

A general framework for rational learning in social networks

MANUEL MUELLER-FRANK

Department of Economics, University of Oxford

This paper provides a formal characterization of the process of rational learning in social networks. Agents receive initial private information and select an action out of a choice set under uncertainty in each of infinitely many periods, observing the history of choices of their neighbors. Choices are made based on a common behavioral rule. Conditions under which rational learning leads to global consensus, local indifference, and local disagreement are characterized. In the general setting considered, rational learning can lead to pairs of neighbors selecting different actions once learning ends while not being indifferent among their choices.

The effect of the network structure on the degree of information aggregation and speed of convergence is also considered, and an answer to the question of optimal information aggregation in networks is provided. The results highlight distinguishing features between properties of Bayesian and non-Bayesian learning in social networks.

KEYWORDS. Learning, social networks, common knowledge, consensus, speed of convergence, optimal information aggregation.

JEL CLASSIFICATION. D82, D83, D85.

1. INTRODUCTION

Social networks have a very important function as a source of information. Individuals constantly communicate with their social peers and use the information obtained through their interactions when forming opinions and making decisions. Within the economic literature, the importance of social networks is widely recognized. The role of social networks for employment outcomes,¹ technology adoption,² models of collective

Manuel Mueller-Frank: manuel.mueller-frank@economics.ox.ac.uk

I am grateful to Eddie Dekel, Alessandro Pavan, and Marciano Siniscalchi for discussions and guidance in the preparation of the paper. I benefited from conversations with Itai Arieli, David Austen-Smith, Andrei Barbos, Daniel Diermeier, Jeff Ely, Douglas Gale, Andrea Galeotti, Florian Herold, Matthew Jackson, Nenad Kos, Margaret Meyer, Mallesh Pai, Jakub Steiner, and Asher Wolinsky. I would also like to thank the co-editor, a number of anonymous referees, and seminar audiences at Bocconi University, University Carlos III, CEU, Ecole Polytechnique, University of Essex, INSEAD, LSE, Northwestern University, University of Ottawa, and University of Oxford.

¹See, for example, Boorman (1975), Calvó-Armengol and Jackson (2004), Ioannides and Loury (2004), Munshi (2003), and Topa (2001).

²See, for example, Besley and Case (1994), Foster and Rosenzweig (1995), Munshi (2004), and Udry and Conley (2001).

political action,³ and bargaining outcomes⁴ has been established.⁵ However, the literature providing proof of the significance of social networks in economic settings is mostly empirical or considers only static environments. Thus there is a lack of understanding of the formal learning process of rational individuals in social networks and how the behavior and opinions of individuals evolve over time if they interact repeatedly. The following key questions are addressed in this paper.

1. In a general setting with rational agents, what are the dynamics of the informational structure and its properties once learning ends?
2. Under which circumstances does rational learning lead to consensus, indifference, or disagreement among the actions chosen? The indifference and disagreement could be local—across neighbors—or global—across all agents in the network.
3. How does the network structure affect both the quality and the speed of information aggregation?

The first of these questions is answered independently by [Rosenberg et al. \(2009\)](#) in a setting where agents select utility maximizing actions in each period. My approach is somewhat different. In the tradition of the papers on knowledge and consensus such as [Cave \(1983\)](#), [Bacharach \(1985\)](#), [Parikh and Krasucki \(1990\)](#), and [Krasucki \(1996\)](#), I do not restrict attention to the expected utility setting, but consider the following framework of repeated choice under uncertainty. Agents are organized in an exogenously given (undirected) network that is common knowledge among all agents. Initially, agents receive private information given by the true cell of their partition and, subsequently, they simultaneously select a choice out of a choice set in each of infinitely many periods. The network structure determines the observability of actions: agents observe the history of choices of their neighbors. In this model, learning refers to a refinement of the information sets of agents, where the information set of an agent is the smallest subset of the state space that the agent knows to contain the true state of the world. Learning is based on the inferences agents make regarding the information sets of their neighbors, depending on the history of actions they observe.

Actions are chosen based on a common behavioral rule that assigns sets of optimal actions to every possible information set. This choice correspondence is required to be union consistent. Union consistency is a behavioral condition across mutually exclusive events and stems from the literature on knowledge and consensus.⁶ If the same action is optimal for each of a disjoint collection of information sets, then union consistency requires this action to be optimal for the information set given by the union of the information sets in the collection.

³See, for example, [Chwe \(2000\)](#). In the sociology literature, [Opp and Gern \(1993\)](#) and [Snow et al. \(1980\)](#) explore the importance of social networks for political participation.

⁴See [Manea \(2011\)](#), and [Wang and Wen \(2002\)](#).

⁵An extensive survey on the relevance of social networks from an economic perspective is given by [Jackson \(2008\)](#).

⁶See, for example, [Bacharach \(1985\)](#), [Cave \(1983\)](#), and [Krasucki \(1996\)](#).

This framework captures, but is not limited to, an expected utility setting, where agents share a common prior over the state space as well as a common utility function without payoff externalities and where, in each period, agents select the action that maximizes their expected utility in that period conditioning on their information.⁷

Rational learning requires all agents in every period to consider the set of possible information sets of all other agents and how their choices impact the information sets of their neighbors in the subsequent period. In increasingly large networks, this becomes an increasingly complex task, especially in incomplete networks.⁸ Due to these complexities, most of the existing literature on learning and evolution of behavior and opinions in social networks assumes boundedly rational agents. Examples can be found in Bala and Goyal (1998), DeMarzo et al. (2003), Golub and Jackson (2010), and Acemoglu et al. (2010).

Despite the practical difficulties and complexities, I show that the learning process and the resulting informational structure can be characterized in a quite simple and intuitive way. The rational learning process modeled here is based on the learning process in Geanakoplos and Polemarchakis (1982). They consider a special case of my general framework, in which agents repeatedly announce their posteriors of an uncertain event in a complete network consisting of two agents.

1.1 *Summary of results*

Both the literature on non-Bayesian learning as well as the literature on knowledge and consensus focus on consensus. In this paper, I show that in a richer setting, allowing for general network structures and choice correspondences, consensus and even local indifference can fail to occur under rational learning and union consistent choice correspondences. Local indifference denotes the case where any action an agent selects once learning ends is optimal for all his neighbors. The main contributions of the paper are presented in the form of four theorems.

The first theorem establishes that if the common behavioral rule is union consistent, then local indifference holds once learning ends. If two agents disagree by choosing different actions once learning ends, they are indifferent and could swap their actions as they would be equally well off. Rational learning in networks can, therefore, lead to heterogeneous choices.

The second theorem provides an asymptotic local indifference result for cases where learning does not end in finite time. The sufficient conditions I characterize in the general framework translate to the expected utility framework in the following way: if all cells of the join of partitions have positive probability and utility functions are bounded, then every action an agent selects infinitely often is optimal for his neighbors in the limit.⁹

⁷The case considered in Rosenberg et al. (2009) is not a special case of my general framework, as they allow for strategic behavior and uncountable signal space.

⁸In complete networks, all agents are neighbors of each other. Incomplete networks are not complete and, therefore, the history of choices is not common knowledge among all agents.

⁹The join of a collection of partitions is the coarsest common refinement of the partitions.

The first two theorems rely on the assumption of common knowledge of strategies—common knowledge of which action every agent selects out of the set of optimal actions for every possible information set. Therefore, the network structure does not affect the validity of the local indifference result under common knowledge of strategies. Relaxing the assumption of common knowledge of strategies and assuming only common knowledge of rationality, i.e., common knowledge of all agents following the same common behavioral rule, I show that the validity of the local indifference result depends on the network structure. In incomplete networks, it generally fails, while [Theorem 3](#) establishes that, under a union consistent choice rule, in complete networks, common knowledge of rationality is sufficient for local indifference. This is the most striking result mentioned so far as it provides a scenario where rational learning can lead to two neighbors selecting different actions once learning ends while not being indifferent among their choices.¹⁰ This is a new insight to the literatures on non-Bayesian learning, Bayesian learning, and knowledge and consensus.

To address the question of the effect of the network structure on the degree of information aggregation and duration to consensus, I consider the special case where agents share a common prior and repeatedly announce their posterior belief of an uncertain event to their neighbors. How precise is the private information of all agents incorporated in the eventual consensus belief? One's intuition might suggest that complete networks should always do at least as well as incomplete networks. However, I show this intuition to be false. I provide an example of a complete network that is dominated by an incomplete network in terms of quality of information aggregation.

[Theorem 4](#) establishes that (i) generically, the private information of all agents is incorporated in the eventual consensus belief and (ii) the duration to consensus is a function of the diameter of the graph: the larger is the diameter, the longer it takes for consensus to occur. Therefore, all connected networks are generically equivalent in terms of quality of information aggregation and differ only in their duration to consensus. [Theorem 4](#) also shows that, generically, agents have no information gain from untruthfully announcing their belief and, therefore, there is no strategic incentive to lie.

1.2 Discussion of results

From a normative perspective, the present paper yields the tools to address the question of optimal information aggregation in institutions. In a followup paper ([Mueller-Frank 2012](#)), the rational learning framework provided in this paper is used to establish superiority of Bayesian communication structures over non-Bayesian communication structures in terms of quality of information aggregation.

Due to the complexity of inferences, it is hard to imagine fully rational learning occurring in the real world, even though, for small networks, it cannot be excluded on the basis of complexity alone. From a positive perspective, the fully rational case analyzed in this paper serves as a counterfactual benchmark. The advantage of the Bayesian analysis in comparison to the existing non-Bayesian learning models is that it is not ad

¹⁰Furthermore, in the expected utility setting, this local disagreement can occur on a subset of the state space that has positive prior probability.

hoc. The weakness of the non-Bayesian literature, as DeMarzo et al. (2003) and Golub and Jackson (2010), is its lack of generality and reliance on a particular functional form: weighted averages. The common ground with the non-Bayesian learning literature is that Bayesian learning generally, but not always, leads to consensus or local indifference. In non-Bayesian models, this occurs based on simple updating rules, while in Bayesian models, consensus is the product of highly complex inferences. But this paper also highlights important distinctions between the Bayesian and non-Bayesian approach as the possibility of heterogeneous choices in the limit for the Bayesian case and the different effects changes of the network have on the speed of convergence.

The rest of the paper is organized as follows. In the next section, I introduce my general framework and briefly describe two special cases. In Section 3, I provide a simple example for the workings of the learning process, and characterize the general learning process and the resulting informational structure. Section 4 presents my local indifference result once learning ends. In Section 5, I analyze the case where learning does not end in finite time, introduce the concept of a dominant set, and present Theorem 2. In Section 6, I contrast common knowledge of rationality with common knowledge of strategies and present Theorem 3. In Section 7 I consider the effect of the network structure on the quality of information aggregation and duration to consensus. Theorem 4 and an example for superiority of incomplete networks, in the sense of information aggregation, is presented. Section 8 compares my framework and results to the existing literature on knowledge and consensus, and Bayesian and non-Bayesian learning in networks. Section 9 concludes. The Appendix presents proofs, lemmas, and examples that are omitted in the main text.¹¹

2. THE FRAMEWORK

2.1 Synopsis

There is a finite set of agents $V = \{1, \dots, v\}$ who face uncertainty, represented by a measurable space $(\Omega, \mathcal{F})$, where Ω is the state space and $\mathcal{F}$ is a σ -algebra of subsets of Ω .¹² At the beginning of time, $t = 1$, one state ω is realized. Each agent i has private information about the realized state given by his partition $\mathcal{P}_i$. If the realized state of the world is ω , then i knows that a state in $\mathcal{P}_i(\omega)$ has occurred. The set of partitions of all players $\{\mathcal{P}_i\}_{i \in V}$ is commonly known.¹³ Time is discrete.

Players form part of a social network G . In each of an infinite number of periods, all players simultaneously select an action out of a choice set $\mathcal{A}$. The network determines the observability of actions: agents observe the history of actions of their neighbors. They select actions as a function of the available information. The information set of an agent plays a crucial role in the subsequent analysis.

¹¹An additional appendix, where I prove statements made in the paper, consider an extension, and provide additional examples, is available in a supplementary file on the journal website, <http://econtheory.org/supp/1015/supplement.pdf>.

¹²I assume $\mathcal{F}$ to be generated by the join of partitions.

¹³Common knowledge of information partitions is not an assumption, but a tautology. See Aumann (1999).

The information set of agent i in period t denotes the smallest subset of the state space that i knows to contain the true state of the world. It is a function of the observables of agent i , the cell of his partition, and the history of choices of his neighbors up to period t . Therefore, the information sets are (weakly) shrinking over time: the information set in period t is a subset of the information set in period $t - 1$. Learning in this setting refers to a refinement of information sets of agents based on the history of actions they observe.

2.2 The social network

The social network is represented by an undirected graph G . A graph is a pair of sets $G = (V, E)$ such that $E \subset [V]^2$. The elements of V are nodes of the graph and the elements of E are the edges of the graph.

Node i in G represents player i . The neighborhood of agent i , N_i , contains all agents who are connected to i by an edge in G :

$$N_i = \{j \in V : ij \in E\}.$$

An undirected graph has the following symmetric property: if agent j is contained in agent i 's neighborhood, then agent i is contained in agent j 's neighborhood. The common neighborhood of two players i and j is denoted by N_{ij} and consists of the set of agents who are neighbors of both i and j :

$$N_{ij} = N_i \cap N_j.$$

A graph G is connected if, for all nodes i, j , there exists a sequence of nodes $k_1, \dots, k_l$, where $k_1 = i$ and $k_l = j$, such that $k_{f+1} \in N_{k_f}$ for $f = 1, \dots, l - 1$. A graph is complete if, for all nodes $i, j \in V$, we have $ij \in E$. A graph is incomplete if it is not complete. I assume that the social network G is common knowledge.

2.3 The common choice correspondence and strategies

Agents select actions based on a common choice correspondence c :

$$c: \mathcal{F} \rightrightarrows A.$$

The choice correspondence assigns a subset of the choice set A to each information set $I \in \mathcal{F}$. The actions assigned to an information set can be thought of as the actions that are optimal given the information set.

A pure strategy s_i for player i is a function that assigns a single action to each information set $I \in \mathcal{F}$, $s_i: \mathcal{F} \rightarrow A$, such that $s_i(I) \in c(I)$ for all $I \in \mathcal{F}$. The choice correspondence assigns a set of optimal actions to each information set, and the strategy s_i selects one of them. The strategies of all players are assumed to be common knowledge.

The history of play at time t is denoted as $h^t = (a^1, \dots, a^{t-1})$, $a^k \in A^v$ for $k = 1, \dots, t - 1$. The history that player i observes in a given period t is denoted as h_i^t , and

consists of the history of actions chosen by i and his neighbors up to period t . The history of actions that both player i and j observe is denoted as h_{ij}^t , and consists of the history of choices up to period t of agents i and j , and of all agents l that are neighbors of both i and j .

2.4 Special cases of the general framework

The general framework captures two prominent settings as special cases: the probability announcement setting and the expected utility setting.¹⁴ For the probability announcement setting, let $\mathcal{P}_i$ be finite for all i and let p be a common probability measure on $\mathcal{F}$. Suppose that all elements of the join of partition have positive probability. Let all agents be concerned with the likelihood of some uncertain event $Q \in \mathcal{F}$, and define $c(\cdot)$ as

$$c(I) = \frac{p(I \cap Q)}{p(I)} \quad \text{for } I \in \mathcal{F}.$$

In this setting, agents announce their conditional probability of the event Q in every period. This is a special case of a choice function rather than a correspondence, as a unique probability is assigned to each information set in $\mathcal{F}$. Note that for a thus defined choice rule, strategic announcements are excluded, i.e., agents are assumed to tell the truth. On first glance, this assumption seems to be restricting, but turns out to matter only for nongeneric probability measures. In [Theorem 4](#), I show that, generically, there is no incentive to lie as no information gain results from it.

For the expected utility setting, let $\mathcal{P}_i$ be countable for all i and let p be a common probability measure on $\mathcal{F}$. Suppose that all agents share a common utility function $u: A \times \Omega \rightarrow \mathbb{R}$ that is bounded and measurable for each $a \in A$. Let the choice correspondence $c(\cdot)$ be defined¹⁵ as

$$c(I) = \arg \max_{a \in A} E[u(a, \omega) | I] \quad \text{for } I \in \mathcal{F}.$$

The inherent assumptions of the thus defined choice correspondence are homogeneous agents in the sense of common values and that, in each period, agents maximize their expected utility of the given period, thus behaving nonstrategically in the way they select actions. An information gain in a later period could occur by misrepresenting one's information in a given period. The assumption of nonstrategic behavior is necessary for tractability reasons in the general framework. However, it seems that a variant of [Theorem 4](#) might hold in the expected utility setting. For generic (bounded) utility functions and probability measures on a finite state space Ω , it appears intuitive that the probability of a given action being optimal for more than one information set converges to zero with the number of actions going to infinity. Hence the incentive to lie vanishes with an increasing set of actions.

¹⁴Note that the probability announcement setting can be captured within the expected utility setting via quadratic loss functions.

¹⁵To assure existence of a maximum for infinite A , A needs to be a compact subset of a topological space $\mathbf{A}$ and u to be continuous in $\mathbf{A}$ for every ω .

In a setting where agents maximize their expected, discounted sum of stage utilities, the incentive to rationally deceive decreases with the value agents assign to future consumption. The more agents discount future payoffs, the less likely they are to deviate from the nonstrategic action. [Rosenberg et al. \(2009\)](#) independently consider the asymptotic properties of behavior in the expected utility setting under Bayesian learning and allow for strategic behavior.¹⁶ They show that the local indifference result holds almost surely.

3. THE LEARNING PROCESS

Agents progressively learn over time through the inferences they make from the history of choices of their neighbors. The information set of an agent in period t is the smallest subset of the state space that the agent knows to contain the true state of the world. It is a function of the private observables of the agent, the cell of his partition, and the history of actions he observed up to period t .

Rational learning requires each agent to consider the set of possible information sets of each other agent in every period and how their choice in a given period impacts the information sets of their neighbors in the subsequent period. Compared to a complete network analysis like in [Geanakoplos and Polemarchakis \(1982\)](#), where the history of choices of all agents is common knowledge, the added difficulty in an incomplete network is that the privately observable component contributing to the information set of a given player is not given only by the cell of his partition, but also by the history of choices he (privately) observes.

I assume fully rational agents who make all possible inferences based on the history they observe. Their inference consists of direct inference regarding the realized partition cell of their neighbors, as well as, over time, indirect inference regarding the realized cell of all other agents. The crucial element of the learning process is the *information set* of an agent. The information set of agent i in period t , $I_i^t \in \mathcal{F}$, is the smallest subset of the state space that the agent knows to contain the true state of the world based on his private observables. Based on the history he observes, the agent learns to dismiss states as not consistent with the history. Therefore, the information set I_i^t is a subset of the information set I_i^{t-1} for every t and every state of the world.

The learning process requires agent i to make inferences regarding the realized partition cells of all agents based on his neighbors actions. To do so, he considers, for each of his neighbors j , their set of possible information sets $\mathcal{I}_j^t$ —the information sets of i that are consistent with the common observables of i and j . This set of possible information sets is a subset of the σ -algebra and hence the strategy of agent j , which is assumed to be common knowledge, assigns a unique action to every information set in $\mathcal{I}_j^t$. Based on the action chosen by agent j in period t , agent i can dismiss all information sets in $\mathcal{I}_j^t$ to which the strategy of j assigns a different action than the one chosen. The set of information sets in $\mathcal{I}_j^t$ that survive the elimination process in period t is denoted as $\mathcal{D}_j^t \subset \mathcal{I}_j^t$ and serves to refine i 's information set in period $t + 1$. In the following discussion, I provide an example of the learning process before presenting its formal structure.

¹⁶Please see [Section 8](#) for a discussion on the relation of Rosenberg et al. and this paper.

ω_1	ω_2
ω_3	ω_4

FIGURE 1. State space of the example.

3.1 Example of the learning process

Prior to providing the formal characterization of the learning process, let us consider a very simple example that eases the understanding of the informational structure and the learning process. Suppose that there are four states of the world, $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, and three agents, $V = \{1, 2, 3\}$. Let us represent the state space with the matrix in Figure 1.

Suppose that agent 1 learns the rows of the matrix, i.e., if state ω_1 or ω_2 is realized, he learns that a state in the first row is realized. Suppose that agent 2 has no private information; he learns only that the true state of the world lies in Ω , while agent 3 learns the columns of the matrix, i.e., if state ω_1 or ω_3 is realized, he learns that the true state of the world lies in the first column. We have

$$\mathcal{P}_1 = \{\omega_1, \omega_2; \omega_3, \omega_4\}$$

$$\mathcal{P}_2 = \{\omega_1, \omega_2, \omega_3, \omega_4\}$$

$$\mathcal{P}_3 = \{\omega_1, \omega_3; \omega_2, \omega_4\}.$$

The first period information sets of agents are given by the realized partition cells. Suppose all agents share an uniform common prior p over Ω and in each round, every agent selects an action $a_i^t \in \{\alpha, \beta\}$. Let the corresponding state dependent utility function be given by

$$u_i(a_i, \omega) = \begin{cases} 1 & \text{if } a_i = \alpha \text{ and } \omega = \omega_1 \\ 2 & \text{if } a_i = \beta \text{ and } \omega = \omega_4 \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that in every round t , each agent selects the action that maximizes his expected utility in round t , conditioning on his private information set $I_i^t \subset \Omega$:

$$c(I_i^t) = a_i^t = \max_{a \in \{\alpha, \beta\}} E[u_i(a, \omega) | I_i^t].$$

In case of indifference, suppose that all agents select action α . Finally, let the agents be organized in a line network with agent 2 being the center agent, $N_i = \{2\}$ for $i = 1, 3$ and $N_2 = \{1, 3\}$. Let us consider the first period action of agents given their information set. Agent 1 selects $a_1^1(I_1^1) = \alpha$ if he observes the first row, $I_1^1 = \{\omega_1, \omega_2\}$, and selects $a_1^1(I_1^1) = \beta$ if he observes the second row, $I_1^1 = \{\omega_3, \omega_4\}$. Agent 2 receives no private information and selects $a_2^1(I_2^1) = \beta$. Agent 3 selects $a_3^1(I_3^1) = \alpha$ if he observes the first column, $I_3^1 = \{\omega_1, \omega_3\}$, and selects $a_3^1(I_3^1) = \beta$ if he observes the second column, $I_3^1 = \{\omega_2, \omega_4\}$. Hence agent 1 and agent 3 reveal their realized partition cell through their first round choice to agent 2, who observes both of their announcements—a fact that is common knowledge among all agents. Suppose that the state ω_1 is realized, which leads to the first period

choice vector $\mathbf{a}^1 = (\alpha, \beta, \alpha)$. From the choice of agent 1, agent 2 learns that either ω_1 or ω_2 is realized, and from the choice of agent 3, he learns that either ω_1 or ω_3 is realized. Combining the two, he learns that the true state of the world is ω_1 . Note that due to the perfect separation of choices across partition cells of both agent 1 and agent 3, agent 2 learns the true state of the world in each state $\omega \in \Omega$ through the first period choice vector. The information sets of the agents in state ω_1 at the beginning of the second round conditional on the observed history $h_i^2(\omega_1)$ are given by

$$I_1^2(\omega_1, h_1^2(\omega_1)) = \{\omega_1, \omega_2\}$$

$$I_2^2(\omega_1, h_2^2(\omega_1)) = \{\omega_1\}$$

$$I_3^2(\omega_1, h_3^2(\omega_1)) = \{\omega_1, \omega_3\}.$$

Since the information sets of neither agent 1 nor 3 changed from the first to the second round, neither will their chosen action. To draw inference from agent 2's second period choice, agents 1 and 3 have to consider the possible information sets of agent 2. The information sets of agent 2 that agent 1 considers possible given the common history h_{12}^2 , i.e., the first period choices of both 1 and 2, are described by the set

$$\mathcal{I}_2^2(h_{12}^2(\omega_1); \omega_1) = \{\omega_1, \omega_2\}.$$

Based on the common history of 1 and 2, it is commonly known among them that the second round information set of agent 2 is either $I_2^2 = \omega_1$ or $I_2^2 = \omega_2$. For agent 1 to draw inference from the second period choice of agent 2, he has to consider the choice 2 makes in each of the possible information sets. But since agent 2 selects α in both states ω_1 and ω_2 , due to the tie-breaking rule, agent 1 can draw no further inference from the second round choice of agent 2. Similarly, agent 3 considers the set of possible information sets of agent 2 based on their common history

$$\mathcal{I}_2^2(h_{23}^2(\omega_1); \omega_1) = \{\omega_1, \omega_3\}.$$

Since agent 2 selects action α in both of the possible information sets (in ω_3 due to the tie-breaking rule), agent 3 can draw no further inference with regards to the realized information set of agent 2. For all $t \geq 2$, we have $\mathbf{a}^t = (\alpha, \alpha, \alpha)$.

3.2 Formal structure of learning process

Let me now present the formal structure of the learning process and introduce the notation used throughout the paper. At the beginning of stage $t = 1$, the information set $I_i^1(\omega)$ of each agent i is given by the true cell of his partition. In period $t = 1$, we have

$$I_i^1(\omega) = \mathcal{P}_i(\omega).$$

It is common knowledge among any pair of agents i and j that the true cell of their meet, $\bigwedge_{l=i,j} \mathcal{P}_l(\omega)$, is realized.¹⁷ Let P_i denote a cell in i 's partition, $P_i \in \mathcal{P}_i$. The set

¹⁷The meet of a pair of partitions is the finest common coarsening of the partitions.

of possible first stage information sets of agent j based on the common information $\bigwedge_{l=i,j} \mathcal{P}_l(\omega)$ of i and j is given by

$$\mathcal{I}_j^1(\omega; i) = \left\{ P_j \in \mathcal{P}_j : P_j \subset \left(\bigwedge_{l=i,j} \mathcal{P}_l(\omega) \right) \right\}.$$

The set of possible information sets of agent j from the perspective of agent i , $\mathcal{I}_j^1(\omega; i)$, consists of the cells P_j of j 's partition $\mathcal{P}_j$ that are contained in the realized cell of the meet of i and j , $\bigwedge_{l=i,j} \mathcal{P}_l(\omega)$. Each agent j selects his first period action a_j^1 according to his information set and strategy. His neighbor i makes the following inference regarding player j 's realized partition cell according to the action a_j^1 chosen:

$$\mathcal{D}_j^1(a_j^1; \omega; i) = \{P_j \in \mathcal{I}_j^1(\omega; i) : a_j^1 = s_j(P_j)\}.$$

The set $\mathcal{D}_j^1(a_j^1; \omega; i)$ is a subset of the set of possible first period information sets, $\mathcal{I}_j^1(\omega; i)$. The set $\mathcal{D}_j^1(a_j^1; \omega; i)$ is also the set of partition cells of j , P_j , contained in the set of possible information sets of j , $\mathcal{I}_j^1(\omega; i)$, to which the strategy, s_j , of agent j assigns the action a_j^1 . Therefore, the sets in $\mathcal{I}_j^1(\omega; i) \setminus \mathcal{D}_j^1(a_j^1; \omega; i)$ are dismissed by agent i based on the additional observable a_j^1 as not consistent with the first period action of agent j .

After observing the first period choices of his neighbors and making inference regarding the realized cells of their partitions, player i takes the intersection of the true cell of his partition with the sets $\bigcup \mathcal{D}_j^1(a_j^1; \omega; i)$ across all his neighbors $j \in N_i$ to compute his second stage information set.¹⁸ Agent i 's information set in period $t = 2$ is denoted as

$$I_i^2(\mathcal{P}_i(\omega), h_i^2(\omega)) = \mathcal{P}_i(\omega) \cap \bigcap_{j \in N_i} \bigcup \mathcal{D}_j^1(a_j^1; \omega; i).$$

In period t , the information set of agent i is given by

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) = \mathcal{P}_i(\omega) \cap \bigcap_{j \in N_i} \bigcup \mathcal{D}_j^{t-1}(a_j^{t-1}; h_{ij}^{t-1}(\omega); \omega).$$

The information set of agent i in period t consists of all states of the world that are contained in his partition cell, $\mathcal{P}_i(\omega)$, and in the union of possible information sets, $\bigcup \mathcal{D}_j^{t-1}(a_j^{t-1}, h_{ij}^{t-1}(\omega); \omega)$, of each of his neighbors j . Any pair of neighbors i and j shares a common history $h_{ij}^t(\omega)$ at the outset of period t given by the history of choices up to period t of i and j , and the history of choices of all agents l that are neighbors of both i and j . The set $\mathcal{I}_j^t(h_{ij}^t(\omega); \omega)$ consists of possible information sets of player j in period t that are consistent with the common observables of agents i and j . The set $\mathcal{I}_j^t(h_{ij}^t(\omega); \omega)$ is common knowledge among them and contains the true information set of player j . We have

$$\mathcal{I}_j^t(h_{ij}^t(\omega); \omega) = \left\{ I_j^t(P_j, \hat{h}_j^t) : \begin{array}{l} \hat{h}_{ij}^t = h_{ij}^t(\omega) \\ \exists I_j^{t-1} \in \mathcal{D}_j^{t-1}(a_j^{t-1}, h_{ij}^{t-1}(\omega); \omega) \text{ s.t. } I_j^t(P_j, \hat{h}_j^t) \subset I_j^{t-1} \end{array} \right\}.$$

¹⁸For a collection of sets $\mathcal{D}$, the notation $\bigcup \mathcal{D}$ denotes the union over all sets in $\mathcal{D}$.

The set of possible information sets of agent j in period t based on the realized common history $h_{ij}^t(\omega)$ of agent i and j consists of all information sets $I_j^t(P_j, \hat{h}_j^t)$ such that the private history of j is consistent with the realized common history $h_{ij}^t(\omega)$, and $I_j^t(P_j, \hat{h}_j^t)$ is consistent with the choice of agent j in period $t - 1$. Player j selects an action a_j^t in period t according to his strategy, which leads to a refinement of the set $\mathcal{I}_j^t(h_{ij}^t(\omega); \omega)$ by his neighbor i ,

$$\mathcal{D}_j^t(a_j^t; h_{ij}^t(\omega); \omega) = \{I_j^t \in \mathcal{I}_j^t(h_{ij}^t(\omega); \omega) : a_j^t = s_j(I_j^t)\},$$

where $\mathcal{D}_j^t(a_j^t; h_{ij}^t(\omega); \omega)$ is common knowledge among i and j , as it relies only on variables that are commonly known among them. At the beginning of stage $t + 1$, agent i processes the inferences made based on the choices of his neighbors in period t , resulting in his private information set in period $t + 1$ given by

$$I_i^{t+1}(\mathcal{P}_i(\omega), h_i^{t+1}(\omega)) = \mathcal{P}_i(\omega) \cap \bigcap_{j \in N_i} \bigcup \mathcal{D}_j^t(a_j^t; h_{ij}^t(\omega); \omega).$$

The alert reader may realize that the inference made by agent i regarding the information set of his neighbor j occurs out of a set of commonly known possible information sets, $\mathcal{I}_j^t(h_{ij}^t(\omega); \omega)$, even though player i 's private information might lead to an exclusion of some elements of $\mathcal{I}_j^t(h_{ij}^t(\omega); \omega)$. In particular, agent i can exclude all those information sets of player j that have an empty intersection with his information set $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))$. While the set of possible information sets that the inferences of agents are based on differs in the private and common observables case, the resulting information sets are the same. For a formal proof, please see the supplementary appendix. The proof is based on an induction argument. An understanding of the intuition behind the equivalence can be gained when considering the interaction in the first round and its effect on the information sets in round two: suppose agent i can exclude some of j 's commonly considered possible partition cells P_j' based on i 's private information. This means that i 's observed partition cell P_i has an empty intersection with P_j' , where P_j' is a subset of the realized cell of the meet of agents i and j . Suppose that P_j' is consistent with the first period action of agent j . However, it is easy to see that, nevertheless, P_j' has an empty intersection with agent i 's second round private information set based on inferences out of pairwise common observables, as the information set I_i^2 equals the intersection of i 's partition cell with the set of possible partition cells of each of i 's neighbors that are consistent with the history. Hence if $P_i \cap P_j' = \emptyset$, then $I_i^2(P_i, h_i^2(\omega)) \cap P_j' = \emptyset$.

Despite the practical difficulties of rational learning in arbitrary incomplete networks, I have shown that the formal structure of the rational learning process has a simple and intuitive form. The learning process defined here can be used in practice to analyze the evolution of information sets of agents and their behavior in specific situations. However, despite the simple and intuitive formal structure of the process, the computational burden can be very large in practice, as each agent effectively has to consider the information sets of all agents in each round and their respective optimal choices.

4. LOCAL INDIFFERENCE ONCE LEARNING ENDS

Having established the formal structure of the rational learning process, one can answer questions with regards to the evolution of behavior in social networks. In particular, does learning lead to convergence of choices? Before I present the first theorem, I define an important property of choice correspondences: union consistency.

DEFINITION 1. The correspondence $c: \mathcal{F} \rightrightarrows A$ is *union consistent* if for all disjoint collections of sets $\mathcal{G}, \mathcal{G} \subset \mathcal{F}$,

$$\bigcap_{G \in \mathcal{G}} c(G) \neq \emptyset \quad \Rightarrow \quad \bigcap_{G \in \mathcal{G}} c(G) = c\left(\bigcup_{G \in \mathcal{G}} G\right).$$

Union consistency stems from the literature on knowledge and consensus, where the property is also known as the sure thing principle.¹⁹ In the literature, the property is applied to choice functions rather than correspondences as in this paper. Union consistency is a behavioral assumption across mutually exclusive events: if an action is optimal for every information set in a collection of disjoint information sets, then it is also optimal for the information set equaling the union over all information sets in the disjoint collection. In addition, only actions that are optimal for all information sets in the disjoint collection are optimal under the information set equaling the union over all information sets in the disjoint collection. Union consistency holds in the probability announcement and expected utility setting if all elements of the join have positive probability and the utility function is bounded and measurable for each action.²⁰

As an example for union consistency, consider the case where a decision maker is uncertain about the weather. She wears a jacket if it rains and if it snows. Union consistency then requires her to wear a jacket if she knows only that it is either raining or snowing.

The following theorem gives my main result with regard to optimal behavior once learning ends.

THEOREM 1. *If the choice correspondence $c(\cdot)$ is union consistent and for a given state ω , there exists a finite t' such that for all $t \geq t'$ and for each pair of neighbors $i \in V, j \in N_i$*

$$\bigcup \mathcal{I}_i^t(h_{ij}^t(\omega); \omega) = \bigcup \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega),$$

then

$$s_j(I_j^t(\mathcal{P}_j(\omega), h_j^t(\omega))) = a_j^t \in c(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)))$$

for all $t \geq t'$.

Theorem 1 establishes that under rational learning in networks, any action an agent selects is optimal for all his neighbors once learning ends. I denote this property as local

¹⁹See, for example, Bacharach (1985), Cave (1983), and Krasucki (1996).

²⁰This can be easily established. A proof of the claim can be found in the supplementary appendix to an earlier version of this paper and is available on request.

indifference. Note that the indifference across actions chosen once learning ends holds only on a local level—across neighbors—rather than on a global level—across all agents in a connected network. A feature of rational learning in networks is that once learning ends, a pair of agents can select different actions while not being indifferent among them, as long as they are not neighbors of each other.²¹ The possibility of heterogeneous choices once learning ends differs from the results in the informational cascades literature: heterogeneous choices once learning ends are consistent with rational learning of homogeneous agents under repeated interaction, while they are not consistent with rational learning of homogeneous agents in the informational cascade model, where once learning ends, i.e., in a cascade, homogeneous agents select the same action.²²

Applying the result of [Theorem 1](#) to the probability announcement setting, we find that within finitely many communication rounds, all agents in a connected network agree on the probability of the uncertain event Q . This corollary constitutes a generalization of the “we can’t disagree forever” result from [Geanakoplos and Polemarchakis \(1982\)](#) to arbitrary connected social networks. Even though posterior probabilities are communicated only on a neighborhood level, rational learning leads to convergence of the posterior of all agents in a connected social network.

In addition to union consistency of the choice correspondence, one condition on the learning process is required: the union over the set of possible information sets that all agents assign to their neighbors remains constant from period t' onward. When informally stating that learning ends, I refer to this condition, which implies that the information sets of all agents remain constant from t' onward. The proof of the theorem is an application of [Proposition 1](#).

PROPOSITION 1. *Let $c: \mathcal{F} \rightrightarrows A$ be union consistent and let $\mathcal{G}, \mathcal{L} \subset \mathcal{F}$ be collections of disjoint sets such that*

$$\bigcup_{G \in \mathcal{G}} G = \bigcup_{L \in \mathcal{L}} L.$$

If there exist actions $a_l, a_g \in A$ such that $a_l \in c(L)$ for all $L \in \mathcal{L}$ and $a_g \in c(G)$ for all $G \in \mathcal{G}$, then

$$a_l \in c(G) \quad \text{for all } G \in \mathcal{G}$$

$$a_g \in c(L) \quad \text{for all } L \in \mathcal{L}.$$

To apply [Proposition 1](#) and thus prove the local indifference result once learning ends, we have to establish that its three conditions are satisfied: (i) for each $i \in V$, $j \in N_i$, we have $\mathcal{I}_i^t(h_{ij}^t(\omega); \omega)$ is a disjoint collection of sets, (ii) for a pair of neighbors i, j , the set of possible information sets $\mathcal{I}_i^t(h_{ij}^t(\omega); \omega)$, $\mathcal{I}_j^t(h_{ji}^t(\omega); \omega)$ comprises partitions of the same set, and (iii) for all agents i , once learning ends, there exists an action that agent i 's

²¹Please see Section 3 of the supplementary appendix for an example.

²²For example, see [Smith and Sørensen \(2000\)](#). In the following discussion, the informational cascade model is sometimes denoted as the sequential social learning model.

strategy assigns to all the information sets considered possible by his neighbors. Conditions (i) and (iii) are established as Lemmas 1 and 2, while condition (ii) is established in the proof of the theorem.²³ Please see the Appendix for the formal statement of the lemmas and the proof of Theorem 1 and Proposition 1.

Geanakoplos and Polemarchakis (1982) consider the probability announcement setting for the special case of a complete network consisting of two agents. In a complete network, the history of announcements is common knowledge and the information set of an agent in a given period equals the intersection of the cell of his partition and the set of states that are commonly known to contain the true state of the world. Therefore, in a complete network, Lemmas 1 and 2 hold trivially. In an incomplete network, the history of announcements is not common knowledge among all agents, so information sets have to be generated in a different manner. As is shown in Proposition 1 of the supplementary appendix, the information set of an agent can be generated from inferences on his neighbors' actions based only on common observables among pairs of neighbors. Once the properties of Lemmas 1 and 2 are established for the learning process, the proof of Theorem 1 follows directly from Proposition 1. Geanakoplos and Polemarchakis implicitly use a version of Proposition 1 applied to choice functions in their proof.²⁴

5. ASYMPTOTIC LOCAL INDIFFERENCE

The previous section provided a local indifference result once learning ends. In the case of finite partitions, learning ends in finite time for all states. In the case of infinite partitions, however, learning does not necessarily end in finite time. In this section, I am going to establish an asymptotic local indifference result that covers the case of infinite partitions. The concept of a dominant set, defined below, plays a crucial role in the asymptotic analysis.

DEFINITION 2. The set $B \in \mathcal{F}$ is a *dominant set* under c if, for all sequences $\{B^t\}_{t=1}^\infty$ in $\mathcal{F}$ such that $B^{t+1} \subset B^t$ and $\bigcap_{t=1}^\infty B^t = B$, the following statement holds: If there exists an infinite subsequence $\{B^{t_k}\}_{k=1}^\infty$ of $\{B^t\}_{t=1}^\infty$ and an action $a \in A$ such that $a \in c(B^{t_k})$ for all $k \in \mathbb{N}$, then $a \in c(B)$.

Set B is dominant under choice correspondence c if, for every sequence of sets that converge to B , every action that is optimal for infinitely many information sets in the sequence is also optimal for the limit set B . In the expected utility setting, every set with positive probability is a dominant set if the utility function is bounded and measurable for every action a .²⁵

²³There exists a more direct way to prove Theorem 1 that does not rely on the characterization of the learning process. The idea is to extend the state space to capture the histories of play. I decided against this approach as no intuition can be gained with regard to how rational learning works and how it relates to the local indifference result.

²⁴I use the term implicitly, as they do not refer to union consistency per se given that their analysis lies in the probability announcement setting where the union consistency property is satisfied.

²⁵Please see Proposition 2 in the supplementary appendix.

In addition to the concept of a dominant set, I need to impose a dominance consistency condition on the choice correspondence c .

DEFINITION 3. The choice correspondence $c(\cdot)$ is *dominance consistent (DC)*, if for all $B, F \in \mathcal{F}$ such that $B \subset F$,

$$B \text{ dominant set} \Rightarrow F \text{ dominant set.}$$

I have stated that in the expected utility setting with a bounded and measurable utility function, any set of positive probability is a dominant set. Thus the dominance consistency conditions holds in this environment as well, as any set that contains a set of positive probability has positive probability.

Before stating the theorem, let me define the limit information set as

$$I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega)) = \bigcap_{t=1}^{\infty} I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)).$$

The limit information set $I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))$ consists of the states of the world that are contained in the information set $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))$ in every period t . Remember that $\mathcal{I}_i^t(h_{ij}^t(\omega); \omega)$ denotes the set of possible information sets of agent i in period t based on the common observables of i and his neighbor j ; $\{\mathcal{I}_i^t(h_{ij}^t(\omega); \omega)\}_{t \in \mathbb{N}}$ denotes a sequence of collections of sets. Let the limit of the sequence be denoted as

$$\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega) = \{I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) : I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')) \in \mathcal{I}_i^t(h_{ij}^t(\omega); \omega) \forall t\}.$$

Note that the limit set $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ is not empty, as the true information set of player i , $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))$, is contained in $\mathcal{I}_i^t(h_{ij}^t(\omega); \omega)$ in every period t , and the true cell of the join is contained in $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))$.

THEOREM 2. *If all elements of the join $\bigvee_{j \in V} \mathcal{P}_j$ are dominant sets under c , and the choice correspondence is union consistent and dominance consistent, then for every state $\omega \in \Omega$, any action that agent i selects infinitely often is optimal for all his neighbors j in their limit information set. Define*

$$A_i^\infty(\omega) = \{a \in A : a = s_i(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))) \text{ for infinite periods } t\}$$

and $j \in N_i$. Then

$$A_i^\infty(\omega) \subset c(I_j^\infty(\mathcal{P}_j(\omega), h_j^\infty(\omega))).$$

Theorem 2 presents conditions under which an asymptotic local indifference result holds; any action an agent selects infinitely often is optimal for his neighbors in their limit information set. Note that a sufficient condition for the asymptotic local indifference result is that the set of actions is finite as in that case $A_i^\infty(\omega)$ is nonempty for all states ω . If the set of actions is infinite, $A_i^\infty(\omega)$ can be empty. Applied to the expected utility setting, the theorem states that if all cells of the join of partitions have positive

probability, then any action an agent selects infinitely often is optimal for all his neighbors in the limit.

Let me give a brief outline of the proof. I establish that every limit information set $I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))$ is a dominant set, relying on the assumption that every cell of the join is a dominant set together with dominance consistency. The limit set being dominant implies that every action in $A_i^\infty(\omega)$ is optimal for agent i in his limit information set. The strategy of agent i assigns, in every period t , the same action to any information set $I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega'))$ that converges to a limit information set in $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ as to the true information set (Lemma 4 in the Appendix):

$$s_i(I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega'))) = s_i(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))).$$

This, together with all limit information sets being dominant, implies that all actions in $A_i^\infty(\omega)$ are optimal for all limit information sets in $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$. Furthermore, $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ is a collection of disjoint sets, and $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ and $\mathcal{I}_j^\infty(h_{ij}^\infty(\omega); \omega)$ partition the same set for neighbors i and j (Lemma 3 in the Appendix). All conditions of Proposition 1 on the informational structure and behavior in the limit are satisfied. Proposition 1 gives the asymptotic local indifference result: any action that agent i selects infinitely often is optimal for all his neighbors in their limit information set.

The objective of this section is to establish conditions on the choice rule that yield an asymptotic local indifference result. Dominant sets and the dominant consistent choice rule play the following role in the proof. Theorem 2 requires one to connect the sequence of actions chosen with optimal actions in the limit. The concept of a dominant set provides such a connection between actions that are optimal for infinitely many information sets in the information set process converging to the limit information set and their optimality in the limit information set. For the local indifference result to be true in all states of the world, all elements of the power set of the join have to be dominant sets. The condition of dominance consistency allows one to require only that the elements of the join be dominant sets, as it implies dominance of all possible limit information sets.

6. COMMON KNOWLEDGE OF RATIONALITY VERSUS COMMON KNOWLEDGE OF STRATEGIES

The analysis so far was based on the assumption of common knowledge of strategies: common knowledge of which action any agent selects out of the set of actions assigned by the choice correspondence for every information set in $\mathcal{F}$. In this section, I assume that only rationality of agents is commonly known, i.e., it is common knowledge that agents select actions according to the common choice correspondence. This forms a departure from a Bayesian point of view, where probabilities should be assignable to strategies, and a departure from a game theoretic view, where (mixed) strategies are common knowledge.

As established in Theorems 1 and 2 above, the local indifference result holds independently of the actual network structure under common knowledge of strategies.

When relaxing the assumption of common knowledge of strategies, however, the network structure determines the validity of the local indifference result. In complete networks, local indifference holds under common knowledge of rationality, while it might fail in incomplete networks.

To establish the claim, consider the rational learning process in complete networks under common knowledge of rationality. In period $t = 1$, the smallest event that is common knowledge (CK) among all players is the true cell of the meet of all players,

$$\text{CK}^1(\omega) = \bigcap_{j \in V} \mathcal{P}_j(\omega).$$

The set of commonly known possible first period information sets of agent j in state ω is given by

$$\mathcal{I}_j^1(\omega) = \{P_j \in \mathcal{P}_j : P_j \cap \text{CK}^1(\omega) \neq \emptyset\}.$$

For a given first period choice a_j^1 of player j , it becomes common knowledge among all agents that the information set of agent j is contained in set $\mathcal{D}_j^1(a_j^1; \omega)$, where

$$\mathcal{D}_j^1(a_j^1; \omega) = \{I_j^1 \in \mathcal{I}_j^1(\omega) : a_j^1 \in c(I_j^1)\}.$$

The set $\mathcal{D}_j^1(a_j^1; \omega)$ consists of the partition cells P_j that are commonly considered possible for agent j , $P_j \subset \text{CK}^1(\omega)$, and for which action a_j^1 is optimal. Through the observation of a_j^1 , it becomes common knowledge among all agents that player j 's realized partition cell has the property that a_j^1 is optimal given the realized cell.

In complete networks, the history of actions is common knowledge among all players. Thus at the beginning of period $t = 2$, it is common knowledge among all agents that a state in $\text{CK}^2(h^2; \omega)$ is realized where

$$\text{CK}^2(h^2; \omega) = \bigcap_{j \in V} \bigcup \mathcal{D}_j^1(a_j^1; \omega).$$

At the outset of period t , we have

$$\text{CK}^t(h^t; \omega) = \bigcap_{j \in V} \bigcup \mathcal{D}_j^{t-1}(a_j^{t-1}; h^{t-1}; \omega).$$

The information set of player i at time t is given by

$$I_i^t(\mathcal{P}_i(\omega), h^t) = \mathcal{P}_i(\omega) \cap \text{CK}^t(h^t; \omega).$$

The set of commonly known possible information sets of agent j in period t is given by

$$\mathcal{I}_j^t(h^t; \omega) = \left\{ P_j \cap \text{CK}^t(h^t; \omega) : \begin{array}{l} P_j \in \mathcal{P}_j \\ P_j \cap \text{CK}^t(h^t; \omega) \neq \emptyset \end{array} \right\}$$

and the inference agents make regarding agent j 's information set based on his choice in period t is given by

$$\mathcal{D}_j^t(a_j^t; h^t; \omega) = \{I_j^t \in \mathcal{I}_j^t(h^t; \omega) : a_j^t \in c(I_j^t)\}.$$

The information set of agent i in period $t + 1$, depending on the choices of all other agents in period t , is given by

$$I_i^{t+1}(\mathcal{P}_i(\omega), h^{t+1}) = \mathcal{P}_i(\omega) \cap \text{CK}^{t+1}(h^{t+1}; \omega),$$

where

$$\text{CK}^{t+1}(h^{t+1}; \omega) = \bigcap_{j \in V} \bigcup \mathcal{D}_j^t(a_j^t; h^t; \omega).$$

Thus, in complete networks, the information set of an agent in period t is given by the intersection of his true partition cell with the set $\text{CK}^t(h^t; \omega)$, which is common knowledge among all agents.

Theorem 3 states the local indifference result in complete networks under common knowledge of rationality.

THEOREM 3. *If the network is complete and the choice correspondence is union consistent, then common knowledge of rationality is sufficient for the following result: For a given state ω , if there exists a time t' such that $\text{CK}^t(h^t; \omega) = \text{CK}^{t'}(h^{t'}; \omega)$ for all $t \geq t'$, then*

$$a_j^t \in c(I_i^t(\mathcal{P}_i(\omega), h^t))$$

for all $i, j \in V$ and $t \geq t'$.

In complete networks, common knowledge of strategies is not necessary to achieve the local indifference result. The weaker assumption of common knowledge of rationality is sufficient. Once common learning ends, any action an agent selects is optimal for all others.

The proof is again an application of **Proposition 1**. I demonstrate that the three conditions of **Proposition 1** are satisfied. Thus local indifference follows. But does local indifference hold for incomplete networks as well?

Generally, the local indifference result fails in incomplete networks. The reason for failure lies in the fact that, in general, the set of possible information sets agents assign to their neighbors is not a collection of disjoint sets. Thus union consistency does not give the local indifference result, as it relies on the information sets being a disjoint collection of sets. To provide intuition, take a pair of neighbors i, j and consider the inference that player i makes based on player j 's first period choice. Whenever there is a cell in player j 's partition to which the choice correspondence assigns more than a single action, the collection of sets $\{\bigcup \mathcal{D}_j^1(a_j; \omega; i)\}_{a_j \in A}$ does not constitute a partition of Ω . The set of possible information sets of agent i in period $t = 2$ from the perspective of his neighbor k , who is not neighbor of j , may fail to be a collection of disjoint sets. Please see the **Appendix** for an example where local disagreement occurs on a set of positive probability in the expected utility setting.

In this section, I have shown that the consensus result of the literature on knowledge and consensus cannot be generalized to union consistent choice correspondences, common knowledge of rationality, and incomplete, undirected networks. Rational

learning in incomplete networks can lead to disagreement among pairs of (mutual) neighbors.

It could be argued that the above analysis assumes boundedly rational agents, as the inferences do not account for the different possible strategies agents might follow; neither do the information sets. It further could be argued that, by definition, the state space contains all possible states of the world and thus each state should include a description of the possible types of all players, represented by their strategies. In the supplementary appendix, I address these points by considering extensions of the state space in cases of lack of common knowledge of strategies and lack of common knowledge of the network structure.

7. OPTIMAL INFORMATION AGGREGATION AND DURATION TO CONSENSUS

In this section, I address the question of how the network structure affects (i) the aggregation of the privately held information and (ii) the duration to consensus. To analyze these questions, I restrict attention to a special case of the general framework analyzed so far: the probability announcement setting.

In this setting, agents share a common prior and repeatedly announce their posterior belief of an uncertain event. [Theorem 1](#) implies that under finite partitions, consensus is established in finite time among all agents in a connected network. The consensus beliefs, however, might differ from network to network. But how does the network structure affect the precision of the consensus belief?

Let me first define a measure for the quality of information aggregation of different network structures. By information setting $((\Omega, \mathcal{F}, p), V, \{\mathcal{P}_i\}_{i \in V}, Q)$, I denote the conjunction of the probability space, the set of agents, their finite partitions of the state space, and the uncertain event $Q \in \mathcal{F}$. The goal is to rank different communication networks by their quality of information aggregation for a fixed information setting. As above, I restrict attention to connected, undirected networks.

Let $\mathcal{P}_{iG}^t$ denote i 's partition of the state space in period t and for network G .²⁶ Consider the meet $\bigwedge_{i \in V} \mathcal{P}_{iG}^{t*}$ of the partitions of all agents once learning ends in all states. For a given cell of the meet, all states in the cell have the same consensus belief, as otherwise learning would continue in a connected network. The quality of information aggregation is defined in terms of the meet $\bigwedge_{i \in V} \mathcal{P}_{iG}^{t*}$: communication network G leads to better information aggregation than network G' if and only if $\bigwedge_{i \in V} \mathcal{P}_{iG}^{t*}$ is finer than $\bigwedge_{i \in V} \mathcal{P}_{iG'}^{t*}$.²⁷ A communication network G leads to optimal information aggregation if the meet of partitions once learning ends equals the join of partitions in the first period. This is a natural definition of optimal information aggregation, as it states that each agent incorporates the private information of all agents in every state of the world.

One's intuition might suggest that complete networks should do at least as well as incomplete networks, as all agents directly communicate with each other. This intuition, however, is incorrect. Please see the [Appendix](#) for an example of an information setting

²⁶See the proof of [Theorem 1](#) for background on the partition.

²⁷Note that not all pairs of networks need be comparable in this manner, as their corresponding meets need not have a finer/coarser relation.

where an incomplete network leads to better information aggregation than the complete network. The example compares a star network and a complete network with the same underlying set of agents, state space, common prior, and partitional information. It is shown that the meet of learning ends partitions of the star network is finer than the meet of learning ends partitions in the complete network. The consensus belief in the star network equals the pooled information posterior—the posterior probability of the event conditioning on the realized cell of the join—in all states of the world, while the consensus belief in the complete network differs from the pooled information posterior in some states. The better information aggregation of the star network comes at the cost of a longer duration to consensus. The main intuition for the better information aggregation of the star network is that in complete networks, agents might jointly jump to conclusions too fast due to the fact that all announcements are common knowledge.

The natural question resulting from the example concerns (i) the characteristics of optimal communication structures as a function of the information setting and (ii) the duration to consensus of different network structures. I do not provide a complete answer to this question, but a generic one, where the common probability measure is treated as a random element.

Let the partitions of all agents be finite and consider the partition of the state space

$$\mathcal{P} = \left\{ \left\{ \bigvee_{i \in V} \mathcal{P}_i(\omega) \cap Q \right\}_{\omega \in \Omega}, \left\{ \bigvee_{i \in V} \mathcal{P}_i(\omega) \cap Q^C \right\}_{\omega \in \Omega} \right\},$$

where Q^C denotes the complement of Q . The partition $\mathcal{P}$ further refines the join $\bigvee_{i \in V} \mathcal{P}_i(\omega)$ of partitions of all agents by partitioning each cell of the join into states that lie in event Q and states that do not. Let k denote the cardinality of the partition $\mathcal{P}$, $K = \{1, \dots, k\}$. Note that for any information setting

$$((\Omega, \mathcal{F}, p), V, \{\mathcal{P}_i\}_{i \in V}, Q),$$

fixing the probabilities of the cells of partition $\mathcal{P}$ determines the posterior announcement process and the consensus probabilities. Then the probability simplex $\Delta(K)$, endowed with a uniform probability measure μ , describes the set of all probability measures to be considered.²⁸ Note that the distance between two agents in the network is the length of the shortest path connecting them. The diameter of the network, $d(G)$, is the greatest distance between any two agents in the network.

THEOREM 4. *For almost every $[\mu]$ probability measure $p \in \Delta(K)$, consensus is achieved in round $t^* \leq d(G) + 1$. In every state $\omega \in \Omega$, the consensus belief equals the conditional probability of event Q , conditioning on the realized cell of the join.*

Returning to the question of optimal information aggregation, [Theorem 4](#) states that, generically, all connected networks lead to optimal information aggregation. The only distinction between the network structures is the duration to consensus. The larger

²⁸Please see the [Appendix](#) for a precise definition.

is the diameter of the graph, the longer is the (maximal) duration to consensus. The theorem has two important implications. The first implication concerns strategic announcements and the assumption that all agents announce their true posterior. In the probability announcement setting, the incentive to lie can derive only from the possibility to learn more based on lying. But [Theorem 4](#) states that, generically, every agent reveals his private information in every communication round. Hence there is no reason to lie. The second implication involves a distinction between Bayesian and non-Bayesian learning in networks. [Golub and Jackson \(2012\)](#) address the question of speed of convergence under weighted average updating function where consensus occurs asymptotically. They show that link density does not affect the speed of convergence, while homophily—the tendency of individuals to associate with similar individuals—slows the rate of convergence. [Theorem 4](#) implies that under Bayesian learning, changing homophily while keeping the diameter of the graph constant has no effect on the speed of convergence. The main distinction with regard to speed of convergence is the following: adding links to a network leads to a (weak) increase in speed of convergence under Bayesian learning, while under non-Bayesian learning, adding homophily increasing links decreases the speed of convergence. Therefore, adding links to a Bayesian network leads to (weakly) less heterogeneity, while adding links to a non-Bayesian network might lead to more heterogeneity.

[Golub and Jackson \(2012\)](#) also discuss a direct contagion process that is characterized by travel along shortest paths. One example of such a process is a contagion process where agents have binary types—infected and noninfected—and a node becomes infected as soon as at least one neighbor is infected. In such a contagion process, the duration to consensus is bounded above by the diameter of the graph plus 1 exactly as under Bayesian learning. From the point of view of duration to consensus, Bayesian learning is hence much closer to a contagion process as opposed to weighted-average based non-Bayesian learning.

[Geanakoplos and Polemarchakis \(1982\)](#) prove a similar result for a complete network with two agents. The difference in analysis beyond complete as opposed to arbitrary connected network structures, lies in me treating the probability measure as random or uncertain, which seems more reasonable from a practical perspective, while they treat the event Q as random. A nice feature of treating the probability measure as opposed to the event as random is that the perfect information aggregation result holds for discrete and nonatomic prior probability measures while under random events nonatomic probability measures are necessary for the result. Formally, Geanakoplos and Polemarchakis show that almost surely both agents reveal their partition cell in the first round of announcements and information is aggregated perfectly. [Theorem 4](#) states that all agents almost surely reveal their private information set in each round of interaction and in any connected network structure.

8. DISCUSSION OF RESULTS IN RELATION TO LITERATURE

This paper is closely related to the literature on knowledge and consensus, and the literature on learning in networks, and has some relation to the sequential social learning

literature. The literature on knowledge and consensus analyzes under which conditions pairwise repeated communication among a finite set of individuals leads to consensus. The most important contributions are by [Aumann \(1976\)](#), [Geanakoplos and Polemarchakis \(1982\)](#), [Cave \(1983\)](#), [Bacharach \(1985\)](#), [Parikh and Krasucki \(1990\)](#), and [Krasucki \(1996\)](#). This literature throughout assumes pairwise communication, decision functions, and finite partitions, whereas I consider simultaneous communication in incomplete networks, a decision correspondence, and finite and countably infinite partitions.

The most closely related papers are [Parikh and Krasucki \(1990\)](#) and [Krasucki \(1996\)](#). They consider repeated pairwise communication among a finite number of individuals in countably many rounds. Agents apply message functions that map information sets to messages. Communication occurs according to a protocol that selects a pair of agents—sender and receiver—in each of countable rounds. In a given round, the sender announces his message to the receiving agent. This differs from the setting considered in this paper where, in each round, all agents simultaneously “communicate” with all their neighbors. The protocol defines a directed graph where an edge from agent i to j exists if agent i communicates with j in infinitely many rounds.

The protocol is denoted as fair if the underlying graph is strongly connected. [Parikh and Krasucki \(1990\)](#) show that consensus occurs in finite time under a fair protocol whenever the message function satisfies a convexity property that is a stronger requirement than union consistency. [Krasucki \(1996\)](#) establishes that union consistency is sufficient for consensus under a fair protocol if the graph contains no cycles.

The main similarity of this paper and the existing literature on knowledge and consensus concerns the nature of the learning process and the corollary result to [Theorem 1](#), which establishes sufficiency of a connected, undirected network and union consistent choice functions for consensus in finite time. The sufficiency of union consistent as opposed to convex choice functions in [Parikh and Krasucki \(1990\)](#), and connected as opposed to a strongly connected network with no cycles in [Krasucki \(1996\)](#), comes from my assumption of an undirected underlying graph.

All four theorems presented in this paper are new contributions to the literature on knowledge and consensus. The main extensions and contributions are the following: (i) the generalization to choice correspondences and the resulting local indifference result allowing for heterogeneous choice; (ii) the extension to infinite partitions, and the resulting sufficient conditions on the information setting and choice correspondence that assure asymptotic consensus; (iii) the distinction between common knowledge of rationality and common knowledge of strategies, and the resulting validity of local indifference in complete networks and the possibility of its failure in incomplete networks. This possibility of local disagreement among pairs of (mutual) neighbors once learning ends is the first example of its kind under union consistent choice correspondences.

In the literature on learning in networks, there are two previous papers that analyze Bayesian learning in networks: [Gale and Kariv \(2003\)](#) and [Rosenberg et al. \(2009\)](#). Both consider the expected utility setting. The closer related one is Rosenberg et al., which independently establishes an asymptotic local indifference result in the expected utility setting, allowing for strategic behavior. Therefore, there is some overlap of their results

and my first two theorems.²⁹ In particular, the overlap concerns the expected utility setting with countable signal support. However, their approach differs significantly from the one taken in this paper due to the different framework considered and the techniques used to establish the result. To establish local indifference, they rely on martingale based and measure theoretic arguments. Another distinction lies in the focus of Rosenberg et al. and this paper: their focus is on establishing a general asymptotic local indifference result in the expected utility setting. The focus of this paper is twofold: First, it lies on a characterization of the learning process in a general setting so as to analyze the occurrence of local indifference, consensus, and local disagreement among agents in social networks. Second, the characterization of the learning process is used as a tool to address questions on optimal aggregation of private information in networks and to establish distinguishing features between non-Bayesian and Bayesian learning.

Within the non-Bayesian learning literature, the most related papers are DeMarzo et al. (2003) and Golub and Jackson (2010). Here agents follow simple updating rules when forming their opinion: the opinion of an agent in round t is the weighted average of the last period opinions of his neighbors (and himself). In the thus defined non-Bayesian setting, asymptotic consensus has been shown to occur in strongly connected networks. In this paper, I have shown that the consensus result carries forward to Bayesian learning settings under choice functions. In the non-Bayesian setting, consensus is the simple consequence of the mechanical updating of agents, while in this paper, it is based on potentially highly complex computations agents make in each round of interaction.

The results presented in this paper also help to distinguish Bayesian and non-Bayesian learning from an observational point of view. First, the local indifference result presented in this paper captures heterogeneous action choices once learning ends (and therefore in the limit), which is inconsistent with weighted-average based non-Bayesian learning in connected, undirected graphs if the weights agents assign to their neighbors and themselves are strictly positive.

Second and more importantly, an interesting distinction with regard to the duration to consensus as a response to changes in the network structure is established. Adding homophily increasing links to a non-Bayesian network leads to a decrease in speed of convergence, and thereby more heterogeneity, while adding links to a Bayesian network leads to a (weak) decrease in duration to consensus, and thereby weakly less heterogeneity.³⁰

Finally, this paper is also related to the sequential social learning literature, with the most prominent contributions being Bikhchandani et al. (1992) and Smith and Sørensen (2000). In these papers, each of (countable) infinitely many (homogeneous) agents observes a private signal about the state of the world and makes a one-time, irreversible choice in a predetermined sequence. In a recent paper, Acemoglu et al. (2011) generalize the sequential social learning model with a social network application in mind. In their

²⁹The results established in my paper go beyond the expected utility setting. For an example of a union consistent choice correspondence that is not representable in expected utility form, see Section 3 in the supplementary appendix.

³⁰See Golub and Jackson (2012) for the non-Bayesian case.

model, agents observe only a stochastically generated subset of their predecessors—their neighborhood—as opposed to all predecessors. The focus of this literature lies on a characterization of conditions on the signal and network structure such that herding, or alternatively perfect learning, where the correct action is chosen in the limit, occurs. Due to the difference in the setting considered, i.e., finite versus infinite set of agents and repeated versus one-time choice, the two approaches are complementary. The results can be compared with regard to the occurrence of herding and the information aggregation properties. [Theorem 1](#) implies a herding result for any network in case of a choice function and finite partitions: there exists a finite time period such that from that period onward all agents in the network select the same action. In the sequential social learning model with homogeneous agents, however, the occurrence of herding depends both on the signal structure, bounded versus unbounded private signals and on the structure of the network.³¹ The implication of such herding for information aggregation differs, however. While herding in the sequential social learning setting precludes perfect learning, [Theorem 4](#) establishes that in the repeated interaction framework analyzed here, herding generically involves perfect learning in the sense of the posterior belief announced that aggregates the private information of all agents.

9. CONCLUSION

This paper analyzes rational learning and the evolution of behavior and opinions in social networks where agents repeatedly take actions. The analysis is motivated by the essential role social networks play in many economic and social settings, and the necessity of a thorough understanding of the underlying rational learning dynamics. Despite the immense (practical) complexities involved in rational learning in incomplete networks, I characterize the informational structure in a simple and intuitive way.

My first result establishes that under common knowledge of strategies, rational learning leads to local indifference: any action an agent selects is optimal for all his neighbors once learning ceases. For the case where learning does not end in finite time, I characterize a set of conditions such that an asymptotic local indifference result holds.

However, rational learning can lead to failure of local indifference. Under the absence of common knowledge of strategies, local indifference might fail in incomplete networks, while it holds in complete networks.

The paper also provides a result on the effect of the network structure on the degree of information aggregation and speed of convergence in the probability announcement setting. Generically, every connected network aggregates the private information of all agents and the duration to consensus increases with the diameter of the graph. This is, however, not generally true. I provide an example of an information setting where the complete network aggregates less information than an incomplete network.

The analysis and the results established in this paper are based on the assumption of the social network being undirected. A natural question to ask then is whether the results hold for more general directed networks. The answer to this question is yes. It is

³¹See Theorems 2 and 4 of [Acemoglu et al. \(2011\)](#).

straightforward to establish the local indifference result of Theorems 1 and 2 for agents who are mutual neighbors of each other, as they then share a common history. Since Theorem 3 requires complete networks, it trivially holds for directed graphs as well.

The assumption of common knowledge of the network structure is made throughout the paper. However, the local indifference result is robust to relaxing assumptions on common knowledge of the network structure and common knowledge of strategies if learning occurs in an extended state space and decisions are made based on information sets within that state space. Please see the supplementary appendix for a formal treatment.

The interested reader might also wonder about the dependence of the local indifference results on the assumption of the network being fixed over time. What if the network structure is allowed to change over time? If the process of networks is common knowledge, then the properties of Lemmas 1 and 3 and Lemmas 2 and 4 are not affected. Therefore, once learning ends, local indifference holds for any pair of neighbors in every period. Asymptotically, local indifference, as stated in Theorem 2, holds for any pair of neighbors who observe each other in infinitely many periods.

APPENDIX A: EXAMPLES

EXAMPLE 1 (Failure of local indifference under common knowledge of rationality). Consider a star network with three agents 1, 2, and 3, $N_1 = \{2\}$, $N_2 = \{1, 3\}$, and $N_3 = \{2\}$. The state space consists of nine states $\Omega = \bigcup_{i=1}^9 \omega_i$. Let the state space be represented by the rectangle in Figure 2.

	c^1	c^2	c^3
r^1	ω_1	ω_2	ω_3
r^2	ω_4	ω_5	ω_6
r^3	ω_7	ω_8	ω_9

FIGURE 2. State space and partitions.

Agents 1 and 3 have private information, while agent 2 has not. Agent 1 learns the rows of the matrix, while agent 3 learns the columns. The partitions are given by

$$\mathcal{P}_1 = \{\omega_1, \omega_2, \omega_3; \omega_4, \omega_5, \omega_6; \omega_7, \omega_8, \omega_9\}$$

$$\mathcal{P}_2 = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6, \omega_7, \omega_8, \omega_9\}$$

$$\mathcal{P}_3 = \{\omega_1, \omega_4, \omega_7; \omega_2, \omega_5, \omega_8; \omega_3, \omega_6, \omega_9\}.$$

The example lies within the expected utility setting. The agents share a common prior p over Ω with $p(\omega_2) = p(\omega_4) = p(\omega_6) = p(\omega_8) = p(\omega_9) = \frac{17}{140}$, $p(\omega_5) = \frac{34}{140}$, and $p(\omega_1) = p(\omega_3) = p(\omega_7) = \frac{7}{140}$. The set of actions consists of odd or even $\mathcal{A} = \{o, e\}$, where action o yields a utility of 1 in the odd indexed states and zero otherwise, while action even yields a utility of 1 in even indexed states and zero otherwise:

$$u(o, \omega_i) = \begin{cases} 1 & \text{if } i = 1, 3, 5, 7, 9 \\ 0 & \text{otherwise} \end{cases} \quad u(e, \omega_i) = \begin{cases} 1 & \text{if } i = 2, 4, 6, 8 \\ 0 & \text{otherwise.} \end{cases}$$

In each round, agents maximize their expected utility conditioning on their information. The choice correspondence c then prescribes the following choices for the possible first period information sets

$$c(\omega_1, \omega_2, \omega_3) = c(\omega_1, \omega_4, \omega_7) = e$$

$$c(\omega_4, \omega_5, \omega_6) = c(\omega_2, \omega_5, \omega_8) = \{o, e\}$$

$$c(\Omega) = c(\omega_3, \omega_6, \omega_9) = c(\omega_7, \omega_8, \omega_9) = o$$

$$c(\omega_1, \omega_2, \omega_4, \omega_5) = c(\omega_2, \omega_3, \omega_5, \omega_6) = c(\omega_5, \omega_6, \omega_8, \omega_9) = c(\omega_4, \omega_5, \omega_7, \omega_8) = o.$$

Let the state ω_3 be realized, which occurs with probability $\frac{7}{140} > 0$. Player 1 then observes the first row and selects e , player 2 selects o , and player 3 observes the third column and selects o . The information sets at the beginning of the second period for players are $I_1^2(\mathcal{P}_1(\omega_3), h_1^2) = \{\omega_1, \omega_2, \omega_3\}$, $I_2^2(\mathcal{P}_2(\omega_3), h_2^2) = \{\omega_2, \omega_3, \omega_5, \omega_6\}$, and $I_3^2(\mathcal{P}_3(\omega_3), h_3^2) = \{\omega_3, \omega_6, \omega_9\}$. The second stage information sets that are common knowledge to be possible for agent i among agents i and j based on the common observables of agents i and j are

$$\mathcal{I}_2^2(h_{12}^2; \omega_3) = \{\omega_1, \omega_2, \omega_4, \omega_5; \omega_2, \omega_3, \omega_5, \omega_6\}$$

$$\mathcal{I}_1^2(h_{12}^2; \omega_3) = \{\omega_1, \omega_2, \omega_3; \omega_4, \omega_5, \omega_6\}$$

$$\mathcal{I}_2^2(h_{23}^2; \omega_3) = \{\omega_2, \omega_3, \omega_5, \omega_6; \omega_5, \omega_6, \omega_8, \omega_9\}$$

$$\mathcal{I}_3^2(h_{23}^2; \omega_3) = \{\omega_2, \omega_5, \omega_8; \omega_3, \omega_6, \omega_9\}.$$

Note that for all information sets of agent 2 considered possible by agents 1 and 3, the choice correspondence prescribes action o . Therefore neither agent 1 nor agent 2 can make any further inference from agent 2's second period choice. At the beginning of the third round (and in every later round as well), agent $i = 1, 3$ has to consider the possibility of an action switch of agent $j \neq i$, $j = 1, 3$, which would perfectly reveal the cell of his partition. Suppose that h^t is such that all agents played the same action in every round as in the first. Then the sets of possible information sets of agent 2 from the perspective of agent 1 are

$$\mathcal{I}_2^t(h_{12}^t; \omega_3) = \{\omega_1, \omega_2, \omega_4, \omega_5; \omega_2, \omega_3, \omega_5, \omega_6; \omega_2, \omega_5\}$$

$$\mathcal{I}_2^t(h_{23}^t; \omega_3) = \{\omega_2, \omega_3, \omega_5, \omega_6; \omega_5, \omega_6, \omega_8, \omega_9; \omega_5, \omega_6\}.$$

But the expected utility maximizing action conditioning on $\{\omega_2, \omega_5\}$ and $\{\omega_5, \omega_6\}$ is action o . Agent 2 optimally selects o for all his possible information sets. In this example, learning ends and the only action chosen by and optimal for agent 1 is e , while for his neighbor, agent 2, the only action chosen by and optimal for is action o . Hence local indifference can fail with positive probability in the expected utility framework if the network is incomplete and only common knowledge of rationality is assumed. $\diamond$



FIGURE 3. The state space Ω and information partitions.

FIGURE 4. Join of partitions.

EXAMPLE 2 (Superiority of information aggregation in incomplete network versus complete network). Consider the example in Figure 3 that compares the quality of information aggregation of a complete three agent network with a three agent star network. Suppose that Ω , the set of all states of the world, is represented by the rectangle in Figure 1.³² All agents share a uniform common prior probability distribution over the rectangle.

Agent 1 partitions Ω horizontally, the cells of his partitions are given by the union of the first two rows, r^1 , the union of rows three and four, r^2 , and the remaining two rows, r^3 , $\mathcal{P}_1 = \{r^1, r^2, r^3\}$. The diagonal line in the rectangle represents the dividing line of agent 2's partition. He partitions the state space diagonally: the cells of his partition are the upper triangle, d^1 , and the lower triangle d^2 , $\mathcal{P}_2 = \{d^1, d^2\}$. Agent 3 partitions the state space vertically. The cells of his partition are given by the union of the first two columns, c^1 , the union of columns three and four, c^2 , and the union of the remaining two columns, c^3 , $\mathcal{P}_3 = \{c^1, c^2, c^3\}$. The event Q with regard to which agents announce their posterior probabilities is given by $Q = \bigcup_{i=1}^{10} Q_i$. From the partitions of the three agents, one can generate their join. Figure 4 represents the cells of the join of all agents and labels them from P_1 to P_{12} .

Consider the complete network, G_1 , first. Applying the learning process defined in Section 3 establishes that learning ends with the second period announcements in all states of the world. The meet of partitions once learning ends equals

$$\bigwedge_{i \in V} \mathcal{P}_{iG_1}^3 = \{P_1 P_2 P_5 P_6, P_3, P_4 P_8, P_9, P_{10}, P_{11}, P_{12}\}.$$

³²There is a continuum of states.

The incomplete network, G_2 , has a star structure with agent 2 as the center node, $N_1 = N_3 = \{2\}$ and $N_2 = \{1, 3\}$. Learning ends with the announcements in period three and the meet of partitions once learning ends equals

$$\bigwedge_{i \in V} \mathcal{P}_{iG_2}^4 = \{P_1, P_2, P_5, P_6, P_3, P_4P_8, P_9, P_{10}, P_{11}, P_{12}\}.$$

For states $\omega \in \bigcup_{i=1,2,5,6} P_i$, the incomplete network leads to more precise information aggregation; in fact, the incomplete network yields common knowledge of the realized cell of the join for

$$\omega \in \bigcup_{i=1,2,5,6} P_i$$

and different consensus posteriors than the complete network, $q_{G_1}^*(\omega) = \frac{1}{2}$ for all

$$\omega \in \bigcup_{i=1,2,5,6} P_i,$$

while

$$q_{G_2}^*(\omega) = \begin{cases} \frac{1}{4} & \text{if } \omega \in P_1 \\ \frac{3}{4} & \text{if } \omega \in P_2 \cup P_5 \\ 0 & \text{if } \omega \in P_6. \end{cases}$$

In all other states of the world, the two networks are equivalent in the quality of information aggregation. The incomplete network G_2 leads to better information aggregation than G_1 as $\bigwedge_{i \in V} \mathcal{P}_{iG_2}^{t*}$ is finer than $\bigwedge_{i \in V} \mathcal{P}_{iG_1}^{t*}$.³³ This example shows that for some information settings, incomplete networks yield better information aggregation than complete networks by restricting communication, even though the duration to consensus is (weakly) longer. $\diamond$

APPENDIX B: PROOFS

PROOF OF PROPOSITION 1. Suppose $c(\cdot)$ satisfies union consistency and there exist actions $a_l, a_g \in A$ such that

$$a_l \in c(L) \quad \text{for all } L \in \mathcal{L}$$

$$a_g \in c(G) \quad \text{for all } G \in \mathcal{G}.$$

By union consistency, we have

$$\bigcap_{G \in \mathcal{G}} c(G) = c\left(\bigcup_{G \in \mathcal{G}} G\right)$$

³³Please see Section 5 in the supplementary appendix for background on the learning process in this example.

and

$$\bigcap_{L \in \mathcal{L}} c(L) = c\left(\bigcup_{L \in \mathcal{L}} L\right).$$

As the union over all sets in $\mathcal{G}$ equals the union over all sets in $\mathcal{L}$, the choice correspondence assigns the same set of actions

$$c\left(\bigcup_{G \in \mathcal{G}} G\right) = c\left(\bigcup_{L \in \mathcal{L}} L\right),$$

implying by union consistency

$$\bigcap_{G \in \mathcal{G}} c(G) = \bigcap_{L \in \mathcal{L}} c(L)$$

and thus

$$\begin{aligned} a_l \in c(L) \forall L \in \mathcal{L} &\Rightarrow a_l \in c(G) \forall G \in \mathcal{G} \\ a_g \in c(G) \forall G \in \mathcal{G} &\Rightarrow a_g \in c(L) \forall L \in \mathcal{L}. \end{aligned}$$

□

LEMMA 1. *For all agents $i \in V$ and all periods t ,*

$$(\mathcal{P}_i(\omega), h_i^t(\omega)) \neq (\mathcal{P}_i(\omega'), \hat{h}_i^t(\omega'))$$

implies

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) \cap I_i^t(\mathcal{P}_i(\omega'), \hat{h}_i^t(\omega')) = \emptyset.$$

PROOF. Let me first consider the case where $\mathcal{P}_i(\omega) = \mathcal{P}_i(\omega')$ and $h_i^t(\omega) \neq \hat{h}_i^t(\omega')$. I use a proof by induction. First I show that $h_i^2(\omega) \neq \hat{h}_i^2(\omega')$ implies

$$I_i^2(\mathcal{P}_i(\omega), h_i^2(\omega)) \cap I_i^2(\mathcal{P}_i(\omega'), \hat{h}_i^2(\omega')) = \emptyset.$$

By definition of $I_i^2(\mathcal{P}_i(\omega), h_i^2(\omega))$,

$$I_i^2(\mathcal{P}_i(\omega), h_i^2(\omega)) = \mathcal{P}_i(\omega) \cap \bigcap_{l \in N_i} \bigcup \mathcal{D}_l^1(a_l^1; \omega; i).$$

The equality $\mathcal{P}_i(\omega) = \mathcal{P}_i(\omega')$ implies $\mathcal{I}_l^1(\omega; i) = \mathcal{I}_l^1(\omega'; i)$. Thus³⁴ $\{\bigcup \mathcal{D}_l^1(a; \omega; i)\}_{a \in A} = \{\bigcup \mathcal{D}_l^1(a; \omega'; i)\}_{a \in A}$ form identical partitions of $\mathcal{I}_l^1(\omega; i)$, which implies for $a_l^1 \neq \hat{a}_l^1$,

$$\left(\bigcup \mathcal{D}_l^1(a_l^1; \omega; i)\right) \cap \left(\bigcup \mathcal{D}_l^1(\hat{a}_l^1; \omega'; i)\right) = \emptyset.$$

The inequality $h_i^2(\omega) \neq \hat{h}_i^2(\omega')$ implies that there exists an agent $l \in N_i$ such that $a_l^1 \neq \hat{a}_l^1$, which implies

$$I_i^2(\mathcal{P}_i(\omega), h_i^2(\omega)) \cap I_i^2(\mathcal{P}_i(\omega'), \hat{h}_i^2(\omega')) = \emptyset.$$

³⁴Where a is such that $\mathcal{D}_l^1(a; \omega; i) \neq \emptyset$.

Suppose now that for all $h_i^{t-1}(\omega) \neq \hat{h}_i^{t-1}(\omega')$, we have

$$I_i^{t-1}(\mathcal{P}_i(\omega), h_i^{t-1}(\omega)) \cap I_i^{t-1}(\mathcal{P}_i(\omega'), \hat{h}_i^{t-1}(\omega')) = \emptyset.$$

Take any $h_i^t(\omega) \neq \hat{h}_i^t(\omega')$. There are two cases to consider. The first case is where the history of actions in the previous period is unequal, $h_i^{t-1}(\omega) \neq \hat{h}_i^{t-1}(\omega')$. By the induction hypothesis, we have

$$I_i^{t-1}(\mathcal{P}_i(\omega), h_i^{t-1}(\omega)) \cap I_i^{t-1}(\mathcal{P}_i(\omega'), \hat{h}_i^{t-1}(\omega')) = \emptyset.$$

This, together with the fact that the information set of player i in period t is a subset of the information set in period $t-1$, i.e.,

$$\begin{aligned} I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) &\subset I_i^{t-1}(\mathcal{P}_i(\omega), h_i^{t-1}(\omega)) \\ I_i^t(\mathcal{P}_i(\omega'), \hat{h}_i^t(\omega')) &\subset I_i^{t-1}(\mathcal{P}_i(\omega'), \hat{h}_i^{t-1}(\omega')), \end{aligned}$$

yields the desired result

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) \cap I_i^t(\mathcal{P}_i(\omega'), \hat{h}_i^t(\omega')) = \emptyset.$$

Now consider the case where $h_i^t(\omega) \neq \hat{h}_i^t(\omega')$ and $h_i^{t-1}(\omega) = \hat{h}_i^{t-1}(\omega')$. This implies that there are some agents l in i 's neighborhood who select a different action in period $t-1$ under $h_i^t(\omega)$ than under $\hat{h}_i^t(\omega')$, i.e., $a_l^{t-1} \neq \hat{a}_l^{t-1}$ for some $l \in N_i$. The equality of histories until time $t-1$, $h_i^{t-1}(\omega) = \hat{h}_i^{t-1}(\omega')$ implies

$$\mathcal{I}_l^{t-1}(h_{il}^{t-1}(\omega); \omega) = \mathcal{I}_l^{t-1}(\hat{h}_{il}^{t-1}(\omega'); \omega')$$

for all $l \in N_i$. Thus the collection of sets $\{\bigcup \mathcal{D}_l^{t-1}(a; h_{il}^{t-1}(\omega); \omega)\}_{a \in A}$, $\{\bigcup \mathcal{D}_l^{t-1}(a; \hat{h}_{il}^{t-1}(\omega'); \omega')\}_{a \in A}$ constitutes identical partitions of $\bigcup \mathcal{I}_l^{t-1}(h_{il}^{t-1}(\omega); \omega)$. For any $a_l^{t-1} \neq \hat{a}_l^{t-1}$, we have

$$\left(\bigcup \mathcal{D}_l^{t-1}(a_l^{t-1}; h_{il}^{t-1}(\omega); \omega)\right) \cap \left(\bigcup \mathcal{D}_l^{t-1}(\hat{a}_l^{t-1}; \hat{h}_{il}^{t-1}(\omega'); \omega')\right) = \emptyset.$$

As the information set $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))$ is given by

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) = \mathcal{P}_i(\omega) \cap \bigcap_{l \in N_i} \bigcup \mathcal{D}_l^{t-1}(a_l^{t-1}; h_{il}^{t-1}(\omega); \omega)$$

and there are some $l \in N_i$ with $a_l^{t-1} \neq \hat{a}_l^{t-1}$, we have

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) \cap I_i^t(\mathcal{P}_i(\omega'), \hat{h}_i^t(\omega')) = \emptyset.$$

Suppose now that $\mathcal{P}_i(\omega) \neq \mathcal{P}_i(\omega')$. By definition of $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))$, we have for all t and ω ,

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) \subset \mathcal{P}_i(\omega),$$

thus implying

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) \cap I_i^t(\mathcal{P}_i(\omega'), \hat{h}_i^t(\omega')) = \emptyset. \quad \square$$

LEMMA 2. For a pair of neighbors $i \in V, j \in N_i$, if

$$\bigcup \mathcal{I}_i^{t'+1}(h_{ij}^{t'+1}(\omega); \omega) = \bigcup \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega),$$

then there exists $a_i^{t'} \in A$ such that

$$a_i^{t'} = s_i(I_i^{t'})$$

for all $I_i^{t'} \in \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega)$.

PROOF. Take any pair of information sets $I_i^{t'}, \hat{I}_i^{t'} \in \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega)$. Suppose

$$s_i(I_i^{t'}) \neq s_i(\hat{I}_i^{t'}).$$

If $I_i^{t'}$ is the true information set of player i , he selects the action $a_i^{t'} = s_i(I_i^{t'})$, leading to $\mathcal{D}_i^{t'}(a_i^{t'}; h_{ij}^{t'}(\omega); \omega)$ with

$$\hat{I}_i^{t'} \notin \mathcal{D}_i^{t'}(a_i^{t'}; h_{ij}^{t'}(\omega); \omega).$$

As $\mathcal{I}_i^{t'+1}(h_{ij}^{t'+1}(\omega); \omega)$ contains only elements $I_i^{t'+1}$ such that there exists an $I_i^{t'} \in \mathcal{D}_i^{t'}(a_i^{t'}; h_{ij}^{t'}(\omega); \omega)$ with $I_i^{t'+1} \subset I_i^{t'}$ and, by Lemma 1, $\mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega)$ is a disjoint collection of sets, we have

$$\bigcup \mathcal{I}_i^{t'+1}(h_{ij}^{t'+1}(\omega); \omega) \cap \hat{I}_i^{t'} = \emptyset,$$

which yields a contradiction to

$$\bigcup \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega) = \bigcup \mathcal{I}_i^{t'+1}(h_{ij}^{t'+1}(\omega); \omega). \quad \square$$

PROOF OF THEOREM 1. By Lemma 2, for all $i \in V$ and their neighbors j , there exists $a_i^{t'}$ such that

$$a_i^{t'} = s_i(I_i^{t'})$$

for all $I_i^{t'} \in \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega)$. Thus $a_i^{t'} \in c(I_i^{t'})$ for all $I_i^{t'} \in \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega)$. By Lemma 1, the set of information sets $\{I_i^t(\omega, h_i^t(\omega))\}_{\omega \in \Omega}$ is a disjoint collection of sets in each period t and for every agent i , implying that $\mathcal{I}_i^t(h_i^t(\omega); \omega)$ is a collection of disjoint sets. Furthermore, the true state ω is contained in $I_i^t(\omega, h_i^t(\omega))$. Therefore, $\{I_i^t(\omega, h_i^t(\omega))\}_{\omega \in \Omega}$ partitions the state space Ω . Denote this partition as $\mathcal{P}_i^t$. Hence for two neighbors i, j , the union over the set of possible information sets $\bigcup \mathcal{I}_i^t(h_{ij}^t(\omega); \omega)$ equals the cell of the meet of partitions in period t ,

$$\bigcup \mathcal{I}_i^t(h_{ij}^t(\omega); \omega) = \bigcap_{l=i,j} \mathcal{P}_l^t(\omega),$$

which implies

$$\bigcup \mathcal{I}_i^t(h_{ij}^t(\omega); \omega) = \bigcup \mathcal{I}_j^t(h_{ij}^t(\omega); \omega).$$

The conditions of Proposition 1 are satisfied and thus for all pairs of neighbors i, j ,

$$s_j(I_j^{t'}(\mathcal{P}_j(\omega), h_j^{t'}(\omega))) = a_j^{t'} \in c(I_i^{t'}(\mathcal{P}_i(\omega), h_i^{t'}(\omega))).$$

In addition, the condition

$$\bigcup \mathcal{I}_i^t(h_{ij}^t(\omega); \omega) = \bigcup \mathcal{I}_i^{t'}(h_{ij}^{t'}(\omega); \omega)$$

for all $t \geq t'$ implies

$$I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) = I_i^{t'}(\mathcal{P}_i(\omega), h_i^{t'}(\omega)).$$

Thus local indifference holds for all periods $t \geq t'$. $\square$

LEMMA 3. *The inequality $(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \neq (\mathcal{P}_i(\omega''), h_i^\infty(\omega''))$ implies*

$$I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \cap I_i^\infty(\mathcal{P}_i(\omega''), h_i^\infty(\omega'')) = \emptyset.$$

PROOF. Take any two distinct limit information sets $I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega'))$ and $I_i^\infty(\mathcal{P}_i(\omega''), h_i^\infty(\omega''))$. For $\mathcal{P}_i(\omega') \neq \mathcal{P}_i(\omega'')$, Lemma 1 implies for all periods t that

$$I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')) \cap I_i^t(\mathcal{P}_i(\omega''), h_i^t(\omega'')) = \emptyset.$$

The definition of the limit set implies

$$I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \subset I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega'))$$

for all t and together we have

$$I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \cap I_i^\infty(\mathcal{P}_i(\omega''), h_i^\infty(\omega'')) = \emptyset.$$

Suppose $\mathcal{P}_i(\omega') = \mathcal{P}_i(\omega'')$ and $h_i^\infty(\omega') \neq h_i^\infty(\omega'')$. The inequality of histories then implies that there exists a $\hat{t}$ such that $h_i^{\hat{t}}(\omega') \neq h_i^{\hat{t}}(\omega'')$. Lemma 1 then yields

$$I_i^{\hat{t}}(\mathcal{P}_i(\omega'), h_i^{\hat{t}}(\omega')) \cap I_i^{\hat{t}}(\mathcal{P}_i(\omega''), h_i^{\hat{t}}(\omega'')) = \emptyset.$$

By definition of the limit set for each period t and each state ω' , we have

$$I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \subset I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')),$$

which implies

$$I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \cap I_i^\infty(\mathcal{P}_i(\omega''), h_i^\infty(\omega'')) = \emptyset. \quad \square$$

LEMMA 4. *If $I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \in \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$, then*

$$s_i(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))) = s_i(I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')))$$

for each period t .

PROOF. Consider a sequence of information sets $\{I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega'))\}_{t \in \mathbb{N}}$ that converges to a limit set $I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega'))$ in $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$. Suppose that for some period t , we have

$$\hat{a} = s_i(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))) \neq s_i(I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')))$$

which implies by [Lemma 1](#) and the definition of $\mathcal{D}_i^t(\cdot)$ that

$$I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')) \cap \left(\bigcup \mathcal{D}_i^t(\hat{a}; h_i^t(\omega); \omega) \right) = \emptyset.$$

By definition, $\bigcup \mathcal{I}_i^{t+1}(h_{ij}^{t+1}(\omega); \omega)$ is a subset of $\bigcup \mathcal{D}_i^t(a_i^t; h_i^t(\omega); \omega)$, implying

$$I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega')) \cap \left(\bigcup \mathcal{I}_i^{t+1}(h_{ij}^{t+1}(\omega); \omega) \right) = \emptyset.$$

As $I_i^{t+1}(\mathcal{P}_i(\omega'), h_i^{t+1}(\omega')) \subset I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega'))$, we have

$$I_i^{t+1}(\mathcal{P}_i(\omega'), h_i^{t+1}(\omega')) \cap \left(\bigcup \mathcal{I}_i^{t+1}(h_{ij}^{t+1}(\omega); \omega) \right) = \emptyset,$$

yielding a contradiction with $I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \in \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$. $\square$

PROOF OF THEOREM 2. Define the set

$$A_i^\infty(\omega) = \{a \in A : a = s_i(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))) \text{ for infinite periods } t\}.$$

The set $A_i^\infty(\omega)$ consists of all actions that agent i selects infinitely often in state ω according to his strategy. First I establish that the set $A_i^\infty(\omega)$ is contained in the set of optimal actions for every $I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega')) \in \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$, where $j \in N_i$. As each information set $I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))$ contains at least one cell of the join, which is dominant by assumption, dominance consistency implies that $I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))$ is a dominant set. By definition of the information set in period t , we have $I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)) \subset I_i^{t-1}(\mathcal{P}_i(\omega), h_i^{t-1}(\omega))$. By definition of the limit information set $I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))$,

$$I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega)) = \bigcap_{t=1}^{\infty} I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega)).$$

Therefore, the sequence $\{I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))\}_{t=1}^{\infty}$ satisfies the conditions of [Definition 2](#). Since the actions in $A_i^\infty(\omega)$ are optimal in infinitely many periods and the limit information set $I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))$ is a dominant set under c , we have

$$A_i^\infty(\omega) \subset c(I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))).$$

For any information set $I_i^\infty(\mathcal{P}_i(\omega'), h_i^\infty(\omega'))$ that is an element of $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ and every period t , we have, by [Lemma 4](#),

$$s_i(I_i^t(\mathcal{P}_i(\omega), h_i^t(\omega))) = s_i(I_i^t(\mathcal{P}_i(\omega'), h_i^t(\omega'))),$$

which implies that all actions in $A_i^\infty(\omega)$ are optimal for every limit information set in $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ due to all limit information sets being dominant sets,

$$A_i^\infty(\omega) \subset c(I_i^\infty)$$

for all $I_i^\infty \in \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$. Applying the same reasoning for agent $j \in N_i$ yields

$$A_j^\infty(\omega) \subset c(I_j^\infty)$$

for all $I_j^\infty \in \mathcal{I}_j^\infty(h_{ij}^\infty(\omega); \omega)$. By [Lemma 3](#), the set of information sets, $\{I_i^\infty(\omega, h_i^\infty(\omega))\}_{\omega \in \Omega}$ is a disjoint collection of sets, implying that $\mathcal{I}_i^t(h_{ij}^t(\omega); \omega)$ is a collection of disjoint sets. Furthermore, the true state ω is contained in $I_i^\infty(\omega, h_i^\infty(\omega))$. Therefore, $\{I_i^\infty(\omega, h_i^\infty(\omega))\}_{\omega \in \Omega} = \mathcal{P}_i^\infty$ partitions the state space Ω . Hence for two neighbors i, j , the union over the set of possible information sets $\bigcup \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ equals the cell of the meet of the limit partitions of i and j ,

$$\bigcup \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega) = \bigwedge_{l=i,j} \mathcal{P}_l^\infty(\omega),$$

which implies

$$\bigcup \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega) = \bigcup \mathcal{I}_j^\infty(h_{ij}^\infty(\omega); \omega).$$

Thus the conditions of [Proposition 1](#) are satisfied for the collection of sets $\mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ and $\mathcal{I}_j^\infty(h_{ij}^\infty(\omega); \omega)$, which implies

$$A_j^\infty(\omega) \subset c(I_i^\infty)$$

for all $I_i^\infty \in \mathcal{I}_i^\infty(h_{ij}^\infty(\omega); \omega)$ and thus for the true limit information set

$$A_j^\infty(\omega) \subset c(I_i^\infty(\mathcal{P}_i(\omega), h_i^\infty(\omega))),$$

concluding the proof. $\square$

PROOF OF THEOREM 3. I have to establish the three conditions of [Proposition 1](#) so as to prove the global indifference result. The set of possible information sets of agent i in period t' from the perspective of all other agents is equal to

$$\mathcal{I}_i^{t'}(h^{t'}; \omega) = \left\{ P_i \cap \text{CK}^{t'}(h^{t'}; \omega) : \begin{array}{l} P_i \in \mathcal{P}_i \\ P_i \cap \text{CK}^{t'}(h^{t'}; \omega) \neq \emptyset \end{array} \right\}.$$

As $\mathcal{P}_i$ is a partition, $\mathcal{I}_i^{t'}(h^{t'}; \omega)$ is a collection of disjoint sets. Furthermore, for all agents $i, j \in V$,

$$\bigcup \mathcal{I}_i^{t'}(h^{t'}; \omega) = \bigcup \mathcal{I}_j^{t'}(h^{t'}; \omega).$$

Thus the first two conditions of [Proposition 1](#) are satisfied.

To establish the third condition of [Proposition 1](#), I need to show that $a_i^{t'}$ is optimal for all possible information sets of i in period t' , $a_i^{t'} \in c(I_i^{t'})$ for all $I_i^{t'} \in \mathcal{I}_i^{t'}(h^{t'}; \omega)$. I do so by contradiction. Suppose there exists an $\hat{I}_i^{t'} \in \mathcal{I}_i^{t'}(h^{t'}; \omega)$ such that $a_i^{t'} \notin c(\hat{I}_i^{t'})$, implying

$$\bigcup \mathcal{D}_i^{t'}(a_i^{t'}; h^{t'}; \omega) \cap \hat{I}_i^{t'} = \emptyset.$$

The definition

$$\text{CK}^{t'+1}(h^{t'+1}; \omega) = \bigcap_{i \in V} \bigcup \mathcal{D}_i^{t'}(a_i^{t'}; h^{t'}; \omega),$$

then yields

$$CK^{t'+1}(h^{t'+1}; \omega) \cap \hat{I}_i^{t'} = \emptyset,$$

contradicting the assumption that common learning ends in period t' , i.e., contradicting

$$CK^{t'+1}(h^{t'+1}; \omega) = CK^{t'}(h^{t'}; \omega).$$

I have now established the three conditions of [Proposition 1](#). Together with the assumption of a union consistent choice correspondence, we have the global indifference result. $\square$

B.1 Proving [Theorem 4](#)

B.1.1 Construction of uniform probability measure Before presenting the proof, let me construct the uniform probability measure μ on the simplex $\Delta(K)$. The simplex is a convex subset of an affine hyperplane in $\mathbb{R}^k$ with dimension $k - 1$. Hence its Lebesgue measure in $\mathbb{R}^k$ is equal to zero. I define a probability measure μ in the following steps. Consider a rotation $r: \Delta(K) \rightarrow \mathbb{R}^k$ such that for all $p, p' \in \Delta(K)$,

$$[r(p)]_1 = [r(p')]_1.$$

Next project each $r(p)$, $p \in \Delta(K)$ on $\mathbb{R}^{k-1}$:

$$\text{proj}_{K \setminus \{1\}}: r(\Delta(K)) \rightarrow \mathbb{R}^{k-1}.$$

Let λ_{k-1} be the Lebesgue measure in $\mathbb{R}^{k-1}$ and denote by $\mathcal{B}(\mathbb{R}^k)$ the Borel σ -algebra in $\mathbb{R}^k$. Define the σ -algebra of subsets of $\Delta(K)$, $\mathcal{F}(\Delta(K))$ as

$$\mathcal{F}(\Delta(K)) = \{M \cap \Delta(K) : M \in \mathcal{B}(\mathbb{R}^k)\}.$$

For the measurable space $(\Delta(K), \mathcal{F}(\Delta(K)))$, define measure λ as

$$\lambda(S) = \lambda_{k-1}(\text{proj}_{K \setminus \{1\}} r(S))$$

for $S \in \mathcal{F}(\Delta(K))$. From the measure λ , construct a probability measure μ in the usual manner; for $S \in \mathcal{F}(\Delta(K))$,

$$\mu(S) = \frac{\lambda(S)}{\lambda(\Delta(K))}.$$

PROOF OF [THEOREM 4](#). Consider the partition

$$\mathcal{P} = \left\{ \left\{ \bigvee_{i \in V} \mathcal{P}_i(\omega) \cap Q \right\}_{\omega \in \Omega}, \left\{ \bigvee_{i \in V} \mathcal{P}_i(\omega) \cap Q^C \right\}_{\omega \in \Omega} \right\}.$$

The cardinality of $\mathcal{P}$ equals k and $K = \{1, \dots, k\}$. The partition $\mathcal{P}$ can be divided into partitions $\mathcal{P}_Q$ and $\mathcal{P}_{Q^C}$, where

$$\begin{aligned} \mathcal{P}_Q &= \{P \in \mathcal{P} : P \cap Q = P\} \\ \mathcal{P}_{Q^C} &= \{P \in \mathcal{P} : P \cap Q^C = P\}. \end{aligned}$$

Each element p of the simplex $\Delta(K)$ is a probability measure over the cells of partition $\mathcal{P}$. Consider the power set of the join $2^{\bigvee_{i \in V} \mathcal{P}_i}$. The information set of each agent is an element of $2^{\bigvee_{i \in V} \mathcal{P}_i}$ for every state of the world and every history of announcements. Consider an information set $I \in 2^{\bigvee_{i \in V} \mathcal{P}_i}$ such that

$$I \cap P = \emptyset \quad \text{for all } P \in \mathcal{P}_Q$$

or

$$I \cap P = \emptyset \quad \text{for all } P \in \mathcal{P}_{Q^c}.$$

For such an information set, the agent knows with certainty whether Q has occurred. The agent does not change his opinion any more and his posterior conditioning on his information is the eventual consensus belief. In the following expression, consider only information sets $I \in 2^{\bigvee_{i \in V} \mathcal{P}_i}$ where agents do not know the occurrence of event Q with certainty, i.e.,

$$I \cap P \neq \emptyset, \quad I \cap P' \neq \emptyset$$

for some $P \in \mathcal{P}_Q$ and $P' \in \mathcal{P}_{Q^c}$. For a given information set I , let $f_I : \Delta(K) \rightarrow [0, 1]$ denote the posterior probability of event Q conditioning on I for a given probability measure:

$$f_I(p) = \frac{\sum_{P \in \mathcal{P}_Q \text{ s.t. } P \subset I} p(P)}{\sum_{P \in \mathcal{P}_Q \text{ s.t. } P \subset I} p(P) + \sum_{P \in \mathcal{P}_{Q^c} \text{ s.t. } P \subset I} p(P)}.$$

The function f_I is continuous. For two different information sets $I \neq I'$, consider the set of probability measures $S_{II'} \in \mathcal{L}_k(\Delta(K))$ that lead to the same conditional probability:

$$S_{II'} = \{p \in \Delta(K) : f_I(p) = f_{I'}(p)\}.$$

Equivalent transformations yield

$$S_{II'} = \left\{ p \in \Delta(K) : \sum_{P_Q \subset I} p(P_Q) \times \left(\sum_{P_{Q^c} \subset I'} p(P_{Q^c}) \right) = \sum_{P_Q \subset I'} p(P_Q) \times \left(\sum_{P_{Q^c} \subset I} p(P_{Q^c}) \right) \right\},$$

where P_Q denotes a typical element of $\mathcal{P}_Q$ and P_{Q^c} denotes a typical element of $\mathcal{P}_{Q^c}$. Define the function

$$f_{II'}(p) = \sum_{P_Q \subset I} p(P_Q) \times \left(\sum_{P_{Q^c} \subset I'} p(P_{Q^c}) \right) - \sum_{P_Q \subset I'} p(P_Q) \times \left(\sum_{P_{Q^c} \subset I} p(P_{Q^c}) \right).$$

Therefore, $f_{II'}(p)$ equals zero if and only if $p \in S_{II'}$. It is easy to see that $f_{II'}$ is surjective. I need to establish that $S_{II'}$ has probability zero in $\Delta(K)$. Consider the following open sets in the codomain: $(-\epsilon, 0)$, $(0, \epsilon)$, and $(-\epsilon, \epsilon)$. As $f_{II'}$ is continuous and surjective, the preimage of each of these sets is an open ball in $\Delta(K)$. Note that open balls in $\Delta(K)$ correspond to open balls in $\mathbb{R}^{k-1}$ in the mapping described in the previous subsection on the construction of the uniform probability measure. The preimage of the open interval $(-\epsilon, \epsilon)$ can be partitioned into the preimages of $(-\epsilon, 0)$, $\{0\}$, and $(0, \epsilon)$. As the preimage

of $(-\epsilon, \epsilon)$ is open, for each $p \in f_{II'}^{-1}(0)$, there exists a radius $\bar{r}_p$ such that for all open balls $B_r(p)$ in $\Delta(K)$ with $r < \bar{r}_p$, we have $p' \in B_r(p)$ implies that p' lies in $f_{II'}^{-1}(0)$, $f_{II'}^{-1}((-\epsilon, 0))$, or $f_{II'}^{-1}((0, \epsilon))$. Therefore, each $p \in S_{II'}$ is either an interior point of $S_{II'}$ or a boundary point of $f_{II'}^{-1}((-\epsilon, 0))$ or $f_{II'}^{-1}((0, \epsilon))$ (or both). But open balls in the Euclidean space are Jordan measurable and, therefore, their boundary has Lebesgue measure zero. Hence, if $\text{proj}_{K \setminus \{1\}}(r(S_{II'}))$ has positive measure in $\mathbb{R}^{k-1}$, its interior is nonempty, which implies that there exists an open ball $B_r(p)$ in $\Delta(K)$ such that $p > \bar{0}$ and for every $\hat{p} \in B_r(p)$, we have $\hat{p} \in S_{II'}$. There are three cases to consider. The first case is given by $I \cup I' \neq \Omega$. Without loss of generality, there exists $P' \in \mathcal{P}$ such that $P' \subset I \setminus I'$. Furthermore, there exists $P'' \in \mathcal{P}$ such that $P'' \subset \{I \cup I'\}^C$. Consider the following probability measure $\hat{p} \in B_r(p)$ where only the probabilities of P' and P'' are adjusted:

$$\hat{p}(P') = p(P') + \epsilon$$

$$\hat{p}(P'') = p(P'') - \epsilon.$$

Then $f_{II'}(p) = 0$ implies $f_{II'}(\hat{p}) \neq 0$, establishing a contradiction. The second case is given by $I \cup I' = \Omega$ and $I' \subset I$, and the third case is given by $I \cup I' = \Omega$, $\neg(I' \subset I)$ and $\neg(I \subset I')$. Arguments similar to those above lead to contradictions with the claim that all probability measures in $B_r(p)$ lie in $S_{II'}$ in both of the remaining cases. As the power set of the join is finite, there are finitely many pairs I, I' involving subsets of the power set of the join. For each of these finitely many pairs, the corresponding set of probability measures that lead to the same posterior announcement has probability zero. Therefore, the union over all sets $S_{II'}$ has probability zero as well, which implies that almost surely $[\mu]$ the probability measure p leads to a different posterior announcement for each information set in the power set of the join.³⁵ Hence all agents reveal their information set almost surely in every period. Consider two agents i, j with distance $d(i, j)$ in the graph. It takes $d(i, j) + 1$ periods for them to mutually (commonly) learn each others' realized partition cell. Let $d(G)$ denote the diameter of the graph. Therefore, consensus and common knowledge of the realized cell of the join are established in $d(G) + 1$ periods. $\square$

REFERENCES

- Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), "Bayesian learning in social networks." *Review of Economic Studies*, 78, 1201–1236. [24, 25]
- Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi (2010), "Spread of (mis)information in social networks." *Games and Economic Behavior*, 70, 194–227. [3]
- Aumann, Robert J. (1976), "Agreeing to disagree." *Annals of Statistics*, 4, 1236–1239. [23]
- Aumann, Robert J. (1999), "Interactive epistemology I: Knowledge." *International Journal of Game Theory*, 28, 263–300. [5]

³⁵More precisely, for all information sets where the occurrence of Q or Q^C is not known.

- Bacharach, Michael (1985), "Some extensions of a claim of Aumann in an axiomatic model of knowledge." *Journal of Economic Theory*, 37, 167–190. [2, 13, 23]
- Bala, Venkatesh and Sanjeev Goyal (1998), "Learning from neighbours." *Review of Economic Studies*, 65, 595–621. [3]
- Besley, Timothy J. and Anne Case (1994), "Diffusion as a learning process: Evidence from HYV cotton." Discussion Paper 174, Research Program in Development Studies, Woodrow Wilson School, Princeton. [1]
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as information cascades." *Journal of Political Economy*, 100, 992–1026. [24]
- Boorman, Scott A. (1975), "A combinatorial optimization model for transmission of job information through contact networks." *Bell Journal of Economics*, 6, 216–249. [1]
- Calvó-Armengol, Antoni and Matthew O. Jackson (2004), "The effects of social networks on employment and inequality." *American Economic Review*, 94, 426–454. [1]
- Cave, Jonathan A. K. (1983), "Learning to agree." *Economic Letters*, 12, 147–152. [2, 13, 23]
- Chwe, Michael S.-Y. (2000), "Communication and coordination in social networks." *Review of Economic Studies*, 67, 1–16. [2]
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003), "Persuasion bias, social influence, and unidimensional opinions." *Quarterly Journal of Economics*, 118, 909–968. [3, 5, 24]
- Foster, Andrew D. and Mark R. Rosenzweig (1995), "Learning by doing and learning from others: Human capital and technical change in agriculture." *Journal of Political Economy*, 103, 1176–1209. [1]
- Gale, Douglas and Shachar Kariv (2003), "Bayesian learning in social networks." *Games and Economic Behavior*, 45, 329–346. [23]
- Geanakoplos, John D. and Heraklis M. Polemarchakis (1982), "We can't disagree forever." *Journal of Economic Theory*, 28, 192–200. [3, 8, 14, 15, 22, 23]
- Golub, Benjamin and Matthew O. Jackson (2010), "Naïve learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2, 112–149. [3, 5, 24]
- Golub, Benjamin and Matthew O. Jackson (2012), "How homophily affects the speed of learning and best response dynamics." Unpublished paper, Stanford University. [22, 24]
- Ioannides, Yannis M. and Linda D. Loury (2004), "Job information networks, neighborhood effects, and inequality." *Journal of Economic Literature*, 42, 1056–1093. [1]
- Jackson, Matthew O. (2008), *Social and Economic Networks*. Princeton University Press, Princeton. [2]

- Krasucki, Paul (1996), "Protocols forcing consensus." *Journal of Economic Theory*, 70, 266–272. [2, 13, 23]
- Manea, Mihai (2011), "Bargaining in stationary networks." *American Economic Review*, 101, 2042–2080. [2]
- Mueller-Frank, Manuel (2012), "Does one Bayesian make a Difference?" Unpublished paper, University of Oxford. [4]
- Munshi, Kaivan (2003), "Networks in the modern economy: Mexican migrants in the U.S. labor market." *Quarterly Journal of Economics*, 118, 549–597. [1]
- Munshi, Kaivan (2004), "Social learning in a heterogeneous population: Technology diffusion in the Indian green revolution." *Journal of Development Economics*, 73, 185–213. [1]
- Opp, Karl-Dieter and Christiane Gern (1993), "Dissident groups, personal networks, and spontaneous cooperation: The East German revolution of 1989." *American Sociological Review*, 58, 659–680. [2]
- Parikh, Rohit and Paul Krasucki (1990), "Communication, consensus, and knowledge." *Journal of Economic Theory*, 52, 178–189. [2, 23]
- Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2009), "Information externalities and emergence of consensus." *Games and Economic Behavior*, 66, 979–994. [2, 3, 8, 23]
- Smith, Lones and Peter Sørensen (2000), "Pathological outcomes of observational learning." *Econometrica*, 68, 371–398. [14, 24]
- Snow, David A., Louis A. Zurcher, and Sheldon Ekland-Olson (1980), "Social networks and social movements: A microstructural approach to differential recruitment." *American Sociological Review*, 45, 787–801. [2]
- Topa, Giorgio (2001), "Social interactions, local spillovers and unemployment." *Review of Economic Studies*, 68, 261–295. [1]
- Udry, Christopher and Timothy Conley (2001), "Social learning through networks: The adoption of new agricultural technologies in Ghana." *American Journal of Agricultural Economics*, 83, 668–673. [1]
- Wang, Ping and Quan Wen (2002), "Network bargaining and organizational hierarchy." Unpublished paper, Vanderbilt University. [2]