theorem 2 states that the players asymptotically learn to estimate the state as if they had access to all the available information. In this sense, Theorem 2 parallels 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 players face payoff externalities in addition to information externalities. It also extends Proposition 4 of Jadbabaie et al. (2012) to the case where the players communicate their conditional estimates of the state, rather than their entire beliefs. Finally, it also relates to the information aggregation result presented in Ostrovsky (2012) for dynamic markets with "separable" securities. The notion of separability rules out fine-tuned payoff and prior beliefs such as the one in Example 2. Our results complement the work in Ostrovsky (2012) in that we show that—even when the players' partition of the state space does not satisfy the separability condition—information aggregation is generically obtained.

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

Example 3. There are $n = 6$ players over a fixed strongly connected social network playing the game with payoffs given by (5) 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 $S_1 = S_2 = \mathbb{R}$. We assume that, conditional on θ, s_1 and s_2 are independent and distributed according to $\mathcal{N}(\theta, 1)$. Note that only the initial players receive informative signals. The evolution of the players' actions over time is depicted in Figure 2 for two realizations of the path of play with $\theta = 0$. The dashed line represents the payoff-maximizing action given the players' private signals—which in the context of this example is equal to the average of the private signals. At $t = 1$, players each receive a private signal and choose the actions that are equal to their private signals. Yet, as time passes, the players' actions converge to the payoff-maximizing action. Moreover, over both of the networks, convergence is complete after a number of time periods equal to the diameter of the graph.[11]

In this example, although the players' signal spaces are not finite, convergence to the optimal action is achieved.[12]

Figure 2. (Color online) Evolution of the agents' actions over time in Example 3.

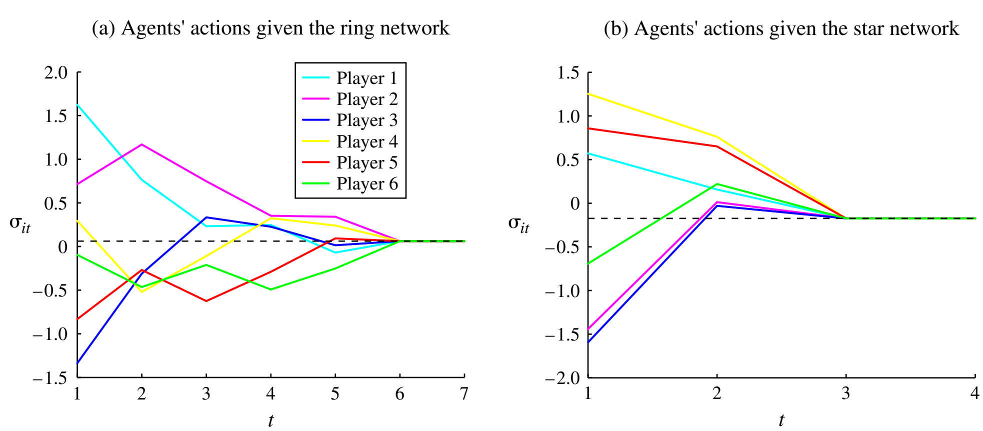

5.3. Extensions

Some extensions have been also provided in §4 for the special case of quadratic payoffs. Theorem 3 states that consensus happens for quadratic utilities even if we allow for time-varying networks and endogenous signals. The allowance for time-varying networks extends our results to cases where interactions between players are not synchronized to a common clock as is assumed in §3. Whenever a pair of players interact, we can define a network with a single edge that joins the pair of players. This definition pigeonholes asynchronous interactions into the time-varying network model of §4. The allowance for endogenous signals includes, as a particular case, situations in which players observe their own payoffs.

In contrast to the consensus results, the information aggregation result in Theorem 2 cannot be extended to the case with endogenous signals. The following example illustrates that the correct notion of information aggregation is unclear when signals are endogenous.

Example 4. Consider a single player who repeatedly plays a game with payoffs as in (5) with $\lambda = 0$. The player's prior for θ is given by the standard normal $\mathcal{N}(0,1)$. His first stage signal is distributed according to $\mathcal{N}(\theta,1)$. The signal s_t observed at $t \geqslant 2$ is distributed as follows:

$$s_t(\theta) \sim \begin{cases} \mathcal{N}(\theta,1) & \text{if } |\theta - a_1| > 1, \\ \mathcal{N}(0,1) & \text{if } |\theta - a_1| \leqslant 1. \end{cases}$$

The player observes an informative signal and chooses an action in the first stage. If his stage one action is not within unit distance of the realized state, he continues to observe informative private signals and asymptotically learns the state with arbitrary precision. However, if the player's first stage action is sufficiently close to the realized state, he does not observe any informative signals after the first stage and thus never learns the state.

In this example, there is an externality associated with the effect of the action of the player in stage one on the distribution of the private signals observed by his successors. If the player 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 players is not fixed; rather it is a function of the realized path of play. Consequently, no well-defined notion of optimal aggregation of information is readily available when signals are endogenous.

Throughout, we assumed that the network is strongly connected over time. When the network is not strongly connected, our results do not continue to hold. Consider a single role i that is disconnected from the rest of the network, and assume that the initial player in each role only observes a single noisy signal of the state. Unless all players happen to be perfectly informed of the state, the players in the disconnected role i will not be in agreement with the rest of the population. Note that this is also true when the other players can observe the actions of the players in role i but the players in role i cannot observe any other player in the network.

Some other extensions are beyond the scope of this paper. For instance, our results are stated for myopic players, but players that optimize for longer time horizons have even stronger incentives to signal their information. Therefore, it is reasonable to expect that all of our theorems hold in this case as well. In fact, it is likely that stronger results can be derived because nonmyopic Markovian players may be able to aggregate information even if the short-run players cannot.

6. Symmetric Supermodular Games

We present four examples of symmetric strictly supermodular games of incomplete information to illustrate the range of models to which our consensus results in §3 are applicable.

6.1. Currency Attacks

Consider investors who attack a currency by short-selling the currency by $a_i \in [0,1]$ amounts. There is a fixed transaction cost of short selling, $-c < 0$, when investor i attacks $a_i > 0$; otherwise, his cost is zero. The strength of the attack is proportional to the average short-selling actions of the investors: $\bar{a} = \sum_i a_i/n$. The government follows a thresholded policy to defend against the investors' attacks based on the fundamentals of the economy θ. That is, if the attack strength is larger than $h(\theta) \in (0,1]$ where $h(\cdot)$ is an increasing function of θ, then the government does not defend; otherwise, it defends. When the government defends, the attack fails and the investors incur the transaction cost. When the government does not defend, the attack succeeds and each investor receives a benefit proportional to his short-selling amount, $B_i(a_i) > 0$, which is a continuous strictly increasing function. However, the investors do not exactly know fundamentals of the economy and only have private information regarding θ. We smooth the government's threshold response by assuming that the likelihood that $\bar{a}$ is larger than $h(\theta)$ is given by $\mathcal{L}(h(\theta);\bar{a})$, which is a continuous and strictly increasing function of $\bar{a}$ given θ. Then the payoff of an investor is summarized as follows.

$$u_i(a_i, a_{-i}, \theta) = \begin{cases} B_i(a_i)\mathcal{L}(h(\theta);\bar{a}) - c & \text{if } a_i > 0, \\ 0 & \text{if } a_i = 0. \end{cases}$$

Under certain assumptions, the utility function above is strictly supermodular. For instance, it is easy to show that the likelihood function $\mathcal{L}(h(\theta);\bar{a}) = \bar{a}^2/(\lambda + h(\theta)^2)$ results in a strictly supermodular utility function for all $\lambda \geqslant 1$. Furthermore, the utility function is symmetric, since each investor's attack contributes equally to the strength of the attack. See Vives (2005) for a variant of this game.

6.2. Bertrand Competition

Consider an oligopoly price competition model where the demand for firm i is determined by the price set by firm i, $a_i \in [0, 1]$, as well as prices of its competitors a_{-i}. That is, firm i's demand function is $D_i(a_i, a_{-i})$. The demand of firm i is decreasing in its own price a_i and increasing with respect to prices of others a_{-i}. The revenue of firm i is its price multiplied by the demand, $a_i D_i(a_i, a_{-i})$. Each firm operates with an identical uncertain cost per production θ. Then the cost of matching demand $D_i(a_i, a_{-i})$ by firm i is θa_i. The payoff of firm i is its net revenue, which is the difference between revenue and cost,

$$u_i(a_i, a_{-i}, \theta) = a_i D_i(a_i, a_{-i}) - \theta a_i.$$

We consider a logistic demand function $D_i(a) = 1/(1 + \sum_{j\neq i} \kappa \exp(\lambda(a_i - a_j)))$ for $\kappa > 0$ and $\lambda > 0$. This demand function yields a symmetric strictly supermodular utility function. See Milgrom and Roberts (1990) for other forms of demand functions that result in supermodular utilities.

6.3. Power Control in Wireless Networks

Consider the problem of power control in wireless network communication; see Altman and Altman (2003). Each user wants to transmit to a base station using the channel designated to himself. User j determines a transmitting power level $a_i \in [0, \hat{a}]$ for some $\hat{a} > 0$. The channel gain of user i transmitting to base station is equal to $h > 0$, which is identical for all the users. Hence, the received signal of user i at the base station is $a_i h$. On the other hand, the transmission of other users interferes with the gain of user i's channel. Given the channel gains h, the signal-to-interference-ratio (SINR) is given by

$$\text{SINR}(a_{-i}) = \frac{h}{h\sum_{j\neq i} a_j + \rho},$$

where $\rho > 0$ is the additive Gaussian noise representing the noise at the base station. Thus the received SINR by user i when it exerts a_i amounts of power is simply $a_i \text{SINR}_i(a_{-i})$. User i incurs a constant uncertain cost θ per unit of power exerted yielding a total cost of θa_i when a_i units of power is exerted. The payoff of user i is the difference between a function of the received SINR $B_i(a_i \text{SINR}_i(a_{-i}))$ and the cost of power consumption,

$$u_i(a_i, a_{-i}, \theta) = B_i(a_i \text{SINR}_i(a_{-i})) - \theta a_i.$$

Under certain conditions on the function $B_i(\cdot)$, the payoff is strictly supermodular. For instance, given $B_i(x) = x^{1-\alpha}/(1-\alpha)$ where $\alpha > 1$, we have $\partial^2 u_i/\partial a_i \partial a_j > 0$. Symmetry of the utility function follows by the definition of the SINR and the unanimity of the channel gain h.

6.4. Arms Race

N countries engage in an arms race; see Milgrom and Roberts (1990). Country i chooses its arms level $a_i \in [0, \hat{a}]$ and incurs a cost of armament that is captured by the cost function $C_i(a_i, \theta)$ that depends on the state of the world θ and own action a_i. The benefit of the armament depends on the distance between self arms, a_i, and the average armament of other countries, $\bar{a}_{-i} = \sum_{j\neq i} a_j/(n-1)$, captured by a strictly concave smooth function $B_i(a_i - \bar{a}_{-i})$. The payoff of country i is given by

$$u_i(a_i, a_{-i}, \theta) = -C_i(a_i, \theta) + B_i(a_i - \bar{a}_{-i}).$$

Since $\partial^2 u_i/\partial a_i \partial a_j = -B_i''(a_i - a_j) > 0$, the game is strictly supermodular. Furthermore, by construction, the utility function is symmetric.

7. Proofs

7.1. Proof of Proposition 2

This proposition is a special case of Proposition 3 proved below.

7.2. Proof of Proposition 3

The proof is constructive. We start at $t = 1$ and inductively construct the unique equilibrium. Let $[\check{a}, \hat{a}]$ be such that $A_i \subseteq [\check{a}, \hat{a}]$.

Consider stage $t = 1$. We show that the game played in the first stage is dominance solvable in the sense that each player has a unique rationalizable strategy.[13] We use the following procedure to iteratively eliminate strictly dominated strategies. The strategy that survives is the unique rationalizable strategy profile and thus also the unique equilibrium. First consider all the possible beliefs that each player can entertain about other players' actions, the payoff-relevant state, and the private signals such that the belief is consistent with the player's prior belief and his observed private signal. Any strategy for player $i1$ that is not a best response to some such belief is eliminated in the first round. Next consider the set of all beliefs that in addition to consistency are restricted to only put positive probability on the strategies that were not eliminated in the previous round and eliminate the strategies that are not best responses to any such belief. We repeat the procedure ad infinitum. The strategy profile surviving this iterated elimination procedure is the unique equilibrium.

Let P_{i1} be the belief player $i1$ entertains about the state, signals, networks, and what other players do in stage 1. P_{i1} is a transition probability from Ω to $\Omega \times A_{-i}$. Player it's best response to this belief is then given by the following first-order condition:

$$\sigma_{i1} = (1-\lambda)E_{i1}[\theta] + \frac{\lambda}{n-1}\sum_{j\neq i} E_{i1}[a_{j1}],$$

where E_{i1} is the expectation operator with respect to P_{i1}. Each player's actions are restricted to belong to the interval $[\check{a}, \hat{a}]$. Thus,

$$(1-\lambda)E_{i1}[\theta]+\lambda\check{a}\leqslant\sigma_{i1}\leqslant(1-\lambda)E_{i1}[\theta]+\lambda\hat{a}.$$

P_{i1} needs to be consistent with the P_1, defined in §4.2, and player $i1$'s observation of signal s_{i1}. Therefore, $E_{i1}[\theta]=E_1[\theta\mid\mathcal{H}_{i1}]$ and so

$$(1-\lambda)E_1[\theta\mid\mathcal{H}_{i1}]+\lambda\check{a}\leqslant\sigma_{i1}\leqslant(1-\lambda)E_1[\theta\mid\mathcal{H}_{i1}]+\lambda\hat{a},$$

where E_1 is the expectation operator with respect to P_1. Thus, the strategies not belonging to the interval $[(1-\lambda)E_1[\theta\mid\mathcal{H}_{i1}]+\lambda\check{a}, (1-\lambda)E_1[\theta\mid\mathcal{H}_{i1}]+\lambda\hat{a}]$ are eliminated in the first round of elimination.

Given that each $j1$'s actions belong to the interval $[(1-\lambda)E_1[\theta\mid\mathcal{H}_{j1}]+\lambda\check{a}, (1-\lambda)E_1[\theta\mid\mathcal{H}_{j1}]+\lambda\hat{a}]$, the support of player $i1$'s belief P_{i1} need to be contained in $[(1-\lambda)E_1[\theta\mid\mathcal{H}_{j1}]+\lambda\check{a}, (1-\lambda)E_1[\theta\mid\mathcal{H}_{j1}]+\lambda\hat{a}]$. Thus, any rationalizable strategy for player $i1$ needs to satisfy the following restrictions:

$$\sigma_{i1}\geqslant(1-\lambda)E_1[\theta\mid\mathcal{H}_{i1}]+\frac{\lambda(1-\lambda)}{n-1}\sum_{j\neq i}E_1[E_1[\theta\mid\mathcal{H}_{j1}]\mid\mathcal{H}_{i1}]+\lambda^2\check{a},$$

$$\sigma_{i1}\leqslant(1-\lambda)E_1[\theta\mid\mathcal{H}_{i1}]+\frac{\lambda(1-\lambda)}{n-1}\sum_{j\neq i}E_1[E_1[\theta\mid\mathcal{H}_{j1}]\mid\mathcal{H}_{i1}]+\lambda^2\hat{a}.$$

This procedure can be repeated ad infinitum. Define $\bar{\theta}^k_{i1}$ recursively as follows: $\bar{\theta}^1_{i1}=E_1[\theta\mid\mathcal{H}_{i1}]$ and $\bar{\theta}^{k+1}_{i1}=\sum_{j\neq i}E_1[\bar{\theta}^k_{j1}\mid\mathcal{H}_{i1}]/(n-1)$. Then for any $k\geqslant 1$,

$$(1-\lambda)\sum_{l=1}^{k}\lambda^{l-1}\bar{\theta}^l_{i1}+\lambda^k\check{a}\leqslant\sigma_{i1}\leqslant(1-\lambda)\sum_{l=1}^{k}\lambda^{l-1}\bar{\theta}^l_{i1}+\lambda^k\hat{a}.$$

The difference between the upper and lower bounds on σ_{i1} in the kth stage of elimination is given by $\lambda^k[\check{a}-\hat{a}]$. Since $\lambda<1$ and $\bar{\theta}^k_i$ belongs to Θ which is a compact set, as k goes to infinity, the upper and lower bounds converge to the same value and a strategy σ_{i1} survives that is unique up to sets of P_1-measure zero:

$$\sigma_{i1}=(1-\lambda)\sum_{k=1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{i1}.$$

The game played in stage 2 is similar to the game played in the first stage with the exception that the information of player $i2$ is now given by $\mathcal{H}^{\sigma^1}_{i2}$. An argument similar to the one above shows that there exists an equilibrium strategy for the stage 2 short-run players that is unique up to sets of $P^{\sigma^1}_2$-probability zero. More generally, by induction, there exists a unique equilibrium that in stage t is given by the following expression:

$$\sigma_{it}=(1-\lambda)\sum_{k=1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{it}, \tag{6}$$

where

$$\bar{\theta}^1_{it}=E^{\sigma^{t-1}}_{t-1}[\theta\mid\mathcal{H}^{\sigma^{t-1}}_{it}]\quad\text{and}$$

$$\bar{\theta}^{k+1}_{it}=\sum_{j\neq i}E^{\sigma^{t-1}}_{t-1}[\bar{\theta}^k_{jt}\mid\mathcal{H}^{\sigma^{t-1}}_{it}]/(n-1).$$

7.3. Proof of Proposition 4

Since θ belongs to a compact set Θ and $\mathcal{H}^\sigma_{it}\uparrow\mathcal{H}^\sigma_i$, by the martingale convergence theorem, $\bar{\theta}^1_{it}\to\bar{\theta}^1_{i\infty}$, P^σ-almost surely, where $\bar{\theta}^1_{i\infty}=E^\sigma[\theta\mid\mathcal{H}^\sigma_{i\infty}]$. Suppose that $\bar{\theta}^k_{it}\to\bar{\theta}^k_{i\infty}$, P^σ-almost surely, for some k, where $\bar{\theta}^{k+1}_{i\infty}=\sum_{j\neq i}E^\sigma[\bar{\theta}^k_{j\infty}\mid\mathcal{H}^\sigma_{i\infty}]/(n-1)$ for all $k\geqslant 1$. Then by the dominated convergence theorem for conditional expectation, $\bar{\theta}^{k+1}_{it}\to\bar{\theta}^{k+1}_{i\infty}$, P^σ-almost surely. Therefore, $\bar{\theta}^k_{it}\to\bar{\theta}^k_{i\infty}$, P^σ-almost surely, for all $k\geqslant 1$.

Define

$$\sigma_{i\infty}=(1-\lambda)\sum_{k=1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{i\infty}. \tag{7}$$

We show that $\sigma_{it}\to\sigma_{i\infty}$ almost surely. Fix some arbitrary $\epsilon>0$. Since $\lambda<1$ and $|\bar{\theta}^k_{it}|\leqslant\max_\Theta|\theta|<\infty$, there exists some $K_1\geqslant 1$ such that

$$\left|(1-\lambda)\sum_{k=K_1+1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{it}\right|\leqslant\frac{\epsilon}{3},$$

for all t. Likewise, there exists some $K_2\geqslant 1$ such that

$$\left|(1-\lambda)\sum_{k=K_2+1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{i\infty}\right|\leqslant\frac{\epsilon}{3}.$$

Let $K=\max\{K_1,K_2\}$. Since $\bar{\theta}^k_{it}\to\bar{\theta}^k_{i\infty}$ almost surely for all $k\geqslant 1$, for P^σ-almost all $\omega\in\Omega$ there exists some T such that $(1-\lambda)\lambda^{k-1}|\bar{\theta}^k_{it}-\bar{\theta}^k_{i\infty}|<\epsilon/3K$ for all $k\leqslant K$. If $t\geqslant T$, then

$$|\sigma_{it}-\sigma_{i\infty}|\leqslant(1-\lambda)\sum_{k=1}^{K}\lambda^k|\bar{\theta}^k_{it}-\bar{\theta}^k_{i\infty}|+\left|(1-\lambda)\sum_{k=K+1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{it}\right|+\left|(1-\lambda)\sum_{k=K+1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{i\infty}\right|\leqslant\epsilon.$$

Thus, since $\epsilon>0$ was arbitrary and convergence is for a set of full measure, $\sigma_{it}\to\sigma_{i\infty}$, P^σ-almost surely.

7.4. Proof of Theorem 3

Let i,j be a pair of roles such that i observes the actions of j infinitely often almost surely. Define the mappings $\sigma'_{it}\colon\Omega\to A_i$ as follows: $\sigma'_{i1}=E^\sigma[\theta\mid\mathcal{H}_{i1}]$, and for $t\geqslant 2$, $\sigma'_{it}=\sigma_{j,t-1}$ if $j\in N_{it}$, and $\sigma'_{it}=\sigma'_{i,t-1}$ otherwise. Since σ is an equilibrium,

$$E^\sigma[u_i(\sigma_{it},\sigma_{-it},\theta)\mid\mathcal{H}^{\sigma_{t-1}}_{it}]\geqslant E^\sigma[u_i(\sigma'_{it},\sigma_{-it},\theta)\mid\mathcal{H}^{\sigma_{t-1}}_{it}]\quad P^\sigma\text{-a.s.}$$

Take expectations of the above inequality with respect to E^σ. For any ω for which $\sigma_{jt}\to\sigma_{j\infty}$ and $j\in N_i$ infinitely often, $\sigma'_{it}\to\sigma_{j\infty}$. Therefore, by the dominated convergence theorem,

$$E^\sigma[u_i(\sigma_{i\infty},\sigma_{-i\infty},\theta)]\geqslant E^\sigma[u_i(\sigma_{j\infty},\sigma_{-i\infty},\theta)].$$

A similar argument can be used to show that

$$E^{\sigma}[u_j(\sigma_{j\infty},\sigma_{-j\infty},\theta)]\geqslant E^{\sigma}[u_j(\sigma_{i\infty},\sigma_{-j\infty},\theta)].$$

By the connectivity assumption, there exists a sequence of roles $i_0, i_1, i_2, \ldots, i_n$ starting and ending with the same role that includes each role other than i_0 exactly once, and such that, for all k, players in role i_k observe the ones in role i_{k+1} infinitely often almost surely. For any k, by the above argument

$$E^{\sigma}[u_i(\sigma_{i_k\infty},\sigma_{-i_k\infty},\theta)]\geqslant E^{\sigma}[u_i(\sigma_{i_{k+1}\infty},\sigma_{-i_k\infty},\theta)]. \quad (8)$$

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

$$\sum_{k=0}^{n-1}E^{\sigma}[u_i(\sigma_{i_k\infty},\sigma_{-i_k\infty},\theta)]\geqslant\sum_{k=1}^{n}E^{\sigma}[u_i(\sigma_{i_k\infty},\sigma_{-i_{k-1}\infty},\theta)].$$

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

$$\sum_{k=0}^{n-1}E^{\sigma}\bigg[\sigma_{i_k\infty}\sum_{j\neq k}\sigma_{i_j\infty}\bigg]\geqslant\sum_{k=1}^{n}E^{\sigma}\bigg[\sigma_{i_k\infty}\sum_{j\neq k-1}\sigma_{i_j\infty}\bigg].$$

Further simplification results in

$$\sum_{k=1}^{n}E^{\sigma}[\sigma_{i_k\infty}\sigma_{i_{k-1}\infty}]\geqslant\sum_{k=1}^{n}E^{\sigma}[\sigma^2_{i_k\infty}]. \quad (9)$$

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

$$\sum_{k=1}^{n}E^{\sigma}[\sigma^2_{i_k\infty}]\geqslant\sum_{k=1}^{n}E^{\sigma}[\sigma_{i_k\infty}\sigma_{i_{k-1}\infty}], \quad (10)$$

with equality if and only if $\sigma_{i_k\infty}=\sigma_{i_{k-1}\infty}$ for all k, P^{σ}-almost surely; Equation (9) implies that (10) indeed holds with equality. Thus, for all i and j and with P^{σ}-probability one, $\sigma_{i\infty}=\sigma_{j\infty}$. Together with Proposition 4, this completes the proof of the theorem.

7.5. Proof of Corollary 2

The corollary immediately follows Proposition 4, Theorem 3, and the assumption that the utility functions are continuous.

7.6. Proof of Theorem 2

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

Lemma 1. *Let $(X,\mathscr{B})$ be a measurable space, and let $(\mathbf{P},d)$ be the metric space where $\mathbf{P}$ is the collection of all probability measures on $(X,\mathscr{B})$ and d is the total variation distance. Let $\mathscr{F}_1$ and $\mathscr{F}_2$ be two arbitrary sub σ-algebras of $\mathscr{B}$, let $\mathscr{F}$ be the σ-algebra generated by the union of $\mathscr{F}_1$ and $\mathscr{F}_2$, and let f be an arbitrary bounded random variable. The set*

$$\mathbf{Q}=\{P\in\mathbf{P}\colon E_P[f\mid\mathscr{F}_1]=E_P[f\mid\mathscr{F}_2]\neq E_P[f\mid\mathscr{F}]\},$$

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

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

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

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

$$\begin{aligned}\int_{B^c} f\,dP&=\int_X f\,dP-\int_B f\,dP\\&=\int_X E_P[f\mid\mathscr{F}_1]\,dP-\int_B E_P[f\mid\mathscr{F}_1]\,dP\\&=\int_{B^c}E_P[f\mid\mathscr{F}_1]\,dP.\end{aligned}$$

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

$$\begin{aligned}\int_{\bigcup_{i=1}^{\infty}B} f\,dP&=\sum_{i=1}^{\infty}\int_{B_i} f\,dP=\sum_{i=1}^{\infty}\int_{B_i}E_P[f\mid\mathscr{F}_1]\,dP\\&=\int_{\bigcup_{i=1}^{\infty}B}E_P[f\mid\mathscr{F}_1]\,dP.\end{aligned}$$

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

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

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

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

$$\mathbf{R}=\{P\in\mathbf{P}\colon E_P[f\mid\mathscr{F}_1]=E_P[f\mid\mathscr{F}_2]\}.$$

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 $Q \in b_\epsilon$,

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

Therefore,

$$\int_{A_1\setminus C} f\, dQ + \int_C f\, dQ = \int_{A_1\setminus C} E_Q[f \mid \mathcal{F}_1]\, dQ + \int_C E_Q[f \mid \mathcal{F}_1]\, dQ. \tag{11}$$

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, $E_{Q_k}[f \mid \mathcal{F}_1]$ converges in Q-probability to $E_Q[f \mid \mathcal{F}_1]$.[14] Therefore, since f is bounded and Q_k converges in total variation distance to Q,

$$\int_{A_2} E_{Q_k}[f \mid \mathcal{F}_1]\, dQ_k \to \int_{A_2} E_Q[f \mid \mathcal{F}_1]\, dQ, \tag{12}$$

and

$$\int_{A_2} f\, dQ_k \to \int_{A_2} f\, dQ. \tag{13}$$

Moreover, for all k,

$$\int_{A_2} f\, dQ_k = \int_{A_2} E_{Q_k}[f \mid \mathcal{F}_2]\, dQ_k = \int_{A_2} E_{Q_k}[f \mid \mathcal{F}_1]\, dQ_k, \tag{14}$$

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 (12)–(14) imply that

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

And hence,

$$\int_{A_2\setminus C} f\, dQ + \int_C f\, dQ = \int_{A_2\setminus C} E_Q[f \mid \mathcal{F}_1]\, dQ + \int_C E_Q[f \mid \mathcal{F}_1]\, dQ. \tag{15}$$

We use (11) and (15) to conclude that $\int_C f\, dQ = \int_C E_Q[f \mid \mathcal{F}_1]\, dQ$ 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 $E_{\hat{Q}_\delta}[f \mid \mathcal{F}_1] = E_Q[f \mid \mathcal{F}_1]$. Let $B \in \mathcal{F}_1$ be arbitrary.

$$\begin{aligned}
\int_B f\, d\hat{Q}_\delta
&= \int_{B\cap A_1} f\, d\hat{Q}_\delta + \int_{B\cap A_1^c} f\, d\hat{Q}_\delta \\
&= (1+\delta Q(A_1^c)) \int_{B\cap A_1} f\, dQ + (1-\delta Q(A_1)) \int_{B\cap A_1^c} f\, dQ \\
&= (1+\delta Q(A_1^c)) \int_{B\cap A_1} E_Q[f \mid \mathcal{F}_1]\, dQ \\
&\quad + (1-\delta Q(A_1)) \int_{B\cap A_1^c} E_Q[f \mid \mathcal{F}_1]\, dQ \\
&= \int_{B\cap A_1} E_Q[f \mid \mathcal{F}_1]\, d\hat{Q}_\delta + \int_{B\cap A_1^c} E_Q[f \mid \mathcal{F}_1]\, d\hat{Q}_\delta \\
&= \int_B E_Q[f \mid \mathcal{F}_1]\, d\hat{Q}_\delta,
\end{aligned} \tag{16}$$

where the third equality follows from the assumption that $E_Q[f \mid \mathcal{F}_1]$ 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 $E_Q[f \mid \mathcal{F}_1]$ is $\mathcal{F}_1$-measurable, Equation (16) proves that $E_Q[f \mid \mathcal{F}_1]$ is a version of $E_{\hat{Q}_\delta}[f \mid \mathcal{F}_1]$. Let $B_1 = A_1 \setminus C$ and $B_2 = A_2 \setminus C$. Equations (11) and (15) imply that

$$\int_{B_1} [f - E_Q[f \mid \mathcal{F}_1]]\, dQ = \int_{B_2} [f - E_Q[f \mid \mathcal{F}_1]]\, dQ. \tag{17}$$

Since $B_1 \cap A_1 = B_1$,

$$\int_{B_1} [f - E_{\hat{Q}_\delta}[f \mid \mathcal{F}_1]]\, d\hat{Q}_\delta = (1+\delta Q(A_1^c)) \int_{B_1} [f - E_Q[f \mid \mathcal{F}_1]]\, dQ. \tag{18}$$

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

$$\int_{B_2} [f - E_{\hat{Q}_\delta}[f \mid \mathcal{F}_1]]\, d\hat{Q}_\delta = (1-\delta Q(A_1)) \int_{B_2} [f - E_Q[f \mid \mathcal{F}_1]]\, dQ. \tag{19}$$

On the other hand, if δ is sufficiently small, $\hat{Q}_\delta \in b_\epsilon$. Therefore, by (11) and (15),

$$\int_{B_1} [f - E_{\hat{Q}_\delta}[f \mid \mathcal{F}_1]]\, d\hat{Q}_\delta = \int_{B_2} [f - E_{\hat{Q}_\delta}[f \mid \mathcal{F}_1]]\, d\hat{Q}_\delta. \tag{20}$$

Equations (17)–(20) imply that

$$\int_{B_1} [f - E_Q[f \mid \mathcal{F}_1]]\, dQ = \int_{B_2} [f - E_Q[f \mid \mathcal{F}_1]]\, dQ = 0. \tag{21}$$

Thus, by (11),

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

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

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

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 $A \in \mathcal{F}$,

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

Together with the fact that $E_Q[f \mid \mathcal{F}_1]$ and $E_Q[f \mid \mathcal{F}_2]$ are both measurable with respect to $\mathcal{F}$, this shows that $E_Q[f \mid \mathcal{F}] = E_Q[f \mid \mathcal{F}_1] = E_Q[f \mid \mathcal{F}_2]$ 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 $\mathbf{P}$. □

7.6.1. Proof of Theorem 2. By the definition of $\bar{\theta}^k_{i\infty}$ and Equation (7),

$$\begin{aligned}\sigma_{i\infty} &= (1-\lambda)\sum_{k=1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{i\infty}\\ &= (1-\lambda)\mathbb{E}[\theta \mid \mathcal{H}^\sigma_{i\infty}]\\ &\quad + (1-\lambda)\lambda\sum_{k=1}^{\infty}\lambda^{k-1}\sum_{j\neq i}\mathbb{E}[\bar{\theta}^k_{j\infty}\mathcal{H}^\sigma_{i\infty}]/(n-1)\\ &= (1-\lambda)\mathbb{E}[\theta \mid \mathcal{H}^\sigma_{i\infty}]\\ &\quad + \lambda\sum_{j\neq i}\mathbb{E}\left[(1-\lambda)\sum_{k=1}^{\infty}\lambda^{k-1}\bar{\theta}^k_{j\infty} \mid \mathcal{H}^\sigma_{i\infty}\right]/(n-1)\\ &= (1-\lambda)\mathbb{E}[\theta \mid \mathcal{H}^\sigma_{i\infty}] + \lambda\sum_{j\neq i}\mathbb{E}[\sigma_{j\infty} \mid \mathcal{H}^\sigma_{i\infty}]/(n-1),\end{aligned}$$

where in the second line the order of integrals are changed using Fubini's theorem and the fact that by compactness of Θ the sums are finite. By Theorem 3, $\sigma_{i\infty} = \sigma_{j\infty}$, $\mathbb{P}$-a.s., for all $i, j \in N$. Therefore, the above equation implies that

$$\sigma_{i\infty} = \mathbb{E}[\theta \mid \mathcal{H}^\sigma_{i\infty}],$$

and so

$$\mathbb{E}[\theta \mid \mathcal{H}^\sigma_{i\infty}] = \mathbb{E}[\theta \mid \mathcal{H}^\sigma_{j\infty}], \quad \forall\, i, j \in N.$$

Let $\mathcal{S}$ be the smallest sub-sigma algebra of $\mathcal{B}$ that makes $s = (s_1, \ldots, s_n)$ measurable. By Lemma 1, for any sub-sigma algebras $\mathcal{S}_1$, $\mathcal{S}_2$ of $\mathcal{S}$, the set

$$\begin{aligned}&\mathbf{M}(\mathcal{S}_1, \mathcal{S}_2)\\ &= \{P \in \mathbf{P}\colon E_P[\theta \mid \mathcal{S}_1] = E_P[\theta \mid \mathcal{S}_2] \neq E_P[\theta \mid \mathcal{S}_1 \vee \mathcal{S}_2]\}\end{aligned}$$

is nowhere dense in $(\mathbf{P}, d)$. Since $\mathcal{S}$ is finite, it has only a finite number of sub-sigma algebras. Therefore, the set

$$\begin{aligned}\mathbf{M} = \big\{&P \in \mathbf{P}\colon E_P[\theta \mid \mathcal{S}_1] = E_P[\theta \mid \mathcal{S}_2] \neq E_P[\theta \mid \mathcal{S}_1 \vee \mathcal{S}_2]\\ &\text{for some } \mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{S}\big\}\end{aligned}$$

is also nowhere dense in $(\mathbf{P}, d)$. Let $\mathcal{H}^\sigma_{ij\infty} = \mathcal{H}^\sigma_{i\infty} \vee \mathcal{H}^\sigma_{j\infty}$. Since the information available to the players in any stage of the game is no more than the information contained in their private signals, $\mathcal{H}^\sigma_{i\infty} \subseteq \mathcal{S}$ for all i and σ. Thus, the set

$$\begin{aligned}\mathbf{N}_{ij} = \{&P \in \mathbf{P}\colon E_P[\theta \mid \mathcal{H}^\sigma_{i\infty}] = E_P[\theta \mid \mathcal{H}^\sigma_{j\infty}]\\ &\neq E_P[\theta \mid \mathcal{H}^\sigma_{ij\infty}] \text{ for some } \sigma\}\end{aligned}$$

is also nowhere dense in $(\mathbf{P}, d)$. On the other hand, we showed above that $E_P[\theta \mid \mathcal{H}^{\sigma^P}_{i\infty}] = E_P[\theta \mid \mathcal{H}^{\sigma^P}_{j\infty}]$ for all $P \in \mathbf{P}$. Therefore, for P in a residual subset of $\mathbf{P}$, $E_P[\theta \mid \mathcal{H}^{\sigma^P}_{i\infty}] = E_P[\theta \mid \mathcal{H}^{\sigma^P}_{ij\infty}]$. By a similar argument, for P in a residual subset of $\mathbf{P}$, $E_P[\theta \mid \mathcal{H}^{\sigma^P}_{i\infty}] = E_P[\theta \mid \mathcal{H}^{\sigma^P}_{ijk\infty}]$, where $\mathcal{H}^\sigma_{ijk\infty} = \mathcal{H}^\sigma_{ij} \vee \mathcal{H}^\sigma_k$. More generally, for P in a residual subset of $\mathbf{P}$, $E_P[\theta \mid \mathcal{H}^{\sigma^P}_{i\infty}] = E_P[\theta \mid \mathcal{H}^{\sigma^P}_{\infty}]$ for all i. This conclusion, together with Theorem 3 and the argument in the first paragraph of this proof, completes the proof of the theorem.

8. Conclusion

This paper studies a dynamic game in which a number of short-run players repeatedly play a symmetric strictly supermodular game of incomplete information. Each short-run player inherits the beliefs of a player playing in the previous stage while also observing the last stage actions of the players in his social neighborhood. Each player's actions reveal information used by other players to revise their beliefs and, hence, their actions. We prove formal results regarding the asymptotic outcomes obtained when agents play the actions prescribed by the Markov Bayesian equilibrium. In particular, we show that players reach consensus in their actions and payoffs if the observation network is connected. We also show that, when the utility functions are quadratic, the consensus action is generically optimal. We also provide extensions of our consensus result to a setting with time-varying and random networks and endogenously generated signals, and we illustrate the logic of our results through examples. Finally, we provide examples of games used in engineering and economics to which our results apply.

The players in this paper are assumed to be short-run and hence myopic. However, we expect our results to generalize to the case of forward-looking agents if attention is restricted to Markovian strategies. In symmetric supermodular games, the players' interests are fully aligned, and so they benefit from sharing the information available to them with the rest of the population. But short-run players cannot capture any of the benefits of sharing their information. Nonetheless, as our results demonstrate, consensus and information aggregation are eventually obtained. With forward-looking agents, the players' incentive to inform their peers provides an additional force that makes consensus and information aggregation, if anything, more likely. We intend to investigate this direction in future research.

Acknowledgments

The authors thank the editors and three anonymous referees for careful reading of the manuscript and helpful comments. They are also grateful to Hamid Sabourian, Alireza Tahbaz-Salehi, Juuso Toikka, Rakesh Vohra, and Muhamet Yildiz for useful feedback and suggestions. The authors acknowledge financial support from the Air Force Office of Scientific Research (MURI Award FA9550-10-1-0567) and the Army Research Office (MURI Award W911NF-12-1-0509).

Endnotes

1. Strategic interactions in which players want to coordinate their actions are best modeled by supermodular games in which the players' actions are strategic complements. Supermodular games have a deep and interesting theory that has been developed by, among others, Milgrom and Roberts (1990), Topkis (1998), Van Zandt and Vives (2007), and Van Zandt (2010). For an excellent survey of some of the theory and applications of supermodular games see Vives (2005).
2. Since the players are short-lived, the set of equilibrium outcomes generated by pure MBEs coincides with the set of outcomes generated by pure perfect Bayesian equilibria (PBE). However, even with myopic players, the set of mixed-strategy MBEs could be strictly smaller than the set of mixed-strategy PBEs since conditioning on history provides the players with more opportunities to coordinate their actions.
3. The imitation principle was first introduced by Gale and Kariv (2003) in a social learning model without strategic interactions.
4. See, for instance, Morris and Shin (2002) and Angeletos and Pavan (2007, 2009), in which the authors study the role of the provision of public information on social welfare, and Calvó-Armengol and Beltran (2009) and Bramoullé et al. (2014), in which equilibria of general quadratic network games are analyzed.
5. If f is twice differentiable, strict supermodularity is equivalent to requiring that $\partial^2 f/\partial x_i \partial x_j > 0$ for all $1 \leqslant i < j \leqslant n$. For more on the theory of supermodular games and their applications in game theory and economics, see Topkis (1998).
6. If $\lambda \geqslant 1$, the game can have multiple equilibria. Yet, it still belongs to the class of symmetric strictly supermodular games, and so by Theorem 1 and Corollary 1, players asymptotically reach consensus in their actions and payoffs.
7. We show that indeed the game can be solved by the iterated elimination of dominated strategies. The iterative elimination is a process that models the thinking process of rational player *it* where *it* recursively refines his belief on the actions that other players could take given the assumption that they are also rational. For the current model, this iterative elimination eventually leads to a single strategy profile being selected.
8. Given a topological space X, a subset A of X is a *first category* or meager set if it can be expressed as the union of countably many nowhere dense subsets of X. The complement of a first category set is called a *residual* set.
9. Given measurable spaces $(X, \mathscr{X})$ and $(Y, \mathscr{Y})$, a function $f: X \times \mathscr{Y} \to [0,1]$ is called a transition probability from X to Y if (i) for any given $x \in X$, $f(x)[\cdot]$ is a probability distribution over $(Y, \mathscr{Y})$; and (ii) given any measurable set $B \in \mathscr{Y}$, the function $x \mapsto f(x)[B]$ is measurable.
10. See, for instance, Theorem 49 in Chapter 4 of Pollard (2002).
11. 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.
12. For a recursive characterization of the players' equilibrium actions in the Bayesian quadratic network games similar to the one studied in Example 3, see the complementary paper by the authors (Eksin et al. 2014).
13. The notion of rationalizability we use is that of interim correlated rationalizability (ICR) introduced by Dekel et al. (2007).
14. This follows a result of Landers and Rogge (1976) (cf. Theorem 3.3. of Crimaldi and Pratelli 2005).

References

Acemoglu D, Jackson MO (2015) History, expectations, and leadership in the evolution of social norms. *Rev. Econom. Stud.* 82(2):423–456.
Acemoglu D, Jackson MO (2014) Social norms and the enforcement of laws. NBER Working Paper 20369, National Bureau of Economic Research, Cambridge, MA.
Acemoglu D, Bimpikis K, Ozdaglar A (2014) Dynamics of information exchange in endogenous social networks. *Theoret. Econom.* 9(1): 41–97.
Acemoglu D, Dahleh MA, Lobel I, Ozdaglar A (2011) Bayesian learning in social networks. *Rev. Econom. Stud.* 78(4):1201–1236.
Altman E, Altman Z (2003) *S*-modular games and power control in wireless networks. *IEEE Trans. Automatic Control* 48(5):839–842.
Amador M, Weill PO (2012) Learning from private and public observations of others' actions. *J. Econom. Theory* 147(3):910–940.
Angeletos GM, Pavan A (2007) Efficient use of information and social value of information. *Econometrica* 75(4):1103–1142.
Angeletos GM, Pavan A (2009) Policy with dispersed information. *J. Eur. Econom. Assoc.* 7(1):11–60.
Angeletos GM, Hellwig C, Pavan A (2007) Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica* 75(3):711–756.
Arieli I, Mueller-Frank M (2015) Inferring beliefs from actions, Social Science Research Network Working paper.
Aumann RJ (1976) Agreeing to disagree. *Ann. Statist.* 4(6):1236–1239.
Bala V, Goyal S (1998) Learning from neighbours. *Rev. Econom. Stud.* 65(3):595–621.
Banerjee A (1992) A simple model of herd behavior. *Quart. J. Econom.* 107(3):797–817.
Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Political Econom.* 100(5):992–1026.
Bonatti A, Cisternas G, Toikka J (2015) Dynamic oligopoly with incomplete information. Working paper, Sloan School of Management, MIT, Cambridge, MA.
Borkar V, Varaiya P (1982) Asymptotic agreement in distributed estimation. *IEEE Trans. Automatic Control* 27(3):650–655.
Bramoullé Y, Kranton R, D'Amours M (2014) Strategic interaction and networks. *Amer. Econom. Rev.* 104(3):898–930.
Calvó-Armengol A, Beltran JdM (2009) Information gathering in organizations: Equilibrium, welfare, and optimal network structure. *J. Eur. Econom. Assoc.* 7(1):116–161.
Çelen B, Kariv S (2004) Observational learning under imperfect information. *Games Econom. Behav.* 47(1):72–86.
Cooper R (1999) *Coordination Games* (Cambridge University Press, New York).
Cooper R, John A (1988) Coordinating coordination failures in Keynesian models. *Quart. J. Econom.* 103(3):441–463.

Crimaldi I, Pratelli L (2005) Two inequalities for conditional expectations and convergence results for filters. *Statist. Probab. Lett.* 74(2): 151–162.
Dekel E, Fudenberg D, Morris S (2007) Interim correlated rationalizability. *Theoret. Econom.* 2(1):15–40.
Eksin C, Molavi P, Ribeiro A, Jadbabaie A (2014) Bayesian quadratic network game filters. *IEEE Trans. Signal Processing* 62(9): 2250–2264.
Foster DP, Young HP (2003) Learning, hypothesis testing, and Nash equilibrium. *Games Econom. Behav.* 45(1):73–96.
Foster FD, Viswanathan S (1996) Strategic trading when agents forecast the forecasts of others. *J. Finance* 51(4):1437–1478.
Gale D, Kariv S (2003) Bayesian learning in social networks. *Games Econom. Behav.* 45(2):329–346.
Jackson MO, Kalai E (1997) Social learning in recurring games. *Games Econom. Behav.* 21(1):102–134.
Jadbabaie A, Molavi P, Sandroni A, Tahbaz-Salehi A (2012) Non-Bayesian social learning. *Games Econom. Behav.* 76(1):210–225.
Jordan JS (1991) Bayesian learning in normal form games. *Games Econom. Behav.* 3(1):60–81.
Jordan JS (1995) Bayesian learning in repeated games. *Games Econom. Behav.* 9(1):8–20.
Kalai E, Lehrer E (1993) Rational learning leads to Nash equilibrium. *Econometrica* 61(5):1019–1045.
Kyle AS (1985) Continuous auctions and insider trading. *Econometrica* 53(6):1315–1335.
Landers D, Rogge L (1976) Convergence of conditional p-means given a σ-lattice. *Ann. Probab.* 4(1):147–150.
Lauermann S, Wolinsky A (2015) Search with adverse selection. *Econometrica.* Forthcoming.
Lobel I, Sadler E (2015a) Preferences, homophily, and social learning. Social science research network. *Oper. Res.*, ePub ahead of print April 7, http://dx.doi.org/10.1287/opre.2015.1364.
Lobel I, Sadler E (2015b) Information diffusion in networks through social learning. *Theoret. Econom.* 10(3):807–851.
Milgrom P, Roberts J (1990) Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* 58(6): 1255–1277.
Morris S, Shin HS (2002) Social value of public information. *Amer. Econom. Rev.* 92(5):1521–1534.
Mossel E, Tamuz O (2010) Efficient Bayesian learning in social networks with Gaussian estimators, arxiv Preprint arXiv:1002.0747v2.
Mossel E, Sly A, Tamuz O (2011) On agreement and learning, arxiv preprint arXiv:1207.5895v1.
Mueller-Frank M (2013) A general framework for rational learning in social networks. *Theoret. Econom.* 8(1):1–40.
Nachbar JH (1997) Prediction, optimization, and learning in repeated games. *Econometrica* 65(2):275–309.
Obstfeld M (1996) Models of currency crises with self-fulfilling features. *Eur. Econom. Rev.* 40(3):1037–1047.
Ostrovsky M (2012) Information aggregation in dynamic markets with strategic traders. *Econometrica* 80(6):2595–2647.
Pollard D (2002) *A User's Guide to Measure Theoretic Probability*, Vol. 8 (Cambridge University Press, New York).
Rosenberg D, Solan E, Vieille N (2009) Informational externalities and emergence of consensus. *Games Econom. Behav.* 66(2):979–994.
Rostek M, Weretka M (2012) Price inference in small markets. *Econometrica* 80(2):687–711.
Shin HS, Williamson T (1996) How much common belief is necessary for a convention? *Games Econom. Behav.* 13(2):252–268.
Smith L, Sørensen P (2000) Pathological outcomes of observational learning. *Econometrica* 68(2):371–398.
Topkis DM (1998) *Supermodularity and Complementarity* (Princeton University Press, Princeton, NJ).
Van Zandt T (2010) Interim Bayesian Nash equilibrium on universal type spaces for supermodular games. *J. Econom. Theory* 145(1):249–263.
Van Zandt T, Vives X (2007) Monotone equilibria in Bayesian games of strategic complementarities. *J. Econom. Theory* 134(1):339–360.
Vives X (2005) Complementarities and games: New developments. *J. Econom. Literature* 43(2):437–479.
Vives X (2010) *Information and Learning in Markets: The Impact of Market Microstructure* (Princeton University Press, Princeton, NJ).
Wolinsky A (1990) Information revelation in a market with pairwise meetings. *Econometrica* 58(1):1–23.

Pooya Molavi is a Ph.D. student in economics at Massachusetts Institute of Technology. He received a Ph.D. in electrical and systems engineering in 2013 from the University of Pennsylvania. His research interests are in the areas of economic theory and macroeconomics.

Ceyhun Eksin is a postdoctoral researcher in the Georgia Institute of Technology affiliated with Schools of Biology, and Electrical and Computer Engineering. He received a Ph.D. in Electrical and Systems Engineering from the University of Pennsylvania. His research interests are in the areas of signal processing, control theory and game theory. His current work focuses on modeling and analysis of networked multi-agent systems in biology, communication, and energy.

Alejandro Ribeiro is the Rosenbluth Associate Professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania. His research interests are in the applications of statistical signal processing to the study of networks and networked phenomena. His current research focuses on wireless networks, network optimization, learning in networks, networked control, robot teams, and structured representations of networked data structures. He received the 2012 S. Reid Warren, Jr. Award presented by the University of Pennsylvania's undergraduate student body for outstanding teaching.

Ali Jadbabaie is the Alfred Fitler Moore Professor of Network Science in the Department of Electrical and Systems Engineering at the University of Pennsylvania. He holds secondary appointments in the Departments of Computer and Information Science and Operations and Information Management. He is currently on leave from the University of Pennsylvania at Massachusetts Institute of Technology as the interim director of the Sociotechnical Systems Research Center, also serving as associate director of a new research and educational effort at MIT focused on information and decision systems, social and institutional behavior, and statistics.