

A general analysis of boundedly rational learning in social networks

MANUEL MUELLER-FRANK

Department of Economics, IESE Business School, University of Navarra

CLAUDIA NERI

Economics & Markets, Zurich Insurance Company

We analyze boundedly rational learning in social networks within binary action environments. We establish how learning outcomes depend on the environment (i.e., informational structure, utility function), the axioms imposed on the updating behavior, and the network structure. In particular, we provide a normative foundation for quasi-Bayesian updating, where a quasi-Bayesian agent treats others' actions as if they were based only on their private signal. Quasi-Bayesian updating induces learning (i.e., convergence to the optimal action for every agent in every connected network) only in highly asymmetric environments. In all other environments, learning fails in networks with a diameter larger than 4. Finally, we consider a richer class of updating behavior that allows for nonstationarity and differential treatment of neighbors' actions depending on their position in the network. We show that within this class there exist updating systems that induce learning for most networks.

KEYWORDS. Social networks, naïve inference, naïve learning, bounded rationality, consensus, information aggregation.

JEL CLASSIFICATION. D83, D85.

1. INTRODUCTION

In situations of uncertainty, when agents have partial and private information, observational learning is a crucial component of human interaction. Among the many possible mechanisms by which individuals learn from others, observational learning describes

Manuel Mueller-Frank: mmuellerfrank@iese.edu

Claudia Neri: claudia.neri@zurich.com

This paper contains and extends results from earlier versions titled “Quasi-Bayesian Updating in Social Networks” and “Social Learning in Networks: Theory and Experiments.” We are grateful to the three anonymous referees for their comments and suggestions. We have benefited from the opportunity to present earlier versions of this paper at University of Berkeley, LBS, WISE 2016 (Caltech), IMEBESS 2016 (Rome), Carlos III, European University Institute, and Technion. We thank seminar participants, Itai Arieli, Arun Chandrasekhar, Adam Szeidl, and Xavier Vives for their comments. Manuel Mueller-Frank is a Fellow of the Public–Private Sector Research Center at IESE Business School and gratefully acknowledges the financial support of the Spanish Ministry of Economy and Competitiveness (Project ECO2015-63711-P). Claudia Neri gratefully acknowledges the financial support of the Swiss National Science Foundation (Grant 100014/135346).

the process by which an individual draws inferences on the information held by other people based on the observation of their behavior. Understanding how individuals update their behavior after observing the behavior of others and which long-run aggregate outcomes such learning generates has important implications for policy. For example, observational learning within a rural community may explain the effectiveness of an information campaign aimed at increasing the adoption of microfinance loans (Banerjee et al. 2013) or of a targeting program aimed at selecting aid beneficiaries (Alatas et al. 2016).

In the literature, two predominant approaches to the study of observational learning have emerged: Bayesian and boundedly rational. In the Bayesian approach, agents are assumed to learn rationally, i.e., they make inferences about the private information of all agents based on the interaction structure and the observed decisions.¹ This is the standard approach in the sequential social learning literature (Bikhchandani et al. 1992, Banerjee 1992, Smith and Sørensen 2000, Acemoglu et al. 2011, Arieli and Mueller-Frank 2019, Lobel and Sadler 2015) and in parts of the literature on repeated interaction in social networks (Gale and Kariv 2003, Rosenberg et al. 2009, Mueller-Frank 2013, Mossel et al. 2015). Despite being a useful benchmark, the Bayesian approach has a severe weakness: the rationality assumption appears unrealistic due to the computational sophistication necessary to make inferences. This is especially true in an incomplete network where agents interact repeatedly. Here every agent has to draw indirect inferences regarding the private information of all agents, based only on the actions that he observes. Hazla et al. (2018) show the conceptual and computational complexity required by Bayesian updating to be extremely high. Such complexity represents a weakness, since experimental evidence suggests that humans are unlikely to engage in such complex cognitive tasks.² To reduce the cognitive complexity inherent in Bayesian updating, the boundedly rational approach assumes instead that agents use simple rules of thumb. For this reason, the boundedly rational approach is especially suitable—and, in fact commonly employed—in complex settings, such as settings of repeated interaction in social networks. DeGroot (1974) provides the standard model within the boundedly rational approach. Its original formulation describes agents who repeatedly communicate beliefs about an underlying state of the world and revise their beliefs to a weighted average of their own and their neighbors' previous beliefs. A more recent formulation, the so-called DeGroot action model, is applied to the study of observational learning in environments with binary states and binary actions where agents, instead of communicating beliefs, observe actions (Chandrasekhar et al. 2020). As a boundedly rational model, the DeGroot model reduces the cognitive complexity of the updating process and its main strength derives from its tractability. However, this comes at the cost of a lack of generality, since the use of a weighted average updating function in the crucial step of belief formation is somewhat arbitrary. Therefore, both the Bayesian approach and the DeGroot model, as the standard formulation of the boundedly rational

¹Throughout the paper, we use the terms “action,” “choice,” and “decision” as synonyms.

²See Kübler and Weizsäcker (2004), Eyster et al. (2015), and March and Ziegelmeyer (2018) for evidence of redundancy neglect, and see Enke and Zimmerman (2019) for evidence of correlation neglect.

approach, have weaknesses that limit their scope. This paper is motivated by such unresolved weaknesses.

This paper focuses on the information aggregation properties of boundedly rational learning in a setting where agents interact repeatedly in a finite network. We assume that all agents share a common prior over the space of possible states of the world and each agent initially receives a private signal drawn from a common set of signals according to a state-dependent distribution. Thereafter, in each of countable periods, all agents simultaneously select an action while observing the previous actions chosen by their neighbors in the network. As is common in the literature, we focus on informational externalities. That is, we assume that all agents share the same continuous utility function, where an agent's utility depends only on his own action and the realized state of the world. In the first period, each agent selects the utility maximizing action conditional on his signal. In the second and all subsequent periods, he updates his action by taking into account the only additional information he receives: the actions he observes. The bounded rationality assumption concerns the way actions are updated. Rather than imposing a particular rule of thumb (or updating function) as in the DeGroot model, we consider a general class of updating behavior defined by three axioms: Markov property, stationarity, and locality. The *Markov property* requires that the updated action of an agent at time t depends only on the actions he observed in period $t - 1$ and is, thus, invariant to the history prior to $t - 1$. *Stationarity* requires that the updated behavior is invariant in the time periods. These axioms jointly imply that the updating behavior of an agent can be described by an updating function that in every period he applies to the previous-period action vector. The Markov property and stationarity axioms are standard in the literature on boundedly rational learning on networks.³ Finally, *locality* requires that the updated action is invariant to the actions of nonneighbors and the structure of the network beyond the neighborhood. Thus, we define the updating behavior of each agent not only for one given network, but rather for all possible networks.

Our analysis focuses on a binary action environment.⁴ A good example for environments that are accurately modeled with binary actions are repeated adoption decisions, such as adhering to a certain diet. We are interested in identifying the environments and the particular updating behavior within our general class that optimize information aggregation. Thus, we approach the question of selecting a particular type of updating behavior from a normative perspective, based on desirable aggregate properties in terms of information aggregation.

Our analysis considers a general state and signal space, and any utility function that is continuous in the state space for either action. We first aim to identify the subclass of stationary local Markov updating systems that optimizes information aggregation, independent of the structure of the network.⁵ Assume that the private signals of agents are

³See, for example, DeGroot (1974), Golub and Jackson (2010, 2012), Jadbabaie et al. (2012), and Chandrasekhar et al. (2020).

⁴Binary actions are a standard assumption in the literature on Bayesian updating; see, for example, Bikhchandani et al. (1992), Smith and Sørensen (2000), Acemoglu et al. (2011), Mossel et al. (2015), and Mossel et al. (2020).

⁵An updating system denotes the tuple of updating functions of all agents.

independent and identically distributed (i.i.d.) conditional on the realized state of the world. Fix an environment described by the set of agents, the state, and signal space, the utility function and the joint probability distribution over the product space of states, and signals across agents. A stationary local Markov updating system satisfies *learning*, if, for every first-period action vector, within finite time, the actions of all agents converge to the action that is optimal, conditional on the realized first-period action vector, in any connected undirected⁶ network.

Our first main result characterizes the environments in which learning may occur. We say that an environment satisfies *Bayesian contagion* if there exists a *contagion action* that is expected to be utility maximizing conditional on every first-period action vector except the action vector that satisfies consensus in the noncontagion action. Intuitively, a Bayesian observer of the first-period action vector would select the contagion action whenever it occurs at least once throughout the network and he would thus be “infected.” [Theorem 1](#) establishes that if a stationary local Markov updating system $\mathbf{f}$ satisfies learning, then the environment satisfies Bayesian contagion. [Theorem 2](#) considers Bayesian contagion environments and identifies a particular type of updating behavior that is necessary for learning under stationary local Markov updating systems: quasi-Bayesian updating.⁷ The concept of *quasi-Bayesian* updating is very simple: When observing a set of actions being chosen by other agents, the observer assumes that each action is optimal given (only) the private information of the agent who chose it. In other words, a quasi-Bayesian updater treats others’ actions as if they are based only on their signal. We emphasize that we are not the first to use such a behavioral assumption. The concept originated in [Eyster and Rabin \(2010\)](#), and is also present in [Bohren \(2016\)](#) and, to some extent, in [Bala and Goyal \(1998\)](#). However, the environment we analyze differs from those considered in such previous works.⁸

Quasi-Bayesian updating addresses successfully the weaknesses of Bayesian updating and DeGroot updating. First, it reduces the complex cognitive tasks required by Bayesian updating, since considerations as to how each observed action might have been affected by other actions—both observed and not observed—are not necessary. Second, it is applicable to action spaces of arbitrary cardinality and allows the quasi-Bayesian updating function to vary with the environment considered.⁹ We next analyze quasi-Bayesian updating in environments where Bayesian contagion is not satisfied. Thus actions fail to converge to the optimal action in some network structures by [Theorem 1](#). We then ask how the structure of the network affects whether learning in a given network succeeds or fails. We focus on *monotone* environments that have the property that the quasi-Bayesian updating satisfies a threshold rule, i.e., action a is optimal if and

⁶In an undirected network, edges are symmetric. A network is connected if there exists a path connecting every pair of agents.

⁷More precisely, [Theorem 2](#) shows that within a subclass of stationary local Markov updating systems, quasi-Bayesian updating is necessary. Proposition 1 then shows that if a stationary local Markov updating system satisfies learning, then the updating functions of all agents coincide with quasi-Bayesian updating for most observed action vectors.

⁸We provide more details on the relation in [Section 2](#).

⁹While our analysis focuses on binary actions, the concept can be directly applied to any social learning environment where agents have an underlying utility function, with an arbitrary action space.

only if at least a certain proportion of observed actions coincide with a . For a fixed network G , we say that a quasi-Bayesian updating system yields *learning in network G* if the actions of all agents converge to the action that is optimal conditional on the initial action vector. [Theorem 3](#) establishes that if the diameter of a network G is larger than 4, then learning in G fails for quasi-Bayesian updating. [Proposition 2](#) shows that learning also fails in networks containing small *cohesive* groups, i.e., groups of agents where each agent has a relatively large share of neighbors within the group.

The negative result stated in [Theorem 3](#) can be traced back to quasi-Bayesian updating failing to satisfy jointly two intuitive properties in the given network structures: information diffusion and information retention. To explain *information diffusion*, first note that since Bayesian contagion is not satisfied, a single occurrence of either action is not sufficient to make this action optimal given the first-period action vector. Thus for the optimal action to eventually spread throughout the network, it is necessary that if any agent is the only one in his neighborhood to select action a , then she updates to action $-a$. *Information retention* requires that if action a is optimal conditional on the first-period action vector, then a is optimal on the action vector of every later period. If the diameter is larger than 4, then for some initial action vectors, the set of agents who initially select the optimal action are distributed throughout the network in such a way that information retention fails. The idea is that one agent, who is isolated from the rest of the agents selecting the optimal action, switches from the optimal to the suboptimal action without swaying the action of any of his neighbors. For some initial action vectors, this is sufficient for the optimal action conditional on the second-period action vector to differ from that of the first period, thus inducing a failure of information retention.

Finally, we return to the general analysis of boundedly rational updating and we relax the locality axiom to weak locality. *Weak locality* allows agents to take into account the overall network structure, and, in particular, the network position of their neighbors, when updating based on the actions of their neighbors. We first show that relaxing locality alone does not improve learning outcomes in monotone environments. [Theorem 4](#) shows that if a stationary weakly local Markov updating system achieves learning, then the environment satisfies Bayesian contagion. Indeed learning fails under such updating systems even in networks such as a star that should be conducive to learning due to the fact that the center agent observes all others. When jointly relaxing stationarity and locality, however, learning can be achieved under a weak condition on the network. Let n denote the size of the network. [Theorem 5](#) shows that if there exists a complete subgroup¹⁰ of size greater or equal to $\frac{\ln n}{\ln 2} + 1$, then there exists a nonstationary weakly local Markov updating system that achieves learning. Here, the complete subgroup overcomes the barrier of information retention by coordinating and retaining all relevant information in its joint action vector. Once the complete subgroup reaches the optimal action, information diffusion can be easily achieved as the updating functions need not be stationary. Thus, when the assumptions of stationarity and locality are relaxed, there exists an updating system that dramatically improves learning outcomes compared to

¹⁰A complete subgroup is a group where each member is a neighbor of each other member of the group.

quasi-Bayesian updating, where, instead, learning fails in any network with a diameter larger than 4.

The paper is organized as follows. In [Section 2](#), we discuss the related literature. In [Section 3](#), we introduce a general model of boundedly rational updating and a particular type thereof: quasi-Bayesian updating. [Section 4](#) provides a characterization of the environments in which learning can occur and a normative foundation for quasi-Bayesian updating. In [Section 5](#), we analyze quasi-Bayesian updating and establish conditions on the network structure that prevent learning. In [Section 6](#), we relax the stationarity and locality assumptions, and provide both a negative and a positive learning result. [Section 7](#) concludes. Most proofs are presented in the [Appendix](#).

2. RELATED LITERATURE

This paper contributes to the theoretical literature on repeated interaction in social networks and is related to previous work both on Bayesian updating and on boundedly rational updating. The behavioral concept of quasi-Bayesian updating originated in and was developed by [Eyster and Rabin \(2010\)](#), who analyze the implication of this assumption, which they call naïve inference, in the context of a sequential social learning model with binary states and uncountable actions.¹¹ [Gagnon-Bartsch and Rabin \(2016\)](#) apply the naïve inference concept to a social learning model with finite actions and states. In each period a generation of agents makes an irreversible choice to observe only the actions of the previous generation. They provide results on social mislearning where some states, despite being realized, are disbelieved over time. Even earlier, the seminal paper by [Bala and Goyal \(1998\)](#) on learning in networks imposes a similar assumption, albeit in a different setting where each agent receives private information in every period. We see our analysis as complementary to the earlier work as we provide a normative foundation for quasi-Bayesian updating in environments of repeated interaction in a social network.

Within the literature on boundedly rational learning in networks, our paper relates closely to [Molavi et al. \(2018\)](#), [Dasaratha and He \(2020\)](#), and [Feldman et al. \(2014\)](#). [Molavi et al. \(2018\)](#) consider local Markov updating functions in a setting where agents repeatedly communicate their beliefs and axiomatize a linear and a log-linear updating function. Under such updating functions, they then provide long-run learning results in a model where each agent receives a private signal in every period. [Dasaratha and He \(2020\)](#) apply the idea of quasi-Bayesian updating to the sequential social learning model with binary states, infinite actions, and Gaussian signals. They provide a necessary and sufficient condition for agents to act optimally in the limit. [Feldman et al. \(2014\)](#) analyze a setting where agents repeatedly update their action according to the observed majority in a network. In their paper, only one agent at a time updates his action, rather than all agents simultaneously doing so in every period, as in our paper. They analyze the process of actions for large networks and are concerned with the structural conditions

¹¹[Bohren \(2016\)](#) also considers a type of naïve inference in the standard sequential social learning model. Here each agent with probability p either observes the actions of all predecessors or observes no history at all. She analyzes the case where agents have misspecified beliefs about p .

on the network such that agents eventually converge to the correct action with large probability. Due to the different setting and the different notion of learning, our results require different proof techniques and are not directly comparable with theirs.¹² They find that in sparse larger-diameter networks, agents eventually learn with high probability, a result that differs sharply from our result on the failure of learning in networks with diameter larger than 4.¹³

Our analysis of quasi-Bayesian updating in a binary action environment mirrors Mossel et al. (2015), who analyze Bayesian learning in a model with binary actions and binary states. They show that despite the action space being coarse, learning occurs as long as the network satisfies bounded out-degree (for infinite networks) and a weak structural condition, which is satisfied in undirected networks. We show that under quasi-Bayesian updating, information aggregation might occur for all networks and agents, even in an environment with coarse actions. We characterize the environments that allow for such information aggregation. While under Bayesian updating, convergence to the optimal action is asymptotic (see Mossel et al. 2015), under quasi-Bayesian updating, information either is aggregated along the shortest path or fails to aggregate.

Our paper is also related to the literature on best-response dynamics in local interaction games, in particular, Morris (2000). Our Proposition 2 borrows the concept of group cohesiveness from Morris (2000) so as to provide a structural network condition necessary for learning.

The concept of quasi-Bayesian updating can be linked to the concept of redundancy neglect in environments where agents take a one-time decision sequentially while commonly observing the actions of their predecessors. In such environments, a fully rational agent who engages in observational learning should realize that other agents, whose actions he observed, are also engaging in observational learning and that, therefore, there is redundancy in their actions, which he should account for. Experimental evidence shows that redundancy neglect is common (Kübler and Weizsäcker 2004), that even mild redundancy neglect can be harmful (Eyster et al. 2015), and that many subjects believe naïvely that each observable choice reveals a substantial amount of another person's private information (March and Ziegelmeyer 2018).

Finally, there are several empirical papers on boundedly rational social learning in networks that are complementary to our work. Chandrasekhar et al. (2020) provide experimental evidence in support of a DeGroot action model where each agent selects the action that in the previous period was chosen by the majority of agents in his neighborhood. In their setting, such updating behavior coincides with quasi-Bayesian updating.¹⁴

¹²Nevertheless, for some environments, local majority functions are indeed quasi-Bayesian.

¹³The sharp difference in results, albeit in different settings, can be partially explained through our more demanding learning notion, which requires convergence to the optimal action conditional on every first-period action vector, as opposed to the probabilistic notion employed in Feldman et al. (2014).

¹⁴Alatas et al. (2016) find support for an extension of the DeGroot model where the weight assigned to a neighbor depends on his distance from the information source. They assume that agents treat each piece of information received from their neighbors as independent signals. Despite differences in setting, this assumption resembles the idea behind quasi-Bayesian updating. Mobius et al. (2015) develop a “streams” model where agents tag information by describing its origin. Quasi-Bayesian updating is loosely related to the imperfect tagging of information originating from a long distance, which they report.

We conclude by emphasizing the connection between our work and the literature on level- k models (Nagel 1995, Camerer et al. 2004, Crawford et al. 2013). The concept of quasi-Bayesian updating is grounded in the concept of limited depth of reasoning. Quasi-Bayesian updating can be loosely interpreted as the behavior of a level-1 player, who receives no private signal and best responds to level-0 players, who take the action based only on their private signal.

3. A GENERAL MODEL OF BOUNDEDLY RATIONAL UPDATING

Consider a finite set of agents N , each of whom faces uncertainty regarding the state of the world $\omega \in \Omega$. The cardinality of N is denoted by n . Agents share a common prior p over the compact metrizable state space Ω endowed with Borel σ -algebra $\mathcal{F}_\Omega$. The state of the world is unknown to the agents, but each agent i observes a private signal s_i drawn independently from a standard Borel space S according to a state-dependent distribution $F_\omega \in \Delta(S)$. We assume that for any two states ω and ω' , the probability measures F_ω and $F_{\omega'}$ are absolutely continuous with respect to each other but not identical. This implies that signals have some information value but do not reveal any a priori noncertain event to be realized with probability 1.

The agents are located in an undirected, connected network $G = (N, E)$.¹⁵ A network is a pair of sets (N, E) such that $E \subset [N]^2$. The elements of N are nodes of the graph, representing the agents, and the elements of E are the edges of the graph, representing the direct connections between agents. The neighborhood of agent i consists of all agents $j \in N$ such that there exists an edge $ij \in E$. In an undirected network, edges are symmetric, that is, if i is a neighbor of j , then j is a neighbor of i . A network is connected if there exists a path connecting every pair of agents.

Interactions take place in discrete time. At time $t = 0$, the state of the world ω is drawn according to the prior p , the signals of agents are drawn independently according to F_ω , and each agent i observes his private signal s_i . At each time $t \geq 1$, all agents $i \in N$ simultaneously take a binary action $a_i \in \mathcal{A} = \{0, 1\}$. We assume that all agents share identical preferences represented by a continuous utility function $u : \mathcal{A} \times \Omega \rightarrow \mathbb{R}$. As is the norm in the observational learning literature, we restrict attention to settings without payoff externalities. In the following discussion, we denote the tuple $(N, \Omega, S, p, F_\omega, u)$ as the (binary action) *environment*.

One good example for environments that are accurately modeled with binary actions are repeated adoption decisions, such as adhering to a certain diet.¹⁶ Binary action environments have been thoroughly analyzed within the literature on Bayesian updating, most commonly to study herding in sequential environments (Bikhchandani et al. 1992, Smith and Sørensen 2000, and Acemoglu et al. 2011), but also to study repeated interaction (Mossel et al. 2015, 2020). Differently from existing models with binary actions, we allow for a general state space Ω , and any continuous utility function on $\mathcal{A} \times \Omega$. In

¹⁵In the following discussion, the terms “graph” and “network” are used interchangeably.

¹⁶Such adoption environments might be appropriately modeled via a binary state space, as is the case in the sequential social learning literature, or a rich state space where a (two-dimensional) state might directly represent the utility of selecting either action.

line with the existing models, we assume that the signals of agents are i.i.d. conditional on the realized state.

We first describe how agents select their action in the first period $t = 1$ when their information consists only of their private signal. We assume that each agent selects a first-period action that maximizes his expected utility conditional on his signal. Formally, a *strategy* σ_i of agent i is a measurable mapping that assigns an action to every realized signal, such that the action maximizes expected utility conditional on the realized signal:

$$\sigma_i(s_i) \in \arg \max_{a \in A} \int u(a, \omega) \Pr(d\omega | s_i).$$

Here $\Pr(d\omega | s_i)$ denotes the posterior distribution over the state space based on the prior p and conditional on the realized signal s_i . In summary, in period $t = 1$, all agents act as a myopic, first-period expected utility maximizing Bayesian agent would.

We make three weak technical assumptions on the environment $(N, \Omega, S, p, F_\omega, u)$ that have important implications for the strategies of agents: non-indifference, nontriviality, and vector non-indifference. So as to define these properties, fix an environment and one agent i , and consider any pair of strategies σ'_i, σ''_i . *Non-indifference* requires that almost surely $\sigma'_i = \sigma''_i$, that is, for a probability 1 subset $S' \subset S$, we have $\sigma'_i(S') = \sigma''_i(S')$. As signals are identically distributed across agents, each agent i is indifferent conditional on his signal with probability 0 and there exists an almost surely unique strategy, which is applied by all agents. Denote this strategy by σ^* and denote by σ^* the corresponding profile of strategies $\langle \sigma_i \rangle_{i \in N}$, where $\sigma_i = \sigma^*$ for each agent $i \in N$. We denote by $a_i^{\sigma^*}$ the random action generated by agent i 's first-period signal and the strategy σ^* . The corresponding random action vector, which takes values in $\{0, 1\}^n$, is denoted by $\mathbf{a}^{\sigma^*}$. Note that since signals are conditionally i.i.d., for any two agents i, j their actions $a_i^{\sigma^*}, a_j^{\sigma^*}$ are conditionally i.i.d. as well.

The second assumption on the environment—*nontriviality*—requires that both actions are chosen with positive probability by each agent, i.e., for $a = \{0, 1\}$, there exists a positive probability subset $S_a \subset S$ such that $\sigma^*(S_a) = a$. Without nontriviality all agents would select the same action independent of their signal, preventing any inference on signals based on observed actions. In such an environment, social learning considerations would be meaningless.

Finally, to introduce the last assumption, consider the Bayesian belief over the state space that is induced by observing a realization of the first-period action vector $\mathbf{a}^1$, drawn from $\mathbf{a}^{\sigma^*}$. *Vector non-indifference* requires that for each possible realization $\mathbf{a}^1 \in A^n$, the corresponding Bayesian belief induces a *unique* expected utility maximizing action for the utility function u . Let $\beta : A^n \rightarrow A$ denote the function that assigns to each first-period action vector the corresponding unique expected utility maximizing action:

$$\beta(\mathbf{a}) = \arg \max_{a \in A} E[u(a, \omega) | \mathbf{a}^{\sigma^*} = \mathbf{a}].$$

We refer to β as the *Bayesian complete observation function*.

3.1 Axioms on boundedly rational updating

Next we describe the procedure by which agents revise their actions in any period $t \geq 2$. We take a general, axiomatic approach to boundedly rational updating in the spirit of Molavi et al. (2018). In particular, we impose three axioms on action updating: the Markov property, stationarity, and locality. We say that the updating behavior of agent i satisfies the Markov property if his updated action in period t depends only on the action vector of period $t - 1$ and is thus invariant to the action vectors of all earlier periods. A straightforward implication of the Markov assumption is that signals play no role for the process of actions beyond determining the action vector in $t = 1$. Stationarity requires that the updating behavior is time-invariant, that is, the updated action is only a function of the previous-period action vector and is invariant to the actual time period. The Markov property and stationarity are standard axioms in the literature on boundedly rational learning on networks. For example, they are explicitly or implicitly imposed in rich action models by DeGroot (1974), Golub and Jackson (2010, 2012), and Jadbabaie et al. (2012) among others, and in a binary action model by Chandrasekhar et al. (2020). Jointly imposing the Markov property and stationarity allows for a more tractable analysis because it implies that the updating behavior of each agent i can be represented by a time invariant *updating function* $f_i : A^n \times \mathcal{G} \rightarrow A$ that assigns an action a_i^t to each pair of action vector $\mathbf{a}^{t-1}$ and network structure G , $a_i^t = f_i(\mathbf{a}^{t-1}, G)$, for every $t \geq 2$.¹⁷ The resulting functional representation is appealing because it reduces complexity vis-à-vis Bayesian updating. Complexity is reduced not only by using the same updating function in all periods, but also by having such an updating function depend only on the previous-period action vector, and not on the full history of actions as is typically the case for Bayesian updating. Thus, the Markov property and stationarity axioms imply that boundedly rational updating boils down to the application of a simple rule of thumb: in our case, the application of a function that assigns one action to the previous-period action vector.

Consider a mapping $\mathbf{f} : A^n \times \mathcal{G} \rightarrow A^n$ such that

$$\mathbf{f}(\mathbf{a}, G) = (f_1(\mathbf{a}, G), \dots, f_n(\mathbf{a}, G)),$$

where the updating functions f_i are the components of $\mathbf{f}$. The mapping $\mathbf{f}$ is denoted an *updating system*. For a given network G , the sequence of action vectors $\{\mathbf{a}^t\}_{t \in \mathbb{N}}$ can now be recursively defined as follows: for all $t = 2, 3, \dots$,

$$\mathbf{a}^t = \mathbf{f}(\mathbf{a}^{t-1}, G) = (f_1(\mathbf{a}^{t-1}, G), \dots, f_n(\mathbf{a}^{t-1}, G)).$$

Hence, for a given first-period action vector $\mathbf{a}^1$, the process $\{\mathbf{a}^t\}_{t \in \mathbb{N}}$ is a deterministic stationary Markov process, determined by the updating system $\mathbf{f}$.

The third axiom, locality, relates the updating behavior to the network structure G by imposing conditions on the updating system $\mathbf{f}$. For agent i and network G , let $N_i(G)$ denote the set of neighbors of i including agent i :

$$N_i(G) = \{j \in N : ij \in E\} \cup \{i\}.$$

¹⁷As is common in the literature, we make the implicit assumption that updating is deterministic. For an analysis of random updating functions, $f_i : A^n \times \mathcal{G} \rightarrow \Delta(A)$, please see Arieli and Mueller-Frank (2017).

Further, for a set of agents $M \subset N$ and a vector $\mathbf{a} \in A^n$, let $\mathbf{a}_M \in A^m$ be the subvector of actions of agents in M , i.e., where $[\mathbf{a}_M]_i = a_i$ for all $i \in M$.

DEFINITION 1. A Markov updating function $f_i : A^n \times \mathcal{G} \rightarrow A$ is *local* if (i) for all pairs $(\mathbf{a}, G)$ and $(\mathbf{a}', G')$ such that $N_i(G) = N_i(G')$ and $\mathbf{a}_{N_i(G)} = \mathbf{a}'_{N_i(G')}$, we have $f_i(\mathbf{a}, G) = f_i(\mathbf{a}', G')$, and (ii) for all G and $j \in N_i(G)$, there exists $\mathbf{a}_{-j} \in A^{n-1}$ and $a \neq a'$ such that

$$f_i((a, \mathbf{a}_{-j}), G) \neq f_i((a', \mathbf{a}_{-j}), G).$$

The first condition for locality requires that the updated action of agent i depends only on the actions of his neighbors (which include his own action), and is invariant to actions of nonneighbors and to the structure of the network as a whole.¹⁸ The second condition states that agent i 's updated action is not invariant in the actions of any of his neighbors. More precisely, a local function requires that for each observed agent $j \in N_i(G)$, there exists a vector of observed actions such that agent j 's action is pivotal for agent i 's updated action. That is, for each neighbor j of i , there exists an action vector $\mathbf{a}_{-j}$ such that changing j 's action changes the updated action of agent i . In other words, we require that agent i does not completely ignore any of the agents he observes.

Note that locality relates the updating function of an agent to his neighborhood structure. Beyond the set of neighbors, however, a given agent's updating function takes no other information regarding the structure of the network into account. This aspect of locality can be motivated in two different ways. First, agents might simply have limited or no knowledge of the structure of the network beyond their neighbors. Intuitively, this is a reasonable assumption for large, real-world social networks where agents have a large number of neighbors. Empirical evidence indeed supports this assumption. [Breza et al. \(2018\)](#) find that even at a village level, network knowledge is low and highly localized. [Krackhardt \(1990\)](#) and [Casciaro \(1998\)](#) find similar patterns in networks of 36 (the former) and 25 members (the latter). These empirical results can be seen as supporting our locality assumption. Second, when agents have to make indirect inferences regarding the private information of non-observed agents, the computational complexity of Bayesian updating (discussed above) comes into play. Boundedly rational agents might simply be unable to incorporate the additional structural information into their updating procedure and instead focus only on their neighborhood. Finally, it is worth highlighting that a myopic Bayesian agent also satisfies locality when updating his action in the second period since the first-period action of each agent is based only on his signal.¹⁹

We see our axioms as one natural way to model boundedly rational updating in a dynamic network environment, but certainly not as the only natural way. In [Section 6](#), we discuss how relaxing the locality axiom, first by itself and then jointly with the stationarity axiom, impacts the long-run information aggregation properties of the induced action process.

¹⁸This condition is naturally satisfied under the assumption that agents have no information about the network beyond their neighbors.

¹⁹A myopic Bayesian agent forms a fully rational belief based on the observed history and selects an action that maximizes his expected utility.

3.2 Definition of learning

Information aggregation lies at the core of our analysis. In particular, we analyze different updating processes in terms of the information aggregation properties of the long-run actions. To understand the type of information aggregation we have in mind, first note that for an updating system $\mathbf{f}$, the initial action vector $\mathbf{a}^1$ determines the subsequent process of action vectors $\{\mathbf{a}^t\}_{t \in \mathbb{N}}$. This implies that, from an information aggregation perspective, the best possible outcome is that eventually all agents act optimally conditional on $\mathbf{a}^1$, rather than, let us say, the realized vector of signals $\mathbf{s}$.²⁰ For a fixed environment, we define two notions of learning: one independent of the network structure and the other for a given network structure. We first state the stronger learning notion.

DEFINITION 2. A stationary local Markov updating system $\mathbf{f}$ yields *learning* if there exists a period t^* such that for every $(\mathbf{a}^1, G) \in A^n \times \mathcal{G}$, every agent $i \in N$, and every $t > t^*$, we have $a_i^t = \beta(\mathbf{a}^1)$.

Thus, learning requires that, in all undirected connected networks and for every initial action vector $\mathbf{a}^1$, eventually the action a_i^t of each agent i coincides with the Bayesian complete observation action $\beta(\mathbf{a}^1)$. Thus, the local information contained in the agents' first-period actions is aggregated over time until all agents eventually select the Bayes optimal action given the first-period action vector $\mathbf{a}^1$. Note that by the vector, non-indifference assumption β is unique. Therefore, learning implies consensus among all agents.

While a strong requirement, notions of learning that require information aggregation for all undirected, connected networks are standard in the Bayesian learning literature. For results establishing information aggregation for any finite, connected, and undirected network, see [Mueller-Frank \(2013, 2014\)](#), and [Arieli and Mueller-Frank \(2017\)](#). We complement the network-independent learning notion with a weaker one that requires information aggregation only for a given network G . We call this property *learning in network G* .

DEFINITION 3. Fix a network $G \in \mathcal{G}$. An updating system $\mathbf{f}$ yields *learning in network G* if there exists a period t^* such that $a_i^t = \beta(\mathbf{a}^1)$ for every $\mathbf{a}^1 \in A^n$, every agent $i \in N$, and every $t > t^*$.

3.3 Quasi-Bayesian updating

The existing literature on boundedly rational learning either imposes a particular functional form (the most prominent example being the DeGroot model) or provides an axiomatization of particular functional forms; see [Molavi et al. \(2018\)](#). Instead, the axioms

²⁰For some environments it is the case that two signal vectors $\mathbf{s}$ and $\mathbf{s}'$ induce the same first-period action vector $\mathbf{a}^1$ but different expected utility maximizing actions

$$\arg \max_{a \in A} E[u(a, \omega) | \mathbf{s}] \neq \arg \max_{a \in A} E[u(a, \omega) | \mathbf{s}'].$$

However, the action process resulting from $\mathbf{s}$ and $\mathbf{s}'$ is identical, making information aggregation relative to signals impossible.

we introduced so far do not impose a specific functional form; neither do they relate the updating procedure to the underlying information structure or utility function.

We now introduce a particular form of updating behavior, satisfying the three axioms, that varies with the information structure and the utility function. Some adaptation of updating behavior to the environment is required, as the focus of our analysis is on information aggregation. Clearly, if the actual updating procedure does not change with the environment, learning is doomed to fail since the optimal action given the first-period action vector does vary with the environment.

The particular adaptive approach we propose in a sense reconciles Bayesian updating with the limitations imposed by our axioms. For an agent i with neighborhood N_i , denote by $\mathbf{a}_{N_i}^{\sigma^*}$ the vector of random actions of agents in N_i given their realized signals and strategy profile σ^* . A Bayesian observation function β_{N_i} assigns an expected utility maximizing action to each realized first-period vector $\mathbf{a}_{N_i}^1 \in A^{|N_i|}$ drawn according to $\mathbf{a}_{N_i}^{\sigma^*}$ in two steps. First, it assigns the Bayesian belief over the state space conditional on $\mathbf{a}_{N_i}^1$. Second, it selects the expected utility maximizing action given the belief. Formally, $\beta_{N_i} : A^{|N_i|} \rightarrow A$.²¹ Thus, one natural way to define a local updating function f_i with a Bayesian foundation is by simply setting $f_i(\cdot, G)$ equal to β_{N_i} .

DEFINITION 4. A stationary local Markov updating function $f_i : A^n \times \mathcal{G} \rightarrow A$ is *quasi-Bayesian* if for all $(\mathbf{a}, G) \in A^n \times \mathcal{G}$, we have

$$f_i(\mathbf{a}, G) \in \arg \max_{a \in A} E[u(a, \omega) | \mathbf{a}_{N_i(G)}^{\sigma^*} = \mathbf{a}_{N_i(G)}].$$

As we show subsequently, quasi-Bayesian updating is crucial for information aggregation within the general class of boundedly rational updating as defined by the three axioms. Before delving into the analysis, however, let us discuss quasi-Bayesian updating in detail and how it compares to Bayesian and DeGroot updating. To start, it is worth highlighting several key features of a quasi-Bayesian updating function. First, a quasi-Bayesian updating function treats each observed action as if the agent taking it were rational *and* best-responding only to his own private signal. Second, a quasi-Bayesian updating function changes with the environment but not with the network structure beyond the agent's neighborhood. Thus, quasi-Bayesian updating abstracts away from the complexities that arise in Bayesian updating, where it is, instead, necessary to make inferences on the private information of unobserved agents based on the actions of observed agents. Third, the nontriviality and conditional independence assumption implies that every possible action vector is consistent with the strategy profile σ^* . Thus, all possible action vectors are on the 'equilibrium path' and this unambiguously pins down the quasi-Bayesian action, but in the case of indifference. Since all agents share the same utility function, for any pair of agents i, j and a network G such that $N_i(G) = N_j(G)$, the corresponding Bayesian updating functions $f_i(\cdot, G)$ and $f_j(\cdot, G)$ coincide, unless possibly where, conditional on the observed vector $\mathbf{a}_{N_i(G)}$, both actions are optimal.

²¹Note that by vector non-indifference, the complete Bayesian observation function $\beta = \beta_{N_i=N}$ is unique. However, for $N_i \neq N$, the Bayesian observation function β_{N_i} need not be unique.

Having clarified the nature of a quasi-Bayesian updating function, it is now instructive to clarify which assumptions a boundedly rational agent applying a quasi-Bayesian updating function needs to satisfy. To start, each agent needs to know the state space Ω , the prior p , the signal generating distributions F_ω , and the utility function u . Moreover, agents must have several computational abilities. First, agents must be able to compute a Bayesian belief over the state space based on (i) their signal s_i and (ii) the first-period action vector of their neighbors, which is drawn according to $\mathbf{a}_{N_i}^{\sigma^*}$. Second, agents must be able to compute the expected utility maximizing action given their belief. We assume that agents are sufficiently rational to do both types of computations.

We next discuss how quasi-Bayesian updating relates to the two predominant approaches in the literature: Bayesian and DeGroot updating. We start by comparing the updating by a quasi-Bayesian agent with that of a Bayesian agent placed in an identical network position. At $t = 1$ and for a given signal, the action chosen by a quasi-Bayesian agent coincides with the action chosen by a myopic Bayesian agent who selects the expected utility maximizing action given his signal. In period $t = 2$ and for a given first-period action vector, the quasi-Bayesian action a_i^2 coincides with the optimal action of a Bayesian agent in environments where the first-period action reveals the agent's signal. In general, however, and already starting in period $t = 2$, the action of a Bayesian agent can deviate from that of a quasi-Bayesian agent, because a Bayesian agent selects his action as a function of his signal and the observed first-period action vector as opposed to the latter only. At any later period $t \geq 3$, quasi-Bayesian updating typically differs also from Bayesian updating, as it drastically reduces complexity by neglecting indirect inferences on the information of unobserved agents and by effectively treating every period t as if it were $t = 2$. To more precisely distinguish quasi-Bayesian updating from Bayesian updating, it is instructive to view quasi-Bayesian updating as a two-step procedure. In the first step, an agent forms a subjective probability distribution over the state space conditional on the realized action vector. In the second step, he selects an action that is expected to be utility maximizing under his subjective probability distribution. The deviation from Bayesian updating lies in the first step and the simplistic way the posterior distribution is formed. In the second step and for a given probability distribution, however, a quasi-Bayesian agent selects the same optimal action as a myopic Bayesian agent.

We now turn to comparing quasi-Bayesian updating to DeGroot updating. In the DeGroot model, the action space $\mathcal{A} \subset \mathbb{R}$ is uncountable, typically either $\mathbb{R}$ or an interval. A DeGroot updating function takes the form of a weighted average, i.e.,

$$f_i(\mathbf{a}, G) = \sum_{j \in N} w_{ij} a_j,$$

such that $w_j = 0$ for all $j \notin N_i(G)$, $w_j > 0$ for all $j \in N_i(G)$, and the sum of weights equals 1. Thus the DeGroot model imposes a particular functional form. While our analysis restricts attention to binary actions, the concept behind quasi-Bayesian updating directly extends to such general action spaces: a quasi-Bayesian updating function simply treats each observed action as if it were optimal among the given set of actions and based

only on the private signal of each respective agent. Clearly, the quasi-Bayesian function does vary with the environment. However, a DeGroot function coincides with a (quasi-)Bayesian function for a particular environment, a fact well known in the literature. Let $\Omega = \mathbb{R}$, and assume that each agent observes a private signal $s_i = \omega + \epsilon_i$, where ϵ_i is i.i.d. across agents and independent of the state. Further assume that ϵ_i is normally distributed with mean zero, and let the prior be improper. If the utility function induces the first-period action of an agent to equal his signal, then the quasi-Bayesian function is indeed a weighted average of the observed first-period actions. For more details, see [De Marzo et al. \(2003\)](#) and [Glaeser and Sunstein \(2009\)](#).

In this section we first introduced a general class of boundedly rational updating as defined by the three axioms of Markov property, stationarity, and locality, and then introduced quasi-Bayesian updating as a specific type thereof. In the following sections, we proceed in a similar fashion. First we analyze the learning properties within the general class of boundedly rational updating. The general results we establish motivate a more detailed study of the learning properties induced by quasi-Bayesian updating.

4. NECESSARY AND SUFFICIENT CONDITIONS FOR LEARNING

In this section, we consider the class of updating behavior as defined by our three axioms (Markov property, stationarity, and locality) and analyze it from the point of view of the induced information aggregation properties. We focus on *learning*, that is, the notion of information aggregation that requires all agents to eventually best-respond to the first-period action vector in every connected network (see [Definition 2](#) above). In particular, we address two questions in this section. First, we aim to understand the properties of the environment $(N, \Omega, S, p, F_\omega, u)$ that are necessary to enable learning. Second, we focus on these environments and characterize the specific subclass of updating systems that guarantees learning. The motivation behind the second approach is to provide a normative foundation for a particular type of updating behavior, i.e., one that achieves learning.

We now turn to identifying the properties the environment needs to satisfy so as to enable learning. Before stating our first result, let us define one property of the environment (and the Bayesian complete observation function), which we call Bayesian contagion. We say that a Bayesian complete observation function $\beta : A^n \rightarrow A$ satisfies *Bayesian contagion* if there exists a *contagion action* a^* that is uniquely expected utility maximizing conditional on all vectors $\mathbf{a}^1 \in A^n$, except for the vector that has consensus on the other action $-a^*$. Formally, $a^* = \beta(\mathbf{a}^1)$ for all $\mathbf{a}^1 \in A^n$ such that $a_i^1 = a^*$ for at least one $i = 1, \dots, n$.²² An environment $(N, \Omega, S, p, F_\omega, u)$ satisfies Bayesian contagion if the corresponding Bayesian complete observation function does. We apply the term “Bayesian contagion” interchangeably to the environment and the Bayesian complete observation function β . The intuitive idea behind the term is that a Bayesian agent is “infected” by the contagion action whenever at least one agent he observes is infected as well. [Theorem 1](#) identifies a necessary condition on the environment for learning to occur under stationary local Markov updating systems.

²²Environments satisfying Bayesian contagion are highly asymmetric; see the example in [Appendix B.1](#).

THEOREM 1. *Consider any stationary local Markov updating system $\mathbf{f}$ on environment $(N, \Omega, S, p, F_\omega, u)$. If $\mathbf{f}$ satisfies learning, then the environment satisfies Bayesian contagion.*

To provide some intuition for the result, let us first discuss a crucial impediment to learning that an updating system $\mathbf{f}$ needs to overcome. By definition, learning requires that eventually all agents best-respond to the first-period action vector in any connected network. However, in most networks no agent observes all others. This implies that the relevant information regarding the first-period action vector is decentralized; each agent observes some, yet not all, actions. For learning to occur, the updating functions need to transmit some information about the first-period action vector that the agents observe. However, since actions are binary, they are too coarse for any agent to encode the action vector he observed or even, say, the number of agents in his neighborhood that selected action $a = 1$. Moreover, since we focus on Markovian updating behavior, information about the first-period action vector cannot be encoded in a sequence of binary actions over some time periods.

Learning requires actions to transmit sufficient information about the locally observed first-period action vector so that eventually consensus on the optimal action given $\mathbf{a}^1$ emerges. The defining feature of a Bayesian contagion environment is that to identify the optimal action, it is enough for the updating system to establish whether at least one agent selected the contagion action in the first period. Thus, the necessary information for learning can be encoded in binary actions, as in “Yes, the contagion action was observed” and “No, the contagion action was not (yet) observed.” [Theorem 1](#) states that Bayesian contagion environments are the only ones where learning may occur.

The proof of [Theorem 1](#) builds upon an auxiliary step. For the auxiliary step, we assume that an updating system $\mathbf{f}$ satisfies learning and show that this has direct implications on the functional form of the individual updating functions f_i . The proof of [Theorem 1](#) then derives a connection between the functional properties of the individual updating functions f_i and the environment as represented by the Bayesian complete observation function β . Indeed the environment needs to satisfy Bayesian contagion for learning to hold.

Let us first provide some details on the auxiliary step that establishes the functional properties.²³ The first property of the updating functions that follows from learning is *unanimity*, i.e., if agreement on one action is observed, then that action continues to be selected. To see why, note that, by vector non-indifference, learning requires that eventually all agents agree on the same action. Consider an action a^* that is uniquely optimal conditional on some first-period action vector. Then learning immediately implies that the agreement vector $\mathbf{a}^*$ is a fixed point of $\mathbf{f}(\cdot, G)$. From this it follows that if agreement on a^* holds in agent i 's neighborhood, then his updated action is a^* . This applies to every f_i and every possible neighborhood $N_i(G)$.

The remaining properties of the updating functions that directly derive from learning concern only neighborhoods of size 2. We show that learning implies that the updating functions f_i of all agents i satisfy anonymity and a common order property whenever

²³The functional properties are stated as lemmas and are proved in the [Appendix](#).

$|N_i(G)| = 2$. We call f_i *anonymous* if, for every $(\mathbf{a}, G)$, permutations of $\mathbf{a}_{N_i(G)}$ induce the same updated action,²⁴ i.e., for $\mathbf{a}, \mathbf{a}'$ such that $a_j = a'_j$ for all $j \notin N_i(G)$ and

$$\sum_{j \in N_i(G)} a_j = \sum_{j \in N_i(G)} a'_j,$$

we have $f_i(\mathbf{a}, G) = f_i(\mathbf{a}', G)$. Anonymity implies that an agent's updated action depends only on the distribution of observed actions, and not also on the identity of agents taking the observed actions. The second property, *common order*, holds if there exists an action a^* such that for every f_i and for all neighborhoods of size 2, $|N_i(G)| = 2$, observed disagreement induces agent i to update to a^* .

We can now present the core argument of the proof of [Theorem 1](#). Consider a star network and assume that learning holds. The argument above states that the updating function of each branch agent satisfies unanimity, anonymity, and the common order property. This implies that there exists an action a^* such that *every* branch agent updates to action a^* if and only if he observes action a^* at least once. Now consider a first-period action vector $\mathbf{a}' \in A^n$, where exactly $k \in \{1, \dots, n-1\}$ branch agents select action a^* . Thus, each branch agent who selected a^* in the first period continues to select a^* in *every* later period. Hence, learning requires that, for every such first-period action vector $\mathbf{a}'$, action a^* is optimal. By definition such an environment satisfies Bayesian contagion.

[Theorem 1](#) establishes Bayesian contagion as a necessary condition on the environment for learning to hold. We now focus on Bayesian contagion environments and turn to the analysis of the updating systems. First note that there exist updating systems that achieve learning in Bayesian contagion environments. For a simple example, consider a setting with three agents, $i = \alpha, \beta, \gamma$. Suppose that action 1 is the contagion action. Assume that each f_i updates to action 1 if and only if action 1 is observed at least once. It is easy to see that learning holds for this updating system.

Our next goal is to identify the properties of the updating system that are necessary for learning. To do so, we first restrict attention to a subclass of local Markov updating systems: those that satisfy anonymity for each f_i . As defined above, anonymous updating functions treat all observed actions equally, independently of the identity of the agent selecting them. While clearly restricting the class of updating systems, anonymity is somewhat natural in a setting where (i) agents have no information of the network beyond their neighbors (inherent in the local property), (ii) agents do not keep track of their private signals over time, and (iii) private signals are conditionally i.i.d. We next present our main result, which identifies a crucial property of updating systems that is necessary for learning.

THEOREM 2. *If an anonymous stationary local Markov updating system $\mathbf{f}$ satisfies learning, then $\mathbf{f}$ is quasi-Bayesian.*

²⁴Recall that for a set of agents $M \subset N$ and a vector $\mathbf{a} \in A^n$, $\mathbf{a}_M \in A^m$ denotes the subvector of actions of agents in M , i.e., $[\mathbf{a}_M]_i = a_i$ for all $i \in M$.

Theorem 2 states that quasi-Bayesian updating is necessary for learning within the subclass of updating behavior considered.²⁵ As such, **Theorem 2** provides a normative foundation for the analysis of quasi-Bayesian updating systems.

We now provide a brief outline of the proof and some intuition for its core argument. Consider an anonymous stationary local Markov updating system $\mathbf{f}$ that satisfies learning. By **Theorem 1**, the underlying environment satisfies Bayesian contagion. Further, as we pointed out in the proof outline of **Theorem 1**, each updating function f_i satisfies unanimity and common order for neighborhoods of size of 2. The proof uses these properties to show that if one agent's updating function is not quasi-Bayesian for some disagreement action vector $\mathbf{a}'$ and some neighborhood, then one can construct a neighborhood-corresponding network where $\mathbf{a}'$ is a fixed point of $\mathbf{f}$. Since learning requires all agents eventually to reach agreement, this contradicts learning.

We describe the argument for the case where an agent deviates from the quasi-Bayesian updating function when observing all other agents. Note that for an agent i who observes all other agents, $N_i(G) = N$, a quasi-Bayesian updating function f_i selects the contagion action a^* if and only if it occurs at least once in $N_i(G)$. Consider a star network G' with agent i in the center and a disagreement action vector $\mathbf{a}' \in A^n$, where agent i selects the noncontagion action $a'_i = \neg a^*$. Since the environment satisfies Bayesian contagion, learning requires that all agents eventually select the contagion action a^* . Assume that $f_i(\mathbf{a}', G')$ deviates from the quasi-Bayesian function, i.e., $f_i(\mathbf{a}', G') = \neg a'$. To see that for such an f_i the vector $\mathbf{a}'$ is a fixed point of $\mathbf{f}$, consider a branch agent j who observes agent i and himself. If the branch agent initially selected $a'_j = \neg a^*$, he remains at $\neg a^*$ in the second period, as his updating function satisfies unanimity. Instead, if the branch agent initially selected a^* , then by the common order property and given that a^* is the contagion action, he remains at a^* in the second period. Hence, $\mathbf{a}'$ is a fixed point of $\mathbf{f}(\cdot, G')$ and learning fails.

To conclude the proof, we need to show that f_i is quasi-Bayesian for all neighborhoods N_i , not only the complete observation case $N_i = N$. This is done in two steps. First we show that in a Bayesian contagion environment, the quasi-Bayesian function satisfies contagion for all neighborhoods.²⁶ We say that an updating function f_i satisfies contagion if it selects a^* whenever at least one agent in the neighborhood N_i selects a^* , for any neighborhood N_i . The second step then shows that f_i being quasi-Bayesian for any neighborhood N_i is necessary for learning. This is done similarly as in the argument for the complete observation case, $N_i(G) = N$. For each neighborhood, one simply constructs a corresponding network and first-period action vector such that f_i failing to satisfy contagion would induce a disagreement fixed point and failure of learning as a result.

Theorem 2 left unanswered the question as to whether non-anonymous and, hence, non-quasi-Bayesian, updating systems may achieve learning in Bayesian contagion environments. For a simple example of such an updating system that indeed achieves

²⁵Note that conditional i.i.d. signals and the non-indifference assumption imply that quasi-Bayesian updating functions satisfy anonymity.

²⁶See **Lemma 5** in the **Appendix**.

learning, consider a setting with three agents, $i = \alpha, \beta, \gamma$. Let f_α and f_γ be quasi-Bayesian. Suppose that action 1 is the contagion action. Let agent β have an updating function f_β that is not quasi-Bayesian. In particular, assume that f_β coincides with the quasi-Bayesian function everywhere on the domain except for the case $N_\beta = \{\alpha, \beta, \gamma\}$, $a_\beta = 1$ and $a_\beta = a_\gamma = 0$. It is easy to see that learning holds. While f_β is not strictly quasi-Bayesian, it does coincide with quasi-Bayesian updating on parts of the domain. As the following result shows, this is a necessary property of any stationary local Markov updating system that achieves learning.

PROPOSITION 1. *Let the environment satisfy Bayesian contagion and let a^* be the contagion action. If a stationary local Markov updating system $\mathbf{f}$ satisfies learning, then for every agent i , $f_i(\mathbf{a}, G)$ is quasi-Bayesian for all $(\mathbf{a}, G)$ with $a_i \neq a^*$.*

Proposition 1 states that *any* stationary local Markov updating system has to coincide with the quasi-Bayesian updating system on a subset of the domain. More precisely, deviations of f_i from the quasi-Bayesian function can only occur if a_i equals the contagion action. The proof of **Proposition 1** follows the same line of reasoning as the proof of **Theorem 2** and is therefore omitted.

Theorem 1 established a necessary condition on the environment for learning to hold. **Theorem 2** and **Proposition 1** have shown that the quasi-Bayesian property plays a crucial role in achieving learning in Bayesian contagion environments. The following corollary, which follows from **Theorem 1** and the proof of **Theorem 2**, integrates these results and provides a necessary and sufficient condition for learning.

COROLLARY 1. *Consider a quasi-Bayesian updating system $\mathbf{f}$. Learning holds if and only if the environment satisfies Bayesian contagion.*

Let us provide some insight into **Corollary 1**. Since a quasi-Bayesian updating system satisfies the Markov, stationarity, and locality axioms, **Theorem 1** immediately implies that the environment needs to satisfy Bayesian contagion for learning to hold. Now consider an environment that satisfies Bayesian contagion. As we argued in the outline of the proof of **Theorem 2**, the quasi-Bayesian updating functions thus satisfy contagion.²⁷ That is, for any neighborhood, every agent updates to the contagion action if and only if the contagion action is observed at least once. It then follows that whenever at least one agent selected the contagion action in the first period, this action spreads through the network along the shortest path. Thus, the period duration to agreement on the optimal action is bounded by the diameter of the network plus 1.

5. STRUCTURAL CONDITIONS FOR LEARNING UNDER QUASI-BAYESIAN UPDATING

Corollary 1 above showed that a Bayesian contagion environment is necessary and sufficient for quasi-Bayesian updating to achieve learning, i.e., optimal information aggregation in *all* connected networks. But Bayesian contagion is a strong requirement and

²⁷See **Lemma 5** in the **Appendix**.

typically will not be satisfied. Hence, it is important to understand the information aggregation properties of quasi-Bayesian updating when the Bayesian contagion property does not hold. Does optimal information aggregation hold in most networks or only in some? How does the structure of a network interplay with the environment in determining the success or failure of learning in that network?

We now turn to answering these questions. The goal is to provide necessary conditions on a network G for quasi-Bayesian updating to achieve *learning in network G* ; see [Definition 3](#). We consider environments $(N, \Omega, S, p, F_\omega, u)$ that do not satisfy Bayesian contagion, as we have shown that instead, under Bayesian contagion, learning in every networks occurs. For tractability, we assume that the environment satisfies a monotonicity condition. Let $\mathbf{a}_k^{\sigma^*}$ denote the random action vector where each component a_i is drawn according to a^{σ^*} . To define monotonicity, consider the Bayesian observation mapping $\hat{\beta}_k$ that assigns to every first-period an action vector $\mathbf{a}^1 \in A^k$, which is drawn according to $\mathbf{a}_k^{\sigma^*}$, the corresponding expected utility maximizing action.²⁸ Formally, $\hat{\beta}_k : A^k \rightrightarrows A$ such that

$$\hat{\beta}_k(\mathbf{a}) = \arg \max_{a \in A} E[u(a, \omega) | \mathbf{a}_k^{\sigma^*} = \mathbf{a}^1].$$

An environment satisfies *monotonicity* if, for any two vectors $\mathbf{b} \in A^k$ and $\mathbf{c} \in A^l$ such that $\hat{\beta}_k(\mathbf{b}) \cap \hat{\beta}_l(\mathbf{c}) \neq \emptyset$, we have $\hat{\beta}_{k+l}(\mathbf{d}) = \hat{\beta}_k(\mathbf{b}) \cap \hat{\beta}_l(\mathbf{c})$ and where $\mathbf{d} = (\mathbf{b}, \mathbf{c}) \in A^{k+l}$. To provide some intuition, note that the action vectors $\mathbf{b}$ and $\mathbf{c}$ can be interpreted as conditional i.i.d. signals. Thus, monotonicity requires that if action a^* is optimal conditional on observing either signal (action vector) $\mathbf{b}$ or $\mathbf{c}$ by itself, then a^* is optimal when observing both signals $\mathbf{b}$ and $\mathbf{c}$ together. One particular example of an environment that satisfies monotonicity is the standard social learning model with binary states and binary actions, where agents achieve a utility of 1 if the action matches the state and 0 otherwise.

We now turn to the analysis of the structural properties of the network that are necessary for learning. It is easy to see that learning in network G holds for complete networks where every agent observes all others. In the complete network, the quasi-Bayesian updating function of each agent coincides with the Bayesian complete observation function β . Thus, in the second period all agents select the optimal action and this consensus vector is a fixed point by monotonicity. Typically, however, learning fails in most networks G . The following theorem bounds the diameter of the network that is necessary for learning in a network to occur.²⁹

THEOREM 3. *Consider a monotone environment where Bayesian contagion fails. If G has a diameter larger than 4, then any quasi-Bayesian updating system $\mathbf{f}$ fails learning in network G .*

²⁸For a group of k agents and the vector of their first-period actions $\mathbf{a}^1 \in A^k$, the Bayesian observation mapping $\hat{\beta}_k$ assigns the expected utility maximizing action conditional on $\mathbf{a}^1$.

²⁹Note that a network of *diameter* k has the property that the length of the shortest path connecting any pair of agents is smaller than or equal to k , and is equal to k for at least one pair of agents.

Theorem 3 states that if learning in network G holds, then its diameter is smaller than or equal to 4. Thus, for quasi-Bayesian updating to induce learning in a given network, the network needs to be sufficiently dense so as to assure a diameter not larger than 4. Since Bayesian contagion is a strong property on the environment and a diameter smaller than or equal to 4 is a strong condition for large networks, quasi-Bayesian updating typically fails to achieve learning. This stands in contrast to the naïve learning result of [Golub and Jackson \(2010\)](#) for DeGroot updating in a rich action environment. They show that as the network size grows to infinity, the asymptotic consensus action accurately aggregates all private information under a weak condition on the sequence of networks (more precisely weight matrices), i.e., that the influence of the most influential agent vanishes to zero.³⁰

Before delving into details on the proof of **Theorem 3**, it is useful to explain informally the role of the monotonicity assumption. An implication of the monotonicity assumption is that the quasi-Bayesian updating function follows a threshold rule. To clarify, consider an action vector $\mathbf{a} \in A^n$ and an agent i . There exists a threshold proportion q_a such that if the proportion of agents in $N_i(G)$ who select action a in vector $\mathbf{a}$ is greater than or equal to q_a , then f_i updates to action a .³¹

To prove **Theorem 3**, we rely on two properties of the quasi-Bayesian updating system that we show to be necessary for learning. The first property, which we call information retention, requires that if action a is optimal conditional on the first-period action vector $\mathbf{a}^1$, then action a is also optimal conditional on all subsequent action vectors $\mathbf{a}^t$. Formally, *information retention* requires that for all initial action vectors $\mathbf{a} \in A^n$, we have that $\beta(\mathbf{a}) = b$ implies $\beta(f(\mathbf{a}, G)) = b$. The second property, information diffusion, is necessary for learning in environments that fail Bayesian contagion. *Information diffusion* requires that no agent, for any neighborhood, ever updates to an action that occurs only once in his neighborhood. Otherwise, one can construct a network and a first-period action vector where all agents but agent i select action a' , but agent i remains at action $\hat{a}$ in all periods. This contradicts learning in the given network, as by assumption the environment does not satisfy Bayesian contagion.

Having established the necessity of these two properties for learning in network G , the proof of **Theorem 3** shows that for any network G of diameter larger than 4, there exists a first-period action vector such that information retention fails. This can be seen when combining the concept of information retention with that of information diffusion. The key insight of the proof is that for any agent i who selects the optimal action a^* in the first period, there needs to be another agent, at a distance of at most 2, who also selects action a^* in the first period. If not, then by information diffusion, all agents in i 's neighborhood, including agent i , select the suboptimal action $\neg a^*$ in the second period. For a first-period action vector that features the minimal number of a^* actions that indeed make a^* optimal, this leads to failure of information retention and, hence, failure of learning in network G . Based on this argument, the proof of **Theorem 3** shows that the diameter of a network G cannot exceed 4 for learning in network G to be satisfied.

³⁰They measure influence by an agent's weight in the stationary distribution of the weight matrix.

³¹For the formal argument, please see [Lemma 6](#) and [Lemma 7](#) and their proof in the [Appendix](#).

We conjecture that a strictly smaller diameter than 4 is necessary for learning, but we leave a strengthening of [Theorem 3](#) for future work. While the exact condition on the diameter is unclear, it is worth noting that some networks with diameter 2 can achieve learning. In the [Appendix](#), we construct a network G of diameter 2 where learning in G holds despite the fact that Bayesian contagion fails.³²

We next turn to another structural property of the network that is necessary for learning. So as to present our next result, we borrow a network property introduced by [Morris \(2000\)](#). Let G be a network on a set of agents N . For a group of agents $M \subset N$, consider for each agent $i \in M$ the proportion of his neighbors that belong to M . Picking the smallest of such proportions among the members of M gives the so-called *group cohesion* of group M :

$$\pi_G(M) = \min_{i \in M} \frac{|N_i(G) \cap M|}{|N_i(G)|}.$$

According to this definition, a group where each member has a large proportion of his neighbors inside the group is highly cohesive.

PROPOSITION 2. *Consider a monotone environment and a quasi-Bayesian updating system $\mathbf{f}$. If in network G there exists a group of agents $M \subset N$ and an action $a \in A$ such that the group cohesion (i) $\pi_G(M) \geq q_a$ and (ii) $|M| < q_a n$, then learning in network G fails.*

[Proposition 2](#) states that learning fails in networks that contain a small highly cohesive group. The underlying reason for failure of learning is that such a network cannot satisfy information diffusion. To see this, note that if all group members initially select the same action, then they continue to select this action in all subsequent periods, independently of the actions of agents outside of the group. If the group consensus action is not optimal conditional on the realized first-period action vector, then learning fails.

This section focused on providing conditions on the network structure that are necessary for quasi-Bayesian updating to lead to learning in a given network. As we have seen, these necessary conditions are fairly strong. These results give rise to the question of how the process of action vectors evolves when the necessary conditions are not satisfied. As explained above, monotonicity implies that quasi-Bayesian updating comes down to the application of a threshold function. This property of quasi-Bayesian updating implies that we can apply a result of [Goles and Olivos \(1980\)](#) on the iterated application of threshold functions with binary range to understand the long-run behavior of the process of actions. Their result applied to our setting implies that the process of action vectors terminates either in a fixed point or a cycle of length 2. This fixed point might be agreement on the optimal action, as it is the case in a complete network for general environments. However, as shown in [Example B.3](#) in the [Appendix](#), in some environments and networks this terminal fixed point can also be the suboptimal action. That is, quasi-Bayesian updating can in certain circumstances lead all agents to agree on the suboptimal action.

³²See [Example B.2](#) in the [Appendix](#).

6. EXTENSION: WEAKLY LOCAL AND NONSTATIONARY UPDATING SYSTEMS

The key take-away from the analysis so far is that learning is very hard to achieve. Our analysis was done under the assumption that the updating behavior of all agents satisfies the Markov property, stationarity, and locality. We now analyze whether learning outcomes can be improved when the locality and stationarity axioms are relaxed. Let us turn to our definition of locality first. Recall that locality, as defined in [Section 3](#), does not allow updating to take into account any information regarding the network structure other than the identities of neighbors. Thus, the updating function cannot condition on the relative network position of neighbors. Instead, let us introduce the following weak form of locality.

DEFINITION 5. A Markov updating function $f_i : A^n \times \mathcal{G} \rightarrow A$ is *weakly local* if, for all pairs $(\mathbf{a}, G)$, $(\mathbf{a}', G)$ such that $\mathbf{a}_{N_i(G)} = \mathbf{a}'_{N_i(G)}$, we have $f_i(\mathbf{a}, G) = f_i(\mathbf{a}', G)$.

Weak locality simply requires that the updated action is invariant in the actions of nonneighbors. A weakly local function allows for observed agents to be treated differentially depending on their position in the network, and additionally allows for the updated action to be invariant in the actions of some or even all neighbors, effectively allowing agents to ignore some or all neighbors.

Consider a binary action environment $(N, \Omega, S, p, F_\omega, u)$ that is monotone as defined in [Section 5](#). As [Theorem 1](#) showed, for a stationary local Markov updating system to satisfy learning (in all networks), the environment needs to satisfy Bayesian contagion. What if we assume weak locality instead of locality? Does the ability to capture the network position of neighbors allow for better learning outcomes? More precisely, are there monotone non-Bayesian contagion environments where stationary weakly local Markov updating systems achieve learning? The following result provides the answer.

THEOREM 4. *Consider a monotone environment. If a stationary weakly local Markov updating system $\mathbf{f}$ satisfies learning, then the environment satisfies Bayesian contagion.*

Thus, in monotone environments, knowledge of the whole network generally does not improve learning outcomes. In fact, even in very simple network structures, such network knowledge does not allow agents to achieve learning. The proof of [Theorem 4](#) considers one such network: the star. The advantage of the star network is that one agent—the center agent—observes all other agents and can identify the optimal action at the end of period $t = 1$. Intuition might suggest that learning outcomes should improve when each branch agent observes the center agent, knowing that he is the center agent. We establish that this is not the case and that this is due essentially to the stationarity of $\mathbf{f}$.

Let us briefly summarize the core argument of the proof of [Theorem 4](#). Consider a star network and an environment that fails Bayesian contagion. A crucial step in the proof is to show that learning implies that every branch agent updates to the previous-period action of the center agent. Consider a first-period disagreement vector and assume that the center agent selected the suboptimal action $-a^*$. In the second period,

thus all branch agents select $\neg a^*$. In contrast, the center agent must select the optimal action a^* in the second period. To see this, note that if the center agent were to select $\neg a^*$ as well, then a fixed point on the suboptimal action would be reached, contradicting learning. Since information retention needs to hold, the action a^* is optimal conditional on the second-period action vector. Since the environment is monotone, the complete Bayesian observation function satisfies a threshold rule, which in this case corresponds to action a^* being optimal if at least one agent selects action a^* . By definition, such an environment satisfies Bayesian contagion, establishing a contradiction.

Thus, relaxing locality alone generally does not improve learning outcomes, not even in a simple network such as the star. Nevertheless, when relaxing both stationarity and locality, there exists an updating system that induces learning in the star network, even when Bayesian contagion fails. Let the updating function f_j^t of the center agent be quasi-Bayesian in all periods t . For the branch agents i , let $f_i^2(\mathbf{a}) = a_i$ for all agents $i \neq j$, and let $f_i^t(\mathbf{a}) = a_1$ for all $t > 2$ and $i \neq 1$. In this updating system, the center agent selects the optimal action in period $t = 2$. Since the branch agents are simply copying the center agent's action from period $t = 2$, all agents select the optimal action in period $t = 3$.

We now turn to relaxing the stationarity assumption. As we show, nonstationary weakly local Markov updating systems achieve learning in much more general network structures than the star network. Let us now consider a nonstationary Markov updating system that is described by a sequence of mappings $(\mathbf{f}^t)_{t \in \mathbb{N}}$, where each mapping $\mathbf{f}^t : A^n \rightarrow A^n$ is *weakly local*. We say that $M \subset N$ forms a *complete subgroup* if M is a complete subgraph, i.e., if each agent in M is a neighbor of all other agents in M . The following theorem presents a sufficient condition for learning in network G to hold.

THEOREM 5. *If G contains a complete subgroup M^* of size greater than or equal to $\frac{\ln n}{\ln 2} + 1$, then there exists a nonstationary weakly local Markov updating system $(\mathbf{f}^t)_{t \in \mathbb{N}}$ such that learning in network G holds.*

By [Theorem 5](#), the existence of a possibly extremely small complete subgroup of agents is sufficient for learning in a given network. For example, in a network of one billion agents, a complete subgroup of 31 agents is sufficient to assure learning for a carefully constructed nonstationary weakly local Markov updating system. This result differs strikingly from [Theorem 3](#), which showed that under quasi-Bayesian updating, learning fails in any network with a diameter larger than 4. However, there are several caveats in regard to the practicality of such an updating system.

Let us provide the core intuition behind [Theorem 5](#) and clarify our statement regarding the practicality. For learning to hold in a given network, agents need to transmit some information about the first-period action vector. At the level of an individual agent, however, this fails, as binary actions are too coarse. Now the nonstationarity of the updating system allows the updating functions to change over time. This in turn allows a core group of agents to jointly encode the relevant information about the first-period action vector in their group-action vector over time. Essentially, a weakly local updating system that induces learning assigns to each agent the role of either an information transmitter or an information retainer. Roughly, the members of the complete

subgraph, which subsequently we denote as the core, retain information while all other agents transmit information regarding the first-period action vector toward the core. Due to the size of the core, their joint action vector can retain all relevant information regarding the distribution of first-period actions. Thus, the process of actions is divided into two phases. In the first phase, information is passed on to and retained by the core. If the core has the minimal critical size, in each period the information regarding the first-period action of only one additional agent is retained. As a result, the first phase is possibly very long. In the second phase, once all information is aggregated, all core agents switch to the optimal action and thereafter the optimal action spreads to all other agents along the shortest path.

We acknowledge that the use of a nonstationary weakly local Markov updating system has two main drawbacks in terms of practicality. First, as highlighted in the outline of the proof, learning requires all agents to forego their utility maximization objective for a potentially very large number of periods. Second, a successful updating system requires maximal coordination among all agents in the network, as each agent needs to know in which period he is required to pass on the first-period action of which neighbor.

7. CONCLUSION

The rise of the internet and online social networks has increased the number of social connections as well as the ease with which behavior, opinions, beliefs, etc. can be observed and shared. A good starting point to understand the resulting aggregate social dynamics is to consider the microlevel, i.e., the individual updating behavior. Complementary to the large literature that considers Bayesian updating, we provide a general analysis of information aggregation in networks with boundedly rational agents. We focus on binary action environments and updating behavior that satisfies the Markov property, stationarity, and locality. We show that learning in all network structures occurs only in rare environments. Within such rare environments, we establish that quasi-Bayesian updating is necessary to achieve learning and, thus, provide a normative foundation for quasi-Bayesian updating. Outside of these rare environments, learning fails in some networks. Focusing on quasi-Bayesian updating, we establish structural properties of the network that are necessary to achieve learning in a given network.

We then consider a richer class of updating heuristics, relaxing the locality axiom and, thus, allowing agents to fine-tune their updating behavior toward their own network position, the network position of their neighbors, and the structure of the network as a whole. We show that, nevertheless, learning fails in very simple networks such as the star, and that learning in all network structures holds in exactly the same rare environments as under locality. Finally, we show that when relaxing both stationarity and locality, there exist boundedly rational updating systems that induce learning for a large set of connected networks. However, the degree of coordination required is extremely high and the time required to reach consensus on the optimal action is potentially very long.

To conclude, we highlight two natural ways to extend our formal analysis. First, one can move away from binary action environments and consider instead rich (i.e., uncountable) action environments, which are indeed the focus of the existing literature on

boundedly rational updating.³³ In Section 3, we highlighted that the concept of quasi-Bayesian updating can directly be extended to a general action space. For Bayesian updating, it is well established that a richer action space enables better inferences on the underlying information held by agents and, hence, leads to better learning outcomes.³⁴ A similar positive result holds for quasi-Bayesian updating if signals are finite and the action space is compact metrizable and perfect. More precisely, for generic³⁵ pairs of continuous utility functions and joint probability measures,³⁶ there exists a quasi-Bayesian updating function that satisfies learning. Moreover, learning occurs along the shortest path in any strongly connected network. More details and a formal proof of this statement can be found in an earlier working paper version, Mueller-Frank and Neri (2015). Extending our analysis in a second direction, one may consider random as opposed to deterministic updating functions. A followup paper, Arieli et al. (2020) considers random updating functions and establishes strong positive learning results.

Finally, this paper provides a theoretical foundation for the analysis of quasi-Bayesian updating. While experimental evidence for quasi-Bayesian updating in simple environments (see, for example, Chandrasekhar et al. 2020) has been documented, an interesting direction for future research would be to investigate the extent to which quasi-Bayesian updating describes the behavior of participants in more complex experimental settings.

APPENDIX A: PROOFS

Note that locality of the updating function $f_i : A^v \times \mathcal{G} \rightarrow A$ implies that f_i is identified by a *reduced function* $\hat{f}_i : A^n \times 2^N \rightarrow A$, where $f_i(\mathbf{a}, G) = \hat{f}_i(\mathbf{a}, N_i(G))$.³⁷ The proofs mostly work directly with the reduced function $\hat{f}_i$ rather than f_i . We say that $\hat{f}_i$ satisfies *unanimity* if, for every $M \subset N$ such that $i \in M$ and $\mathbf{a} \in A^n$ such that $a_j = a$ for all $j \in M$, we have $\hat{f}_i(\mathbf{a}, M) = a$.

A.1 Auxiliary results for Theorem 1

Prior to introducing the auxiliary results, we define one additional property of the Bayesian complete observation function $\beta : A^n \rightarrow A$. We call β *dominant* if there exists a *dominant* action a^* such that $\beta(\mathbf{a}) = a^*$ for all $\mathbf{a} \in A^n$.

LEMMA 1. *Consider an environment where the Bayesian complete observation function $\beta : A^n \rightarrow A$ satisfies dominance. Then for every $M \subset N$, $|M| = m$, the quasi-Bayesian function $\beta_M : A^m \rightarrow A$ satisfies dominance.*

³³For examples, see De Marzo et al. (2003), Banerjee et al. (2019), Golub and Jackson (2010), and Molavi et al. (2018).

³⁴See, for example, Lee (1993) and Arieli and Mueller-Frank (2017).

³⁵We use the standard notion of topological genericity, i.e., a *residual set*: a set that can be represented by a countable intersection of open dense sets.

³⁶By joint probability measure, we refer to the probability measure over the product of state and signal spaces, $p^* \in \Delta(\Omega \times S^n)$.

³⁷The range of 2^N is limited to sets that contain i .

PROOF. Since signals are conditional i.i.d., note that if β_M satisfies dominance for M for some $M \subset N$, then it also satisfies dominance for all other groups M'' that have the same cardinality as M . Without loss of generality (WLOG), let 1 be dominant for observation set N and suppose that $M \subset N$ is the largest group such that 1 fails dominance, i.e., there exists $\mathbf{a} \in A^m$ such that $\beta_M(\mathbf{a}) \neq 1$. For example, let the observed vector $\mathbf{a}' \in A^m$ induce 0 as the optimal action:

$$E[u(1, \omega)|\mathbf{a}'] \leq E[u(0, \omega)|\mathbf{a}'].$$

As $m+1$ has dominance in 1, we have

$$E[u(1, \omega)|(\mathbf{a}', a)] > E[u(0, \omega)|(\mathbf{a}', a)]$$

for all $a \in \{0, 1\}$, establishing a contradiction via the law of total expectation. $\square$

For the remaining lemmas, we assume that the environment satisfies nontriviality, non-indifference, and vector non-indifference. Note that the nontriviality assumption on the environment implies that for $n = 1$, the Bayesian complete observation function $\beta : A \rightarrow A$ is not dominant. This argument is employed in the following lemma.

LEMMA 2. *Learning of $\mathbf{f}$ implies that $\hat{f}_i : A^n \times 2^N \rightarrow A$ satisfies unanimity for all $i \in N$.*

PROOF. Consider the Bayesian complete observation function β . As the environment satisfies vector non-indifference, learning implies that for every initial vector $\mathbf{a}^1 \in A^n$, consensus on $\beta(\mathbf{a}^1)$ emerges within finitely many rounds. The proof is established in two steps.

Step 1. We first show that for each action $a \in A$, there exists at least one corresponding vector $\mathbf{a}' \in A^n$ such that $a = \beta(\mathbf{a}')$. Suppose not, i.e., assume that β satisfies dominance. By Lemma 1, it follows that β_M satisfies dominance as well for all $M \subset N$. However, nontriviality implies nondominance for $|M| = 1$, establishing a contradiction. Therefore, for each action $a \in A$, there exists at least one corresponding vector $\mathbf{a}' \in A^n$ such that $a = \beta(\mathbf{a}')$.

Step 2. Learning implies the emergence of consensus on $\beta(\mathbf{a}^1)$. As either action might be optimal, this requires that, for any network $G \in \mathcal{G}$, the respective consensus vectors are fixed points of $\mathbf{f}$, which implies unanimity of $\hat{f}_i(\mathbf{a}, M)$ for all $M \subset N$ and for all $i \in N$. $\square$

LEMMA 3. *Let $\mathbf{f}$ be a local Markov updating system that induces learning. Then $f_i(\mathbf{a}, G)$ is anonymous whenever $|N_i(G)| = 2$ for all $i \in N$.*

PROOF. Let $N_i(G) = 2$. By Lemma 2, learning of $\mathbf{f}$ implies unanimity of $\hat{f}_i$. Abusing notation slightly, assume that $\hat{f}_i(0, 1) \neq \hat{f}_i(1, 0)$. Without loss of generality, let $\hat{f}_i(0, 1) = \hat{f}_i(0, 0) = 0$ and $\hat{f}_i(1, 0) = \hat{f}_i(1, 1) = 1$. If so, then the updated action depends only on the first component and, thus, one agent in the observation set is never pivotal for the updated action. This contradicts our definition of locality. $\square$

As defined in the main text, we say that the updating system $\mathbf{f}$ satisfies *common order* if there exists an action a^* such that for every i and every network G with $|N_i(G)| = 2$, we have $\hat{f}_i(a, N_i(G)) = a^*$ if and only if $a_j = a^*$ for at least one $j \in N_i$.

LEMMA 4. *Let $\mathbf{f}$ be a local Markov updating system that induces learning. Then $\mathbf{f}$ satisfies common order.*

PROOF. By Lemma 3, learning implies that for each $i \in N$, f_i is anonymous for any $|N_i| = 2$. Assume to the contrary that the statement of Lemma 4 is incorrect. Then there exist two agents $j \neq l$, and matching neighbors j_1 and l_1 such that f_j assigns a' under disagreement among j and j_1 , while f_l assigns $a'' \neq a'$ in case of disagreement among l and l_1 . Consider a nonconsensus vector $\mathbf{a}'$ such that a' is uniquely optimal conditional on $\mathbf{a}'$, $\beta(\mathbf{a}') = a'$. Let G_l be a network with $N_l = \{l, l_1\}$ and let $a'_l = a''$. Since f_l assigns a'' in case of disagreement, agent l remains at a'' in all periods $t \geq 1$, contradicting learning. $\square$

A.2 Proof of Theorem 1

Let a^* denote the action that is preferred under the common order for neighborhoods of size 2 when observing disagreement (see Lemma 4). Consider a star network G' with agent i in the center and consider any first-period disagreement vector $\mathbf{a}^1 \in A^n$. There exists at least one branch agent $j \neq i$ with $N_j(G) = \{i, j\}$ who observes a^* . By unanimity (Lemma 2) and the common order property (Lemma 4), any such agent j continues to select a^* in all periods $t \geq 2$. This implies that in the star network, the agents fail to reach agreement in the action $\neg a^*$ for every first-period disagreement action vector $\mathbf{a}^1$. However, learning under vector non-indifference implies convergence to agreement. Thus, we need to have that $\beta(\mathbf{a}^1) = a^*$. Since $\mathbf{a}^1$ was chosen as an arbitrary disagreement vector, learning implies that β satisfies Bayesian contagion.

A.3 Auxiliary result for Theorem 2

For a subgroup of agents $M \subset N$, $|M| = m$, a Bayesian observation function β_M assigns to every first-period action vector $\mathbf{a}^1 \in A^m$ a corresponding expected utility maximizing action. Let $\mathbf{a}_M^{\sigma^*}$ denote the random vector with realizations in A^m , where the action of each agent i is drawn according to $a_i^{\sigma^*}$. Formally, a Bayesian observation function $\beta_M : A^m \rightarrow A$ satisfies

$$\beta_M(\mathbf{a}) \in \arg \max_{a \in A} E[u(a, \omega) | \mathbf{a}_M^{\sigma^*} = \mathbf{a}]$$

for all $\mathbf{a} \in A^m$. To establish that f_i is quasi-Bayesian for all neighborhoods, we rely on the following auxiliary result.

LEMMA 5. *If the Bayesian observation function $\beta : A^n \rightarrow A$ satisfies Bayesian contagion, then for every $M \subset N$, $|M| = m$, every Bayesian observation function $\beta_M : A^m \rightarrow A$ satisfies Bayesian contagion.*

PROOF. By assumption, $\beta_N(\mathbf{a}) = \beta(\mathbf{a})$ satisfies Bayesian contagion, and WLOG, let 1 be the contagion action. Suppose that for some $M \subset N$, the Bayesian function fails to satisfy contagion. Failure of contagion of β_M implies that one or both of the following two properties have to be satisfied by β_M : Either there exists a vector $\mathbf{a}' \in A^m$ that does not satisfy consensus on 0 such that

$$E[u(1, \omega)|\mathbf{a}'] \leq E[u(0, \omega)|\mathbf{a}']$$

or the vector $\mathbf{0} \in A^m$ that features consensus on 0 satisfies

$$E[u(1, \omega)|\mathbf{0}] \geq E[u(0, \omega)|\mathbf{0}].$$

Let us consider the former case first. Since $\mathbf{a}'$ is not an agreement vector, contagion at size $m + 1$ implies

$$E[u(0, \omega)|(\mathbf{a}', a)] < E[u(1, \omega)|(\mathbf{a}', a)]$$

for all $a \in A$, contradicting the law of total expectation. Suppose the latter case holds. Following the result for the former case, the latter would imply dominance of action 1, i.e., $\beta_M(\mathbf{a}) = 1$ for all $\mathbf{a} \in A^m$. By Lemma 1, it follows that $\beta_{M'}$ satisfies dominance as well for all $M' \subset M$. However, nontriviality implies nondominance for $|M| = 1$, establishing a contradiction. $\square$

A.4 Proof of Theorem 2

By Theorem 1, learning requires Bayesian contagion. By Lemma 5, in a Bayesian contagion environment, the Bayesian observation function $\beta_M : A^m \rightarrow A$ satisfies Bayesian contagion for all $M \subset N$. Let a^* denote the contagion action. By Lemma 4, we have that f_i updates to a^* for any neighborhood of size 2 whenever a^* is observed at least once. Thus, each f_i coincides with the Bayesian observation function for neighborhoods of size 2. The theorem is established in three steps. The first step considers the neighborhood that equals the set of all agents, and shows that learning implies that for each agent i , $\hat{f}_i(\cdot, N)$ is quasi-Bayesian, i.e., coincides with the Bayesian complete observation function β . The second step considers neighborhoods of sizes at least 3 and establishes that learning implies that $\hat{f}_i$ assigns the contagion action a^* whenever in i 's neighborhood (i) agreement on $\neg a^*$ fails and (ii) there are at least two occurrences of $\neg a^*$. The third step then establishes that for every $N_i \subset N$, we have $\hat{f}_i(\mathbf{a}, N_i) = a^*$ if $a_j = a^*$ for some $j \in N_i$. Hence, we have $\hat{f}_i(\cdot, N_i) = \beta_{|N_i|}$ for all i and, thus, each $\hat{f}_i$ is quasi-Bayesian. We now formally establish the steps outlined above.

Step 1. Suppose that there exists an agent i whose updating function f_i does not coincide with the Bayesian function β_N when i 's observation set $N_i(G) = N$, i.e., there exists a disagreement action vector $\mathbf{a}' \in A^n$ such that $\hat{f}_i(\mathbf{a}', N) = \neg a^* \neq \beta(\mathbf{a}') = a^*$. As $\mathbf{a}'$ is a disagreement vector, without loss of generality we can assume that $a'_i = \neg a^*$ (by anonymity of f_i). Consider a star network G_i with i in the center. We have $\hat{f}_i(\mathbf{a}', N) = \neg a^*$. Note that $\mathbf{a}'$ is a fixed point of $\mathbf{f}(\cdot, G)$. To see this, note that by Lemma 2, all branch agents j with $a'_j = \neg a^*$ remain at $\neg a^*$. All agents j such that $a'_j = a^*$ observe disagreement

and remain at a^* by Lemma 4. As $\mathbf{a}'$ is a fixed point of $\mathbf{f}(\cdot, G_i)$, learning fails, establishing a contradiction. Hence, $\hat{f}_i(\mathbf{a}, N) = \beta(\mathbf{a})$ for all agents $i \in N$ and for all action vectors $\mathbf{a} \in A^n$.

Step 2. Consider an agent i and a neighborhood $\hat{N}_i \neq N$, $|\hat{N}_i| \geq 3$, and an action vector $\mathbf{a}^1$, such that at least two but not all agents j in the neighborhood $\hat{N}_i$ select $a_j = -a^*$. Assume that $\hat{f}_i(\mathbf{a}^1, \hat{N}_i) = -a^*$. As $\hat{f}_i$ is anonymous, we have $a_i^1 \neq a^*$ without loss of generality. Let $N_i^{a^*}$ denote the set of agents in $\hat{N}_i$ with $a_j^1 = a^*$. Construct a connected undirected network G' such that $N_i(G') = \hat{N}_i$ and $N_j(G') = \{j, i\}$ for all agents $j \in N_i^{a^*}$. Next consider an action vector $\hat{\mathbf{a}}^1$ that agrees with $\mathbf{a}^1$ on $\hat{N}_i$, i.e., $\hat{a}_j^1 = a_j^1$ for all $j \in \hat{N}_i$ and where $\hat{a}_j^1 = -a^*$ for all $j \notin N_i^{a^*}$. Note that due to unanimity (Lemma 2) of the agents not in $N_i^{a^*}$ and by Lemma 4, it follows that $\hat{\mathbf{a}}^1$ is a fixed point of $\mathbf{f}(\cdot, G')$, contradicting learning of $\mathbf{f}$.

Step 3. Consider an agent i , a neighborhood $\hat{N}_i \neq N$, $|\hat{N}_i| \geq 3$, and an action vector $\mathbf{a}^1$, such that in $\hat{N}_i$, exactly one agent $j \in \hat{N}_i$ selects the noncontagion action, $a_j = -a^*$. Assume that $\hat{f}_i(\mathbf{a}^1, \hat{N}_i) = -a^*$. Again, as $\hat{f}_i$ is anonymous, we have $a_i^1 = -a^*$ without loss of generality. Construct a network G' such that (i) $N_i(G') = \hat{N}_i$ and (ii) for every $j \in \hat{N}_i$, $j \neq i$, we have $N_j(G') = N$. Next consider an action vector $\hat{\mathbf{a}}^1$ that (a) agrees with $\mathbf{a}^1$ on $\hat{N}_i$, i.e., $\hat{a}_j^1 = a_j^1$ for all $j \in \hat{N}_i$, and (b) $\hat{a}_l^1 = a^*$ for all $l \notin \hat{N}_i$. By Step 1, all agents $j \in \hat{N}_i$, $j \neq i$ satisfy $f_j(\hat{\mathbf{a}}^1, G') = a^*$. By unanimity (Lemma 2), all agents $l \notin \hat{N}_i$ satisfy $f_j(\hat{\mathbf{a}}^1, G') = a^*$. If $f_i(\hat{\mathbf{a}}^1, N_i(G')) = -a^*$, then $\hat{\mathbf{a}}^1$ is a disagreement fixed point of $\mathbf{f}$, contradicting learning.

A.5 Auxiliary results for Theorem 3

We first introduce some further concepts and notation necessary for the subsequent exposition of the results. Let $\#_a^n : A^n \rightarrow \mathbb{N}$ assign to each action vector $\mathbf{a} \in A^n$ the number of agents selecting action $a = \{0, 1\}$:

$$\begin{aligned}\#_1^n(\mathbf{a}) &= \sum_{i \in N} a_i \\ \#_0^n(\mathbf{a}) &= n - \#_1^n(\mathbf{a}).\end{aligned}$$

Similarly, for a given network G and agent i , let $\#_a^i : A^n \rightarrow \mathbb{N}$ count the number of agents in the neighborhood $N_i(G)$ of agent i who select action a . Thus, we have $\#_0^i(\mathbf{a}) = |N_i(G)| - \#_1^i(\mathbf{a})$.

We say that β satisfies a threshold rule if there exists a $\#^*$ such that $\#_1^n(\mathbf{a}) \geq \#^*$ implies $\beta(\mathbf{a}) = 1$ and $\#_1^n(\mathbf{a}) < \#^*$ implies $\beta(\mathbf{a}) = 0$.

LEMMA 6. *For any $n \in \mathbb{N}$, monotonicity implies that the Bayesian complete observation function β satisfies a threshold rule.*

PROOF. As signals are conditional i.i.d., the Bayesian function β is anonymous. Consider a disagreement vector $\mathbf{a}' \in A^n$ such that action 1 is optimal, $\beta(\mathbf{a}') = 1$. As $\mathbf{a}'$ satisfies disagreement there exists an agent j with $a'_j = 0$. By monotonicity, it follows that 1 is optimal conditional on $\mathbf{a}'_{-j}$. Further, monotonicity implies that $(\mathbf{a}'_{-j}, 1)$ induces 1

as a uniquely optimal action. Therefore, if 1 is optimal for an action vector $\mathbf{a}'$, then it is uniquely optimal conditional on any action vector $\mathbf{a}'' \in A^n$ with $\#_1^n(\mathbf{a}'') > \#_1^n(\mathbf{a}')$. $\square$

The following analysis relies on the minimal proportion of agents in N inducing either action as optimal. Let q_a denote the minimal proportion of agents in N necessary to induce action a as optimal, i.e., $q_1 = \frac{\#^*_1}{n}$ and $q_0 = \frac{n - \#^*_1 + 1}{n}$. Further, let q^s denote the smaller of the two, $q^s = \min\{q_0, q_1\}$.³⁸ If $q^s = q_1$, we call action 1 the *stronger* action and call action 0 the *weaker* action. Let a^s denote the stronger action and let a^w denote the weaker action. The following lemma shows that the same threshold proportion that induces a as uniquely optimal given the population vector of actions with cardinality n also applies to action vectors of any smaller cardinality $m < n$.

LEMMA 7. *Consider a monotone environment and a quasi-Bayesian function f_i . Then $\frac{\#_i^a(\mathbf{a})}{|N_i(G)|} \geq q_a$ implies that $f_i(\mathbf{a}, G) = a$.*

PROOF. Let $\hat{\beta}_k : A^k \rightrightarrows A$ denote the Bayesian observation mapping for a network of size $k \leq n$. We first establish by induction that $\frac{\#_i^a(\mathbf{a})}{|N_i(G)|} \geq q_a$ implies that $\hat{\beta}_k(\mathbf{a}) = a$. By nontriviality, the claim is true for $m = 1$. Suppose the claim is true for all $k \leq m$. That is, for all $k \leq m$, $\mathbf{a} \in A^k$, and $\frac{\#_i^a(\mathbf{a})}{k} \geq q_a$, we have $\hat{\beta}_k(\mathbf{a}) = a$. Consider $\hat{\beta}_{m+1}$ and $\mathbf{b} \in A^{m+1}$ such that $\frac{\#_1^{m+1}(\mathbf{b})}{m+1} > q_1$. For the given vector $\mathbf{b} \in A^{m+1}$, consider a vector $\mathbf{a}' = (\mathbf{b}, \mathbf{c}) \in A^n$. Let $l = n - (m+1)$. Suppose that for all $\mathbf{c} \in A^l$, we have $\frac{\#_1^n(\mathbf{a}')}{n} \geq q_1$, which by Lemma 6 and monotonicity implies that $\hat{\beta}_{m+1}(\mathbf{b}) = 1$. Otherwise consider a $\mathbf{c} \in A^l$ such that $\frac{\#_1^n(\mathbf{a}')}{n} = q_1$, which implies $\frac{\#_1^l(\mathbf{c})}{l} < q_1$. Assume that $m+1 > \frac{n}{2}$. There are two cases to consider.

Case 1. Suppose that $\frac{\#_1^l(\mathbf{c})}{l} \leq \frac{\#^*_1 - 1}{n}$, which implies $\hat{\beta}_l(\mathbf{c}) = 0$ by the induction hypothesis. Then $\hat{\beta}_{m+1}(\mathbf{b}) = 1$, as otherwise monotonicity yields $\hat{\beta}(\mathbf{a}') = 0$, leading to a contradiction with $\frac{\#_1^n(\mathbf{a}')}{n} = q_1$.

Case 2. Suppose instead that $\frac{\#_1^l(\mathbf{c})}{l} \in (\frac{\#^*_1 - 1}{n}, \frac{\#^*_1}{n})$ and, by contradiction, that $\hat{\beta}_{m+1}(\mathbf{b}) = 0$. If $0 \in \hat{\beta}(\mathbf{c})$, then, by monotonicity, $\hat{\beta}(\mathbf{a}') = 0$, contradicting $\frac{\#_1^n(\mathbf{a}')}{n} > q_1$. Hence, assume $\hat{\beta}_l(\mathbf{c}) = 1$. Since $m+1 > \frac{n}{2}$, the dimensionality of $\mathbf{c}$ is strictly smaller than that of $\mathbf{b}$. As $\frac{\#_1(\mathbf{b})}{m+1} > q_1$, there exists a vector $\mathbf{d} \in A^k$, where $k = 2(m+1) - n$, and $\mathbf{b}' = (\mathbf{c}, \mathbf{d})$ such that $\frac{\#_1^k(\mathbf{d})}{k} > q_1$, $\frac{\#_1^{m+1}(\mathbf{b}')}{m+1} > q_1$, and $\frac{\#_1^{m+1}(\mathbf{b}')}{m+1} \leq \frac{\#_1^{m+1}(\mathbf{b})}{m+1}$. By the induction hypothesis, $\hat{\beta}_k(\mathbf{d}) = 1$ and, thus, by monotonicity $\hat{\beta}_{m+1}(\mathbf{b}') = 1$. But then, by Lemma 6, we also have $\hat{\beta}_{m+1}(\mathbf{b}) = 1$, establishing a contradiction.

Note that the case $\frac{\#_1^{m+1}(\mathbf{b})}{m+1} = q_1$ and the case $m+1 < \frac{n}{2}$ follows along the same lines. For the latter, we use the largest $k \in \mathbb{N}$ such that $k(m+1) < n$ and a vector $\mathbf{d} \in A^{k(m+1)}$ that consists of k copies of $\mathbf{b}$. To conclude the proof, note that whenever $\hat{\beta}_k$ is single-valued, any quasi-Bayesian function f_i coincides with $\hat{\beta}_k$ for any $|N_i(G)| = k$ for all $k = 1, \dots, n$. $\square$

³⁸Note that in environments with an odd number of agents, q_0 equals q_1 whenever β selects the action with majority.

Let $\delta^+(G)$ denote the cardinality of the largest neighborhood in G , $\delta^+(G) = \max_{i \in N} |N_i(G)|$, and let $\delta^-(G)$ denote the cardinality of the smallest neighborhood in G , $\delta^-(G) = \min_{i \in N} |N_i(G)|$. The following lemma formalizes the information diffusion argument made in the text by relating the degree distribution to the environment.

LEMMA 8. *Consider a monotone environment where Bayesian contagion fails. If $q^s \leq \frac{1}{\delta^-(G)}$, then learning in network G fails.*

PROOF. Consider any agent i with $N_i(G) = \delta^-(G)$, and consider a first-period action vector $\mathbf{a}' \in A^n$ such that $a'_i = a^s$ and $a'_j = a^w$ for all $j \neq i$. Since Bayesian contagion fails, we have $\beta(\mathbf{a}') = a^w$. Assume that $q^s \leq \frac{1}{\delta^-(G)}$. By Lemma 7, agent i updates to a^s and remains at a^s in all subsequent periods, contradicting learning in network G . $\square$

A.6 Proof of Theorem 3

Suppose that monotonicity holds and Bayesian contagion fails. Consider a network G with n agents. Suppose that learning holds in G . By Lemma 6, it follows that for every initial vector $\mathbf{a}^1 \in A^n$ such that $\#_1^n(\mathbf{a}^1) \geq \#^*$, consensus on action 1 emerges, and if $\#_1^n(\mathbf{a}^1) < \#^*$ then consensus on action 0 emerges. Let $N^1(\mathbf{a})$ denote the set of agents selecting action 1 in $\mathbf{a}$. The claim is established in three steps.

Step 1. The first step establishes the following necessary structural condition on G for learning in G to hold: For every group $N^1(\mathbf{a}^1) \subset N$ with $\#_1^n(\mathbf{a}^1) = \#^*$ and for every $i \in N^1(\mathbf{a}^1)$, the shortest path in G from i to $N^1(\mathbf{a}^1) \setminus i$ is smaller than or equal to 2. This claim is established by contradiction in several steps.

- (a) Recall that learning in G implies information retention, i.e., for every $\mathbf{a}^1 \in A^n$ with $\#_1^n(\mathbf{a}^1) < \#^*$, the resulting process of actions $\langle \mathbf{a}^t \rangle_{t \in \mathbb{N}}$ satisfies $\#_1^n(\mathbf{a}^t) < \#^*$ for every $t \geq 1$.
- (b) Suppose that $\#_1^n(\mathbf{a}^1) = \#^*$ and there exists an agent $i \in N^1(\mathbf{a}^1)$ such that the shortest path from i to every other agent in $N^1(\mathbf{a}^1)$ is larger than 2. We show that all agents in N_i select action 0 in period $t = 2$. To see this, note that at the beginning of period $t = 2$, all agents in N_i observe one action 1, while all other observed actions are 0. Therefore, by Lemma 8, in period $t = 2$, all agents in N_i select action 0.
- (c) Denote the set of agents in $N^1(\mathbf{a}^1)$ without i by N_{-i}^1 , i.e., $N_{-i}^1 = N^1(\mathbf{a}^1) \setminus i$. Consider the neighborhood of N_{-i}^1 , i.e., all agents in N_{-i}^1 plus the set of all agents who are neighbors of at least one agent in N_{-i}^1 . Denote this set by $(N_{-i}^1)^1$. Note that $(N_{-i}^1)^1 \cap N_i = \emptyset$. By Step 1(b) and unanimity, in period $t = 2$, all agents who do not belong to $(N_{-i}^1)^1$ select action 0. Learning in G implies by Step 1(a) that the cardinality of agents in $(N_{-i}^1)^1$ who select action 1 in period $t = 2$ is bounded above by $\#^* - 1$. If not, there exists an initial action vector $\hat{\mathbf{a}}^1 \in A^n$ with $\#_1^n(\hat{\mathbf{a}}^1) = \#^* - 1$ and $N^1(\hat{\mathbf{a}}^1) = N_{-i}^1$ such that $\#_1^n(\hat{\mathbf{a}}^2) \geq \#^*$, contradicting Step 1(a) and, therefore, contradicting learning for G . Therefore, for every group $N^1(\mathbf{a}^1) \subset N$ with $\#_1^n(\mathbf{a}^1) = \#^*$ and for every $i \in N^1(\mathbf{a}^1)$, the shortest path in G from i to $N^1(\mathbf{a}^1) \setminus i$ is smaller than or equal to 2.

Step 2. Let $d_G(i, j)$ denote the distance of agent j from agent i in G , i.e., the length of the shortest path from i to j . Let ϑ be the diameter of G and select an agent $i^* \in N$ such that $d_G(i^*, j) = \vartheta$ for some $j \in N$. Assume without loss of generality that $a^s = 1$. Select a vector $\hat{\mathbf{a}}^1 \in A^n$ that satisfies two properties: (i) $\#_1^n(\hat{\mathbf{a}}^1) = \#^* - 1$, and (ii) if $\hat{\mathbf{a}}_l^1 = 1$ and $d_G(i^*, l) = m$, then $\hat{\mathbf{a}}_k^1 = 1$ for all $k \in N$ with $d_G(i^*, k) < m$. Let $\hat{d}_{i^*}$ denote the maximal distance of an agent in $N^1(\hat{\mathbf{a}}^1)$ from i^* , i.e.,

$$\hat{d}_{i^*} = \max_{j \in N^1(\hat{\mathbf{a}}^1) \setminus j^*} d_G(i^*, j).$$

By Step 1, for every agent $j \in N$, we have $d_G(i^*, j) \leq \hat{d}_{i^*} + 2$, which implies that the diameter satisfies $\vartheta \leq \hat{d}_{i^*} + 2$.

Step 3. Next consider a subset $M_{j^*} \subset N$ that has size $\#^* - 1$ and that maximizes the distance to agent i^* . As $a^s = 1$, there exists such a set M_{j^*} with $M_{j^*} \cap N^1(\hat{\mathbf{a}}^1) = \emptyset$. As $M_{j^*} \cap (N^1(\hat{\mathbf{a}}^1)) = \emptyset$, it follows that the minimal distance from i^* to M is bounded below by $\hat{d}_{i^*}$, implying by Step 1 that $\hat{d}_{i^*} \leq 2$. By Step 2, it then follows that for every agent $j \in N$, we have $d_G(i^*, j) \leq 4$. As i^* was chosen such that the distance of at least one agent $j \in N$ from i^* is equal to the diameter, it follows that if Bayesian contagion fails and learning occurs in G , then the diameter of G is smaller than or equal to 4.

A.7 Proof of *Proposition 2*

Suppose there exists an action a and a group $M \subset N$ in G such that the group cohesion $\pi_G(M)$ is greater than or equal to q_a . Without loss of generality, let $a = 1$. Consider a vector $\hat{\mathbf{a}} \in A^n$ such that $N^1(\hat{\mathbf{a}}) = M$, i.e., an action vector where all agents in M select action $a = 1$ and all others select action $a = 0$. Consider any agent $i \in M$. As $\pi_G(M) \geq q_1$ and $N^1(\hat{\mathbf{a}}) = M$, the proportion of agents in N_i who select action $a = 1$ is greater than or equal to q_1 . By *Lemma 7*, agent i therefore remains at action $a = 1$ in the second period and so do all other agents in M . The same reasoning applied inductively implies that all agents in M remain at $a = 1$ in every period $t \in \mathbb{N}$. Action $a = 1$ is optimal conditional on $\hat{\mathbf{a}}$ if the cardinality of M is greater than or equal to $\#_1^*(n)$. But by assumption, the cardinality of M is smaller than $nq_1 = \#_1^*(n)$. Therefore, learning in G fails.

A.8 Auxiliary results for *Theorem 4*

First recall that *learning* implies that *learning in network* G holds for every connected undirected network. To show that Bayesian contagion is a necessary property of the environment for learning to hold, we show that for a star network G' , learning in network G' fails unless the environment satisfies Bayesian contagion.

First we state a version of *Lemma 2* for weakly local updating systems.

LEMMA 9. *Consider a stationary weakly local Markov updating system $\mathbf{f}$. Learning of $\mathbf{f}$ implies that f_i satisfies unanimity for all $i \in N$.*

The proof of this lemma is omitted, as it follows from the exact same argument as the proof of Lemma 2.

Consider a network G' and an agent i such that $N_i(G') = \{i, j\}$. We say that $f_i(\cdot, G')$ is *self-following* if $f_i(\mathbf{a}, G') = a_i$ for all $\mathbf{a} \in A^n$. We say that $f_i(\cdot, G')$ is *other-following* if $f_i(\mathbf{a}, G') = a_j$ for all $\mathbf{a} \in A^n$. Recall that $f_i(\cdot, G')$ is anonymous if there exists an action $\hat{a} \in \{0, 1\}$ such that for every vector $\mathbf{a}' \in A^n$ with $a'_i \neq a'_j$, we have $f_i(\mathbf{a}', G') = \hat{a}$.

LEMMA 10. Consider a stationary weakly local Markov updating function f_i and a network G' such that $|N_i(G')| = 2$. Assume that $f_i(\cdot, G')$ satisfies unanimity. Then $f_i(\cdot, G')$ is either self-following, other-following, or anonymous.

PROOF. Let $j \in N_i(G')$. Since f_i is weakly local, it is invariant in the actions of all agents but i and j . Assume $a'_i = a'_j$. Then by unanimity, we have $f_i(\mathbf{a}', G') = a'_i = a_j$, which is consistent with self-following, other-following, and anonymity. Thus, we need to focus on the two possible cases where $a'_i \neq a'_j$. The possible updates are summarized by Table 1. The first two possible updating functions are anonymous, while the third is self-following and the fourth is other-following. □

LEMMA 11. Consider a stationary weakly local Markov updating system $\mathbf{f}$ and a star network G' . Assume that the environment fails Bayesian contagion. If $\mathbf{f}$ satisfies learning in network G' , then every branch agent is not self-following and not anonymous.

PROOF. As the environment does not satisfy Bayesian contagion, there exist two disagreement vectors $\mathbf{a}', \mathbf{a}'' \in A^n$ such that $\beta(\mathbf{a}') = 0$ and $\beta(\mathbf{a}'') = 1$. Recall that since signals are conditional i.i.d. and the environment satisfies non-indifference and vector non-indifference, any permutation of $\mathbf{b}'$ of $\mathbf{a}'$ satisfies $\beta(\mathbf{b}') = 0$, and similarly for any permutation $\mathbf{b}''$ of $\mathbf{a}''$, we have $\beta(\mathbf{b}'') = 1$. Consider the star network G' . Let i be a branch agent and let j be the center agent in G' . We prove the claim in two steps.

Step 1. Assume that i is self-following. Since $\mathbf{a}'$ is a disagreement vector, there exists a permutation $\mathbf{b}'$ with $b'_i = 1$, and $b'_j = 0$ and $\beta(\mathbf{b}') = 0$. Thus, agent i continues to select action 1 in all periods and learning fails.

Step 2. Assume that i is anonymous. Assume that he updates to action 1 when he observes disagreement. Consider a permutation $\mathbf{b}'$ of $\mathbf{a}'$ with $b'_i = 1$, and $b'_j = 0$ and $\beta(\mathbf{b}') = 0$. Again agent i continues to select action 1 in all periods and learning fails. Assuming he updates to action 0 when observing disagreement, the same arguments can be applied to a permutation $\mathbf{b}''$ of $\mathbf{a}''$, again leading to a contradiction. □

	$f_i(\mathbf{a}', G')$	$f_i(\mathbf{a}', G')$	$f_i(\mathbf{a}', G')$	$f_i(\mathbf{a}', G')$
$a'_i = 1, a'_j = 0$	1	0	1	0
$a'_i = 0, a'_j = 1$	1	0	0	1

TABLE 1. Possible updating functions $f_i(\mathbf{a}', G')$

A.9 Proof of Theorem 4

Consider a stationary weakly local Markov updating system $\mathbf{f}$ that achieves learning. To prove the statement of the theorem, we prove its negation. That is, assume that the environment does not satisfy Bayesian contagion and consider one particular network G' , a star network. We show that learning in network G' fails and, hence, by definition, learning fails.

Consider the star network G' and any branch agent i . By Lemma 10, f_i is either anonymous, self-following, or other-following. Since the environment fails Bayesian contagion and, by assumption, learning holds by Lemma 11, each branch agent is other-following, i.e., each branch agent i copies the last period action of the center agent j^* . Consider a first-period disagreement vector $\mathbf{b}$ such that $\beta(\mathbf{b}) = 0$ and $b_{j^*} = 1$. Let $\mathbf{b}^2$ denote the second-period action vector, $\mathbf{f}(\mathbf{b}) = \mathbf{b}^2$. By the argument above, we have $b_i^2 = 1$ for all agents $i \neq j^*$. Since $\mathbf{f}$ is stationary and since learning holds, we have $\beta(\mathbf{b}^2) = 0$. Furthermore, $b_{j^*}^2 = 0$, as otherwise unanimity induces a contradiction with learning. Since the environment is monotone, the complete Bayesian observation function β satisfies a threshold rule by Lemma 6. In particular, by the argument above, if at least one agent in N selects action 0, then 0 is optimal, contradicting the assumption that Bayesian contagion fails.

A.10 Proof of Theorem 5

Consider a complete group M^* of size m . We first show that if $m \geq \frac{\ln n}{\ln 2} + 1$, then for each $k = m, m+1, \dots, n-1$, there exists an injective mapping $g_k : \{0, 1, 2, \dots, k\} \rightarrow A^{m-1}$, where the number corresponds to the number of 1 actions in a group of size k . For a group of size $m-1$, the number of distinct action vectors $\mathbf{a} \in A^{m-1}$ equals 2^{m-1} . Thus, we require

$$2^{m-1} \geq n$$

$$m \geq \frac{\ln n}{\ln 2} + 1.$$

Before we use these injective mappings to construct an updating system $\langle \mathbf{f}' \rangle_{t \in \mathbb{N}}$ in several steps, let us first lay out the principal idea. Note that the members of the core M^* jointly observe their action vector $\mathbf{a}_{M^*} \in A^m$. In the first phase of the updating process, in every period t , one core agent introduces new information into the core while the remaining core members retain the current amount of information. Partition the set of noncore agent as follows: select a noncore agent j_1 with maximal distance to the core and consider one shortest path to the core including the respective core agent. Assign all agents on the path including j_1 but excluding the core agent to one partition cell. We denote the cell corresponding to the core agent i as P_i . Any noncore agents whose unique shortest path overlaps with this cell is assigned to the same cell and, hence, to the same core agent. Proceed in the same manner until all noncore agents are assigned to one core agent, and order the core agents from 1 to m . Let $\mathbf{f}'$ denote the updating mapping applied in period $t+1$.

Step 1. Let $\mathbf{f}^1$ be as follows: For the agents in $M_1^* = M^* \setminus 1$, let

$$\mathbf{f}_{M_1^*}^1(\mathbf{a}) = g_m \left(\sum_{i \in M^*} a_i \right).$$

For agent 1, consider one path $\pi \in P_1$ originating at agent 1, and order the agents belonging to π from lowest to largest distance to 1 and denote them by $\pi_1, \pi_2, \dots, \pi_k$. Let $f_1^1(\mathbf{a}) = a_{\pi_1}$ and $f_{\pi_j}^1(\mathbf{a}) = a_{\pi_{j+1}}$ for all $j = 1, \dots, k-1$, and let $f_l^1(\mathbf{a}) = a_l$ for $l = \pi_k$ and all $l \in N \setminus M^* \setminus P_1$. Then set $\mathbf{f}^2$ as follows: For the agents in M_1^* ,

$$\mathbf{f}_{M_1^*}^2(\mathbf{a}) = g_{m+1} \left(g_m^{-1}(\mathbf{a}_M) + a_1 \right).$$

Again let $f_1^1(\mathbf{a}) = a_{\pi_1}$ and $f_{\pi_j}^1(\mathbf{a}) = a_{\pi_{j+1}}$ for all $j = 1, \dots, k-2$, and let $f_l^1(\mathbf{a}) = a_l$ for $l = \pi_k, \pi_{k-1}$ and all $l \in N \setminus M^* \setminus P_1$. Construct $\mathbf{f}^r$ in a similar way to pass down the information from the agents in the path π , all other paths in P_1 , as well as all other cells P_i corresponding to the core agents $i \in M^*$.

Step 2. Let $\hat{\beta} : \{0, 1, \dots, n\} \rightarrow \mathcal{A}$ assign the Bayesian optimal action as a function of the number of 1 actions in the network. Note that since signals are conditional i.i.d., the distribution of actions is sufficient to determine the optimal action, i.e.,

$$\beta(\mathbf{a}) = \hat{\beta} \left(\sum_{i \in N} a_i \right)$$

for all $\mathbf{a} \in \mathcal{A}^n$. Let t^* be the last period where a core agent i^* passes information to $M_{i^*}^*$, the remainder of the core. The action vector of all agents in M^* now reveals the first-period distribution of actions and is sufficient for each agent in M^* to select the optimal action in period $t^* + 1$. We have

$$\mathbf{f}_{M^*}^{*+1}(\mathbf{a}) = \hat{\beta} \left(g_{n-1}^{-1}(\mathbf{a}_M^{t^*}) + a_{i^*} \right).$$

For all agents in M^* , we have $f_i^t(\mathbf{a}) = a_i$ for all $t \geq t^* + 2$. Thus, from $t^* + 1$ and onward, information only needs to diffuse from the core to the noncore agents. For all first degree neighbors j of core agent i simply set $f_j^{t^*+2}(\mathbf{a}) = a_i$ and equal to their own last period action in all later periods. Similarly, construct the updating system for k th degree neighbors. Thus, from $t^* + 2$, the optimal action spreads along the shortest path in the network.

APPENDIX B: EXAMPLES

B.1 Environments satisfying Bayesian contagion are highly asymmetric

Consider an environment with binary states, binary actions, and binary signals, $\Omega = \mathcal{A} = \mathcal{S} = \{0, 1\}$. Both states are a priori equally likely. Each agent achieves a utility of 1 if his action matches the realized state and achieves 0 otherwise. Thus, action $a = 0$ ($a = 1$) is uniquely optimal if and only if the conditional probability of state $\omega = 0$ ($\omega = 1$) is greater than 0.5. Assume that the environment is highly asymmetric in that the dis-

tribution of signals conditional on states is highly uneven: $\Pr(s = 0 | \omega = 0) = \frac{99}{100}$ and $\Pr(s = 1 | \omega = 1) = \frac{1}{10}$. Assuming a prior probability $\Pr(\omega = 1) = \mu$, the posterior probabilities conditional on s are $\Pr(\omega = 1 | s = 1) = \frac{10\mu}{10\mu + (1-\mu)}$ and $\Pr(\omega = 1 | s = 0) = \frac{90\mu}{90\mu + (1-\mu)99}$. Such values of the posterior probability have two implications. First, for a uniform prior $\mu = 0.5$, the action $a = s$ is, conditional on s , uniquely expected-utility maximizing. Thus, the chosen action reveals the underlying signal. Second, the posterior probability of state $\omega = 1$ is strictly larger than the prior probability if $s = 1$ and strictly smaller than the prior probability if $s = 0$. As a consequence, unanimity is satisfied, i.e., if there is consensus on an action, then that action is optimal conditional on the consensus action vector. Next consider a vector $\hat{\mathbf{a}}^n \in A^{n+1}$ such that exactly one agent selects action $a = 1$, i.e., $\sum_{i=1}^{n+1} \hat{a}_i^n = 1$. Under the uniform prior, the posterior probability of state $\omega = 1$ is $\Pr(\omega = 1 | \hat{\mathbf{a}}^n) = 1/[1 + \frac{1}{10}(\frac{99}{90})^n]$. For $n = 0, \dots, 30$, therefore, action $a = 1$ is uniquely expected-utility maximizing conditional on $\hat{\mathbf{a}}^n$. In other words, in this example, Bayesian contagion holds for network sizes up to 31.

B.2 *Learning in a network with diameter 2 despite failure of Bayesian contagion*

Consider a network G with five agents, $N = \{1, 2, 3, 4, 5\}$. Suppose that the environment is symmetric, i.e., $q_1 = q_0 = \frac{3}{5}$. The network has the following properties. Agents $i = 1, 2, 3$ observe all members of the network, i.e., $N_i = N$ for $i = 1, 2, 3$. We denote the set of the first three agents as the core of the network. The remaining agents $j = 4, 5$ do not observe each other, i.e., $N_j = \{1, 2, 3, j\}$, and are called periphery agents. Note that in period $t = 2$, each core agent selects the action that is optimal conditional on $\mathbf{a}^1$. As $q_1 = q_0 = \frac{3}{5}$, the core agents remain at the optimal action in all periods $t \geq 2$. This implies that each periphery agent sees a proportion of at least $\frac{3}{4}$ of the agents in his observation set selecting the optimal action from period $t = 2$ onward. As $q_1 = q_0 = \frac{3}{5}$, they select the optimal action from period $t = 3$ onward. Therefore, consensus on the optimal action occurs by period $t = 3$ for any first-period action vector. Note that analogous examples can be constructed for any arbitrary large network as long as there exists a core group of agents making up a majority and such that each of them observes all agents in the network.

B.3 *Quasi-Bayesian updating leading to consensus on the suboptimal action*

Consider a monotone environment where Bayesian contagion fails. Consider a network G such that the maximal neighborhood size $\delta^+(G)$ satisfies $q^s \leq \frac{1}{\delta^+(G)}$. This implies that every agent in the network selects the stronger action whenever he observes at least one occurrence of the stronger action in his observation set. Thus, the stronger action, whenever it is selected at least once in the first period, spreads throughout the network along the shortest path, resembling the process of actions in a Bayesian contagion environment. For all initial action vectors but consensus on the weaker action, consensus on the stronger action emerges. However, since the environment does not satisfy Bayesian contagion, this outcome is not optimal for all first-period action vectors.

REFERENCES

- Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), “Bayesian learning in social networks.” *Review of Economic Studies*, 78, 1201–1236. [318, 319, 324]
- Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken (2016), “Network structure and the aggregation of information: Theory and evidence from Indonesia.” *American Economic Review*, 106, 1663–1704. [318, 323]
- Arieli, Itai, Yakov Babichenko, and Manuel Mueller-Frank (2020), “Naive learning through probability matching.” Unpublished paper, SSRN 3338015. [342]
- Arieli, Itai and Manuel Mueller-Frank (2017), “Inferring beliefs from actions.” *Games and Economic Behavior*, 102, 455–461. [326, 328, 342]
- Arieli, Itai and Manuel Mueller-Frank (2019), “Multi-dimensional social learning.” *Review of Economic Studies*, 86, 913–940. [318]
- Bala, Venkatesh and Sanjeev Goyal (1998), “Learning from neighbours.” *Review of Economic Studies*, 65, 595–621. [320, 322]
- Banerjee, Abhijit V. (1992), “A simple model of herd behavior.” *Quarterly Journal of Economics*, 107, 797–817. [318]
- Banerjee, Abhijit V., Emily Breza, Arun G. Chandrasekhar, and Markus Mobius (2019), “Naive learning with uninformed agents.” Working paper, No. 25497. [342]
- Banerjee, Abhijit V., Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson (2013), “The diffusion of microfinance.” *Science*, 341. [318]
- 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. [318, 319, 324]
- Bohren, J. Aislinn (2016), “Informational herding with model misspecification.” *Journal of Economic Theory*, 163, 222–247. [320, 322]
- Breza, Emily, Arun Chandrasekhar, and Alireza Tahbaz-Salehi (2018), “Seeing the forest for the trees? An investigation of network knowledge.” Working paper, NBER Working Paper No. w24359. [327]
- Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong (2004), “A cognitive hierarchy model of games.” *Quarterly Journal of Economics*, 119, 861–898. [324]
- Casciaro, Tiziana (1998), “Seeing things clearly: Social structure, personality, and accuracy in social network perception.” *Social Networks*, 20, 331–351. [327]
- Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri (2020), “Testing models of social learning on networks: Evidence from two experiments.” *Econometrica*, 88, 1–32. [318, 319, 323, 326, 342]

Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri (2013), “Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications.” *Journal of Economic Literature*, 51, 5–62. [324]

Dasaratha, Krishna and Kevin He (2020), “Network structure and naive sequential learning.” *Theoretical Economics*, 15, 415–444. [322]

De Marzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel (2003), “Persuasion bias, social influence, and unidimensional opinions.” *Quarterly Journal of Economics*, 118, 909–968. [331, 342]

DeGroot, Morris H. (1974), “Reaching a consensus.” *Journal of the American Statistical Association*, 69, 118–121. [318, 319, 326]

Enke, Benjamin and Florian Zimmerman (2019), “Correlation neglect in belief formation.” *Review of Economics Studies*, 86, 313–332. [318]

Eyster, Erik and Matthew Rabin (2010), “Naive herding in rich-information settings.” *American Economic Journal: Microeconomics*, 2, 221–243. [320, 322]

Eyster, Erik, Matthew Rabin, and Georg Weizsäcker (2015), “An experiment on social mislearning.” Unpublished paper, SSRN 2704746. [318, 323]

Feldman, Michal, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg (2014), “Reaching consensus via non-Bayesian asynchronous learning in social networks.” Unpublished paper, arxiv:1408.5192. [322, 323]

Gagnon-Bartsch, Tristan and Matthew Rabin (2016), “Naive social learning, mislearning, and unlearning.” Unpublished paper, Department of Economics and Business School, Harvard University. [322]

Gale, Douglas M. and Shachar Kariv (2003), “Bayesian learning in social networks.” *Games and Economic Behavior*, 45, 329–346. [318]

Glaeser, Edward and Cass R. Sunstein (2009), *Extremism and Social Learning*, volume 1, 263–324. [331]

Goles, Eric and J. Olivos (1980), “Periodic behavior of generalized threshold functions.” *Discrete Mathematics*, 30, 187–189. [338]

Golub, Benjamin and Matthew Jackson (2012), “How homophily affects the speed of learning and best-response dynamics.” *Quarterly Journal of Economics*, 127, 1287–1338. [319, 326]

Golub, Benjamin and Matthew O. Jackson (2010), “Naïve learning in social networks and the wisdom of crowds.” *American Economic Journal: Microeconomics*, 2, 112–149. [319, 326, 337, 342]

Hazla, Jan, Ali Jadbabaie, Elchanan Mossel, and M. Amin Rahimian (2018), “Reasoning in Bayesian opinion exchange networks is PSPACE-hard.” Unpublished paper, arXiv:1809.01077. [318]

- Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012), “Non-Bayesian social learning.” *Games and Economic Behavior*, 76, 210–225. [319, 326]
- Krackhardt, David (1990), “Assessing the political landscape: Structure, cognition, and power in organizations.” *Administrative Science Quarterly*, 35, 342–369. [327]
- Kübler, Dorothea and Georg Weizsäcker (2004), “Limited depth of reasoning and failure of cascade formation in the laboratory.” *Review of Economic Studies*, 71, 425–441. [318, 323]
- Lee, In Ho (1993), “On the convergence of informational cascades.” *Journal of Economic Theory*, 61, 395–411. [342]
- Lobel, Ilan and Evan Sadler (2015), “Information diffusion in networks through social learning.” *Theoretical Economics*, 10, 807–851. [318]
- March, Christoph and Anthony Ziegelmeyer (2018), “Excessive herding in the laboratory: The role of intuitive judgments.” Unpublished paper, SSRN 3140365. [318, 323]
- Mobius, Markus, Tuan Phan, and Adam Szeidl (2015), “Treasure hunt: Social learning in the field.” [323]
- Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018), “A theory of non-Bayesian social learning.” *Econometrica*, 86, 445–490. [322, 326, 328, 342]
- Morris, Stephen (2000), “Contagion.” *Review of Economic Studies*, 67, 57–78. [323, 338]
- Mossel, Elchanan, Manuel Mueller-Frank, Allan Sly, and Omer Tamuz (2020), “Social learning equilibria.” *Econometrica*, 88, 1235–1267. [319, 324]
- Mossel, Elchanan, Allan Sly, and Omer Tamuz (2015), “Strategic learning and the topology of social networks.” *Econometrica*, 83, 1755–1794. [318, 319, 323, 324]
- Mueller-Frank, Manuel (2013), “A general framework for rational learning in social networks.” *Theoretical Economics*, 8, 1–40. [318, 328]
- Mueller-Frank, Manuel (2014), “Does one Bayesian make a difference?” *Journal of Economic Theory*, 154, 423–452. [328]
- Mueller-Frank, Manuel and Claudia Neri (2015), “A general model of boundedly rational observational learning: Theory and evidence.” Unpublished paper, SSRN 2566210. [342]
- Nagel, Rosemarie C. (1995), “Unraveling in guessing games: An experimental study.” *American Economic Review*, 85, 1313–1326. [324]
- Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2009), “Information externalities and emergence of consensus,” *Games and Economic Behavior*, 66, 979–994. [318]
- Smith, Lones and Peter Sørensen (2000), “Pathological outcomes of observational learning.” *Econometrica*, 68, 371–398. [318, 319, 324]

Manuscript received 19 August, 2017; final version accepted 26 June, 2020; available online 11 July, 2020.