

A GENERAL MODEL OF BOUNDEDLY RATIONAL OBSERVATIONAL LEARNING: THEORY AND EXPERIMENT

Manuel Mueller-Frank

Claudia Neri

The Public-Private Sector Research Center is a Research Center based at IESE Business School. Its mission is to develop research that analyzes the relationships between the private and public sectors primarily in the following areas: regulation and competition, innovation, regional economy and industrial politics and health economics.

Research results are disseminated through publications, conferences and colloquia. These activities are aimed to foster cooperation between the private sector and public administrations, as well as the exchange of ideas and initiatives.

The sponsors of the Public-Private Sector Research Center are the following:

- Ajuntament de Barcelona
- Departament d' Economia i Coneixement de la Generalitat de Catalunya
- Departament d' Empresa i Ocupació de la Generalitat de Catalunya
- Diputació de Barcelona
- EVERIS
- Fundació AGBAR
- Institut Català de les Empreses Culturals (ICEC)
- PricewaterhouseCoopers
- Sanofi

The contents of this publication reflect the conclusions and findings of the individual authors and not the opinions of the Center's sponsors.

A General Model of Boundedly Rational Observational Learning: Theory and Experiment*

Manuel Mueller-Frank and Claudia Neri[†]

February 2015

Abstract

This paper introduces a new general model of boundedly rational observational learning: Quasi-Bayesian updating. The approach is applicable to any environment of observational learning and is rationally founded. We conduct a laboratory experiment and find strong supportive evidence for Quasi-Bayesian updating. We analyze the theoretical long run implications of Quasi-Bayesian updating in a model of repeated interaction in social networks with binary actions. We provide a characterization of the environment in which consensus and information aggregation is achieved. The experimental evidence is in line with our theoretical predictions. Finally, we establish that for any environment information aggregation fails in large networks.

Keywords: social networks, naive learning, bounded rationality, experiments, consensus, information aggregation.

JEL codes: C91, C92, D83, D85.

*We have benefitted from the opportunity to present this paper at IESE Business School. We are grateful to seminar participants and to Philipp Kircher for their comments. This research has been supported in part by Swiss National Science Foundation grant 100014/135346.

[†]Manuel Mueller-Frank, corresponding author: Department of Economics, IESE Business School, University of Navarra. Phone +34 932 534 200, Email: mmuellerfrank@iese.edu. Claudia Neri: Department of Economics, University of St.Gallen, Varnbuelstrasse 19, St.Gallen 9000 Switzerland. Phone: +41 71 224 2757, Email: claudia.neri@unisg.ch.

1 Introduction

Social learning is a crucial component of human interaction. Among the many possible mechanisms by which individuals learn from others, observational learning describes 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 make and update their behavior after observing the behavior of others and which long run aggregate outcomes such learning generates has important implications for economic policy. For example, the role of observational learning taking place over social networks such as Facebook may explain the effectiveness of social advertising relative to standard display advertising (Mueller-Frank and Pai 2015)¹. Also, recently there has been a substantial interest for observational learning in development economics. In particular, observational learning taking place over the network of relationships within a rural village may explain the effectiveness of an information campaign aimed at increasing the adoption of microfinance loans (Banerjee, Chandrasekhar, Duflo and Jackson 2013) or of a community-based targeting program aimed at selecting aid beneficiaries (Alatas, Banerjee, Chandrasekhar, Hanna and Olken 2014).

In the literature there are two predominant approaches to study observational learning: one based on Bayesian updating and one based on boundedly rational updating. 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. This is the standard approach in the sequential social learning literature² and in parts of the literature on repeated interaction in social networks.³ Despite being a very useful benchmark, the Bayesian approach has a severe weakness: the rationality assumption is 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.⁴

The boundedly rational approach is particularly common in the literature on repeated interaction in social networks, due to the complexity of inferences. DeGroot (1974) is the standard model within this approach. The DeGroot model addresses the weakness of the Bayesian approach by assuming that agents use a simple heuristic (rule of thumb), revising their decision to a weighted average of their own and their neighbors' previous decisions. However, the specification of a weighted average updating function, while making the analysis of the long run evolution of actions tractable,

¹In social advertising individuals who purchased a specific product are highlighted to their network friends.

²See Bikhchandani, Hirshleifer and Welch (1992), Banerjee (1992), Smith and Sorensen (2000), Acemoglu, Dahleh, Lobel and Ozdaglar (2011), Arieli and Mueller-Frank (2014).

³See Bala and Goyal (1998), Gale and Kariv (2003), Rosenberg, Solan and Vieille (2009), and Mueller-Frank (2013a), Mossel, Sly and Tamuz (2012), Lobel and Sadler (2015). For an excellent survey of the literature see Goyal (2014).

⁴Throughout the paper, we use the terms action, choice, and decision as synonyms.

is somewhat arbitrary. Moreover, it has been shown that DeGroot updating might lead to very undesirable long run outcomes (Mueller-Frank 2014).⁵ Finally, the DeGroot model applies only to infinite real-numbered action spaces.⁶ Therefore, both the Bayesian approach and the boundedly rational approach have weaknesses which severely limit their scope. This paper is motivated by such unresolved weaknesses.

In this paper we propose a new general model of boundedly rational observational learning, which is based on the concept of Quasi-Bayesian updating. The concept 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. This assumption reduces complexity compared to Bayesian updating because considerations as to how each observed action might have been affected by other actions, both observed and not observed, are not necessary.

The Quasi-Bayesian approach addresses the weaknesses of the Bayesian and DeGroot approaches. First, it is applicable to any environment of observational learning, without restrictions on the nature of the utility functions, state space, action space or signal space.⁷ Additionally, for each given environment, Quasi-Bayesian updating is not arbitrary but has a rational foundation. While these advantages might make Quasi-Bayesian updating appealing from a theoretical standpoint, the key question is whether it can accurately describe how individuals update their actions. We conduct a laboratory experiment and find strong supportive evidence for Quasi-Bayesian updating.

We also consider the theoretical long run implications of Quasi-Bayesian updating in a model of repeated interaction in social networks. We contribute to the existing literature by showing that under Quasi-Bayesian updating the predictions under coarse (i.e. binary) actions differ substantially from those established in the literature under uncountable actions. First, consensus is hard to achieve as it requires highly asymmetric environments. Second, consensus and information aggregation generally coincide. Third, information aggregation is achievable in finite networks but not in infinite networks. Our results are relevant since many real world environments are better described by binary than by uncountable actions. In our laboratory experiment we find that consensus often fails to arise, but that when it does occur, it occurs almost always with information aggregation, in line with our theoretical predictions.

⁵Mueller-Frank (2014) considers a model in which agents receive initial private information and then repeatedly announce their conditional probability of an uncertain event to their neighbors. He shows that the long run consensus probability fails to represent the private information of agents, in any finite network and for any possible weighted average function of each agent.

⁶From a technical point of view, real vector spaces are admissible as well.

⁷Only minimal technical assumptions are required to assure existence of optimal actions, joint and regular conditional probability measures.

Theoretical model We consider an observational learning model with finitely many agents who share a common prior over a general state space. Each agent receives an i.i.d. private signal belonging to a general signal space. The distribution over the signal space depends on the realized state of the world. After observing his private signal, each agent takes an expected utility maximizing action out of a compact metrizable action set. The utility of each agent depends only on his action and the state of the world. We assume that all agents share the same utility function.⁸

An individual signal strategy for an agent denotes a mapping from his signal realization to expected utility maximizing actions. The Quasi-Bayesian updating function assigns an action to any observed subset of agents and all possible corresponding observed action vectors such that the assigned action is Bayes optimal conditional on the realized set of actions and assuming that each individual observed action is selected according to the corresponding individual signal strategy. In other words, Quasi-Bayesian updating is equal to Bayesian updating assuming that each agent selects an action based only on his private signal.

However, in most observational learning models, agents choose their action based not only on their private signal but also on information inferred from the actions chosen by others. In the sequential social learning model, each agent chooses his action based on his private signal and the history of actions of the agents that decided before him. In the models of repeated interaction in social networks, each agent considers the history of actions of his neighbors, in addition to his private signal. Therefore, in observational learning models Quasi-Bayesian updating functions typically do not coincide with Bayesian updating functions. Nevertheless, Quasi-Bayesian updating functions are applicable to any observational learning model and have the additional advantage of being rationally founded.

Theoretical results We analyze the implications of Quasi-Bayesian updating for the model of repeated interaction in social networks. Here every agent takes an action in each of countable rounds and observes the history of actions of his neighbors. We allow for a general signal and state space, but restrict attention to binary actions, which is a common assumption in the observational learning literature.⁹ We further assume that the environment satisfies two properties, reverse symmetry and monotone likelihood, which we show are also satisfied in the setting commonly considered in the social learning literature. We assume that every agent updates his action according to a Quasi-Bayesian updating function. That is, in each round he assigns a Quasi-Bayesian action to the action vector observed in the previous round, which consists of the actions of his neighbors and himself.¹⁰

⁸This assumption is not necessary and is made only to simplify the exposition.

⁹See for example Bikhchandani, Hirshleifer and Welch (1992), Smith and Sorensen (2000) and Acemoglu, Dahleh, Lobel and Ozdaglar (2011).

¹⁰The focus on Markov revision functions is standard in the analysis of non-Bayesian learning in networks. The seminal paper by DeGroot (1974) gave rise to parallel literatures in computer science, electrical engineering and

We provide a characterization of the necessary and sufficient conditions for consensus and information aggregation. Consensus is defined as the convergence towards agreement on one action among all agents in the network, for every strongly connected network and for every initial action vector.¹¹ This notion of consensus coincides with the one established for Bayesian updating in networks, for DeGroot updating and its generalization.¹² Our first theorem establishes that under Quasi-Bayesian updating consensus holds if and only if there exists an action a' which is optimal conditional on every action vector (of all agents) except for the action vector where all agents choose the other action a'' .¹³ Therefore, the condition is very strong, since, in order for it to be satisfied, the environment needs to be highly asymmetric.

Information aggregation is defined as convergence of the actions of all agents towards an action that is optimal conditional on the vector of first round actions of all agents, for every such vector and for every strongly connected network. This definition is weaker than the one in the literature on Bayesian learning in networks, which requires convergence to an action that is optimal conditional on the realized signals of all agents. Under Quasi-Bayesian updating, the sequence of action vectors is determined by the first round action vector. If actions fail to reveal the underlying signals, then the stronger notion of information aggregation is unachievable. Our second theorem establishes that information aggregation holds if and only if there exists an action a' which is optimal conditional on every action vector (of all agents) except for the action vector where all agents choose the other action a'' . Therefore, under binary actions and Quasi-Bayesian updating consensus and information aggregation coincide. Our third theorem shows that under Quasi-Bayesian updating and binary actions the naive learning result of Golub and Jackson (2010) does not hold. That is, for every environment there exists a finite network size n^* such that information aggregation (and consensus) fails for all networks of size larger than n^* .

The theoretical results on Quasi-Bayesian updating display some striking differences compared to the existing literature on boundedly rational updating. First, we show that under binary actions consensus is hard to achieve as it requires highly asymmetric environments. In contrast, the existing literature shows that consensus in the DeGroot model and its generalization requires only that each agent takes his own previous round action into account when updating. Also, consensus in these models is robust to changes of the updating functions over time. This difference in result is mainly driven by the difference in the set of actions, binary in our model while uncountable in the DeGroot model. Second, we show that information aggregation is achievable in finite networks under Quasi-

economics, all sharing the Markov assumption. For some recent examples see Lobel, Ozdaglar and Feijer (2011), Blondel, Hendrickx and Tsitsikilis (2009) and Golub and Jackson (2010).

¹¹A network is strongly connected if there exists a directed path from every agent to every other agent.

¹²See Rosenberg, Solan and Vieille (2009) and Mueller-Frank (2013a) for Bayesian updating, DeMarzo, Vayanos and Zwiebel (2003) for DeGroot updating and Mueller-Frank (2013b) for a generalization of DeGroot updating.

¹³To capture indifference case, additionally we require the Quasi-Bayesian revision functions of all agents to select a' in case of indifference.

Bayesian updating. To the best of our knowledge no such result has been established in the literature on boundedly rational learning in networks.¹⁴ Third, while in the DeGroot model information aggregation in finite networks is not achievable, the conditions for information aggregation in infinite networks are relatively weak. In contrast, in our model the opposite holds. Information aggregation might occur in finite but not in infinite networks.

Experimental results The experiment consists of a series of urn-guessing games played by subjects connected by a network structure. The design is adapted from the urn-guessing experiment, which is standard in the experimental literature on observational learning, and was previously applied both to sequential social learning (Anderson and Holt 1997, Hung and Plott 2001, and Kübler and Weizsäcker 2004) and to social learning in networks (Choi, Gale and Kariv 2005, and Grimm and Mengel 2014). The game is played with urns, each containing a different composition of colored balls. After the computer has randomly selected one urn, each agent privately observes one draw (with replacement) from the selected urn. In each of several decision rounds, every agent is asked to select the urn that he thinks is more likely to have been used. After observing his neighbors' decisions in the previous round, each agent is thereafter asked again to select the urn that he thinks is more likely to have been used. At the end of the experiment each participant earns a positive payoff if he correctly selected the urn that was used, and nothing otherwise. Within the experiment we vary network size (5 or 7 agents), choice-set size (2 or 4 urns), and network connections.

We find that the rate at which choice behavior is consistent with the definition of Quasi-Bayesian updating is extremely high across network size, choice-set size, and network connections, whether we pool the data of all participants together or allow for heterogeneity across participants. Pooling the data of all participants together, choice behavior is consistent with Quasi-Bayesian updating in more than 90 percent of observations. Allowing for heterogeneity across participants, and computing for each participant an individual score of Quasi-Bayesian updating, the median individual score is higher than 92 percent. In terms of long-run properties of Quasi-Bayesian updating, we find that consensus often fails to arise, but that when it arises, it occurs almost always with information aggregation, which is in line with our theoretical predictions.

Our paper is closely related to Choi, Gale and Kariv (2005), Möbius, Phan and Szeidl (2010), Chandrasekhar, Larreguy and Xandri (2012), and Grimm and Mengel (2014). Choi, Gale and Kariv (2005) test the predicted aggregate outcomes of the Bayesian model of Gale and Kariv (2003) in 3-agent networks and find evidence of a strong tendency toward consensus and a high efficiency of information aggregation. Möbius, Phan and Szeidl (2010) study learning in a field experiment on Facebook and propose a streams model in which, in contrast to the DeGroot model, agents process information that originates within the network by tagging its source and thus avoid double-counting

¹⁴Recall that we require information aggregation to hold in every strongly connected network.

information. Chandrasekhar, Larreguy and Xandri (2012) and Grimm and Mengel (2014) perform an horse-race between the Bayesian model and the DeGroot model in 7-agent network experiments. Chandrasekhar, Larreguy and Xandri (2012) find evidence supporting the DeGroot model, including simple majority rules.¹⁵ Grimm and Mengel (2014) find that the DeGroot model outperforms the Bayesian model in explaining individual decisions, but that aggregate properties (consensus and information aggregation) are only partially consistent with the DeGroot model. They find evidence that experimental participants use an heuristic, which, differently from the DeGroot model, takes into account, at least partially, for correlations in neighbors’ actions.

Despite being closely related to the works mentioned above, our paper differs from them in focus and objective. We aim at proposing Quasi-Bayesian updating as an alternative to DeGroot updating with the boundedly rational approach, yielding a model which is at the same time theoretically tractable, rationally founded and empirically relevant.

Organization of the paper The rest of the paper is organized as follows. In Section 2, we propose a general model of boundedly rational observational learning and introduce the concept of Quasi-Bayesian updating. In Section 3, we apply the concept of Quasi-Bayesian updating to repeated interaction in social networks. In Section 4 we describe our experimental design and in Section 5 we analyze whether the behavior of experiment participants is consistent with Quasi-Bayesian updating. In Section 6, we characterize the long-run properties of Quasi-Bayesian updating in terms of consensus and information aggregation, and present our theoretical results. In Section 7 we report experimental evidence on consensus and information aggregation. Finally, Section 8 concludes. All proofs are provided in Appendix A. The experiment instructions are reported in Appendix B. Additional tables and figures are included in Appendix C.

2 A general model of boundedly rational observational learning

In this section we introduce a general approach towards modeling boundedly rational observational learning. Our approach captures a wide range of environments and structures of decision making. Among others, it is applicable to repeated interaction as well as to interaction occurring in a strictly sequential or multi-dimensional order. Consider a countable set of agents N , each of which faces uncertainty regarding the state of the world ω . The cardinality of N is denoted by n , which might be finite or infinite. All agents share a common prior p over a Polish state space Ω with Borel σ -algebra

¹⁵In a binary action setting, Chandrasekhar, Larreguy and Xandri (2012) and Grimm and Mengel (2014) employ a DeGroot ‘action model’ in which agents revise their beliefs over the two actions according to a DeGroot rule, then choose the action whose corresponding belief is higher than 0.5, and finally announce their choice. They distinguish the ‘action model’ from a ‘communication model’, in which beliefs are directly announced.

\mathcal{F} . The state of the world is drawn according to the common prior and each agent receives a private signal $s_i \in S$, where S is a standard Borel space.¹⁶ The distribution F_ω , according to which the signal s_i is drawn, depends on the realized state of the world, $F_\omega \in \Delta(S)$ for $\omega \in \Omega$.¹⁷ 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 are not perfectly informative regarding the realized state of the world. Conditional on the realized state of the world, signals are identically distributed and independent among agents.

All agents $i \in N$ need to take an action $a_i \in A$, where A is a compact metrizable space of actions. Suppose that all agents share identical preferences represented by a continuous utility function $u : A \times \Omega \rightarrow \mathbb{R}$. As is the norm in the observational learning literature, we restrict attention to settings without payoff externalities. Let σ_i^s be the strategy of agent i that assigns the expected utility maximizing action to each private signal realization $s_i \in S$,

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

where $\Pr(d\omega | s)$ is the posterior distribution over the state space based upon the prior p and conditional on the realized signal s_i . We use the term individual signal strategy for the strategy σ_i^s and denote the individual signal strategy profile by $\sigma^s = \langle \sigma_i^s \rangle_{i \in N}$. For a set of agents $N' \subseteq N$ let $\mathbf{a}_{N'}^s$ denote the realization of the random action vector of agents in N' given that each follows his individual signal strategy.

Observational learning entails drawing an inference on the realized private information of an agent from observing the action chosen by that agent. In most models of observational learning considered in the literature, however, an agent's action generally is based not only on his signal realization and the common prior, but also on additional information gained observing the actions of other agents. As an example, consider the literature on sequential social learning, where agents act in a strict sequential order. The first agent acts according to his individual signal strategy, while all subsequent agents update their prior belief based upon the public history of actions and then form their posterior belief given their signal using Bayes rule. As another example, consider the literature on repeated interaction in social networks, where each agent takes an action in every of infinitely many rounds. In the first round all agents follow the individual signal strategy. From the second round onward, instead, they update their belief based upon the previous actions taken by their neighbors and select their expected utility maximizing action based upon the computed posterior distribution.

We now introduce our concept of boundedly rational observational learning. A Quasi-Bayesian updating function assigns an action to each possible action vector. The central idea of the Quasi-

¹⁶Formally, we assume that both Ω and S are Polish spaces with respective Borel σ -algebra.

¹⁷We assume that the mapping F that assigns a probability distribution over S to each state of the world generates a Markov Kernel.

Bayesian approach is to abstract away from the structure of interaction and from the indirect inferences the actions of other agents might be based upon.

Definition 1 *For an individual signal strategy profile σ^s , a **Quasi-Bayesian updating function** $\beta : A^n \times 2^N \rightarrow A$ satisfies the following properties*

1. $\beta(\mathbf{a}, N') = \beta(\mathbf{a}', N')$ for all \mathbf{a}' such that $\mathbf{a}'_{N'} = \mathbf{a}_{N'}$, and
2. $\beta(\mathbf{a}, N') = \arg \max_{a \in A} E[u(a, \omega) | \mathbf{a}_{N'}^s]$ where $\mathbf{a}_{N'}^s = \mathbf{a}_{N'}$.

A Quasi-Bayesian updating function treats the observed actions of a group of individuals as if each action were based only on the private signal realization of each agent. That is, for an observed action vector $\mathbf{a}_{N'}$, the observed action of each agent i in N' is treated as if it resulted from the individual signal strategy σ_i^s . This approach departs from Bayesian learning and, in doing so, drastically reduces the complexity of observational learning.¹⁸

The crucially appealing feature of Quasi-Bayesian updating functions is that they are generally applicable to any environment of observational learning, whether in sequential or repeated interaction settings, and that they have a limited rational foundation.¹⁹ Naturally, the specific functional form of the Quasi-Bayesian updating function as well as its general properties vary among environments. In order to further the understanding of Quasi-Bayesian updating functions we analyze their implications for aggregate behavior in a model of repeated interaction in social networks.

We are aware that the proposed Quasi-Bayesian approach needs to be assessed not only in terms of its theoretical appeal but also in terms of its empirical relevance. In Section 5 we report evidence from a laboratory experiment which supports the relevance of Quasi-Bayesian updating.

3 Quasi-Bayesian updating and repeated interaction in social networks

A finite set of agents V is organized in a strongly connected graph $G = (V, E)$.²⁰ A graph is a pair of sets (V, E) such that $E \subset [V]^2$. The elements of V are nodes of the graph, representing

¹⁸For example, observational learning in a network environment might require to draw indirect inferences on the private information of unobserved neighbors of neighbors from the observed actions of neighbors.

¹⁹Note that in general a Quasi-Bayesian function need not be unique for a given environment, in particular if the image of the individual signal strategies is not equal to A . This paper considers only settings where the image of the individual signal strategies is equal to A .

²⁰In the following the terms graph and network are used interchangeably.

the agents, and the elements of E are the edges of the graph, representing the direct connections between agents. Let $N_i(G)$ denote the neighborhood of agent i in network G

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

A directed graph is strongly connected if there exists a directed path among every pair of agents.

All agents face uncertainty exactly as described in the general model of Section 2. That is, they share a common prior p over (Ω, \mathcal{F}) . The state of the world is drawn according to the prior in time $t = 0$ and all agents observe a private signal $s_i \in S$ that is independently drawn according to $F_\omega \in \Delta(\Omega)$, at the beginning of round $t = 1$.²¹ In each round $t \in \mathbb{N}$ every agent i takes an action $a_i^t \in A$ where A is a compact metrizable set. The vector of actions taken in period t is denoted as $\mathbf{a}^t \in A^v$. All agents share identical preferences represented by a stage utility function $u : A \times \Omega \rightarrow \mathbb{R}$. Each agent i follows his individual signal strategy in the first round, i.e. his first period action maximizes his expected utility conditional on his signal. Each agent observes only the history of actions of each of his neighbors. Stage utility realizations are not observed.

The above model has been analyzed for the case of Bayesian agents who draw fully rational inferences on the private information of all agents based on the history of actions they observe.²² As we argue in the introduction, a fully Bayesian model is a useful benchmark but not a good representation of the real world, due to the complexity of indirect inferences in large networks. We now introduce an alternative approach to action updating based upon our definition of Quasi-Bayesian updating functions. Following the non-Bayesian literature on learning in networks, we restrict attention to local Markov updating functions.²³ To be more precise, let \mathcal{G} denote the set of all strongly connected graphs on the set of nodes V . The Markov updating function $f_i : A^v \times \mathcal{G} \rightarrow A$ of agent i assigns an action a_i^t to each possible pair of action vector \mathbf{a}^{t-1} and network structure G . Under a local updating function the updated action of agent i depends only on the last round actions of his neighbors and himself and not also on the structure of the network as a whole.²⁴ Formally, the updating function f_i is local if, for all pairs $(\mathbf{a}, G), (\mathbf{a}', G')$ such that $N_i(G) = N_i(G')$ and $\mathbf{a}_j = \mathbf{a}'_j$ for all $j \in N_i(G) \cup \{i\}$, we have $f_i(\mathbf{a}, G) = f_i(\mathbf{a}', G')$. For expositional purposes the term ‘local’ is omitted in the remainder of the paper.

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

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

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

²¹Where $F_\omega, F_{\omega'}$ are absolutely continuous with respect to each other but not identical.

²²See for example Gale and Kariv (2003), Rosenberg, Solan and Vieille (2009), and Mueller-Frank (2013a).

²³See for example DeMarzo, Vayanos and Zwiebel (2003), Golub and Jackson (2010), and Mueller-Frank (2013b)

²⁴This condition is naturally satisfied under the assumption that agents know only the identity of their neighbors, and not the structure of the network as a whole, which is the case in our experiment.

for all $t = 2, 3, \dots$

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

Hence, the process $\{\mathbf{a}^t\}_{t \in \mathbb{N}}$ is a (deterministic) stationary Markov process. A Quasi-Bayesian Markov updating function is defined as follows.²⁵

Definition 2 *A local Markov updating function $f_i : A^v \times \mathcal{G} \rightarrow A$ is **Quasi-Bayesian** if for all $(\mathbf{a}, G) \in A^v \times \mathcal{G}$ we have*

$$f_i(\mathbf{a}, G) = \beta(\mathbf{a}, N_i(G) \cup i).$$

Under a Quasi-Bayesian updating function, an agent updates his action in every round in a Bayesian manner, but assumes that the last round actions of each of his neighbors and himself are based only on the respective private signal realization. Hence, the Quasi-Bayesian updating function reduces the complexity of indirect inferences by treating the action vector in each round as if it were the first.

In Section 4 we present the design of a laboratory experiment that fits within the network-based observational learning model described above. In Section 5 we report evidence of Quasi-Bayesian updating among the participants in the experiment.

4 Experimental design

The experimental task consists of an urn-guessing game played by subjects connected by a network structure. The design is adapted from the urn-guessing experiment, which is standard in the experimental literature on observational learning, and was previously applied both to sequential social learning (Anderson and Holt 1997, Hung and Plott 2001, and Kübler and Weizsäcker 2004) and to social learning in networks (Choi, Gale and Kariv 2005, and Grimm and Mengel 2014). The instructions are reported in Appendix B.

We conducted three sessions with networks with 5 agents and three sessions with networks with 7 agents. In total, 198 subjects participated: 100 in 5-agent networks and 98 in 7-agent networks. Each subject participated in only one session.²⁶ The experiment starts with the computer randomly

²⁵To define such Quasi-Bayesian Markov updating functions, the individual signal strategies of agents need to be known, or the posterior distributions need to induce unique optimal actions.

²⁶The three 5-agent-network sessions had 35, 30, and 35 participants, respectively. The three 7-agent-network sessions had 35, 28, and 35 participants, respectively. The sessions were conducted in December 2012 at the Decision Science Lab at ETH Zürich. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). Recruitment was implemented using the Online Recruitment System ORSEE (Greiner 2004) among the university student community and achieved a number of participants per session equal to either a multiple of 5 or a multiple of 7. The Decision Science Lab at ETH Zürich can accommodate up to 36 participants.

assigning a label to each participant in the laboratory. The labels are $\{A, B, C, D, E\}$ in the 5-agent-network treatment and $\{A, B, C, D, E, F, G\}$ in the 7-agent-network treatment. An equal number of participants is assigned to each label. Then the computer randomly forms groups such that each member of a group has a different label. Each participant’s label and the composition of each group remained constant throughout the experiment. In the course of the experiment, participants become informed (via a table displayed on the computer screen) of the choices made by some or all members of their group; if agent i is informed of the choice of agent j , then j is informed of i ’s choice.²⁷ In other words, i and j are connected (i.e. neighbors). The choices of non-connected group members are displayed instead as N/A (not available).

Within each group, participants play an urn-guessing game. The game is played with 2 or 4 urns containing colored balls, and each participant is always informed of whether the 2-urn game or the 4-urn game is played. Figure B.1 depicts the urns and their composition in both the 2-urn game and the 4-urn game. In the 2-urn game, there are a White Urn and a Black Urn, each containing 3 colored balls. The White Urn has 2 white and 1 black balls. The Black Urn has 2 black and 1 white balls. In the 4-urn game, the urns are Red, Yellow, Green and Blue. Each urn contains 7 colored balls. The Red Urn has 4 red, 1 yellow, 1 green, and 1 blue balls. The Yellow Urn has 4 yellow, 1 red, 1 green, and 1 blue balls. The Green Urn has 4 green, 1 red, 1 yellow, and 1 blue balls. The Blue Urn has 4 blue, 1 red, 1 yellow, and 1 green balls. The 2-urn game and the 4-urn game follow the same rules and differ only in the number and composition of the urns.

The computer starts by randomly selecting an urn, each of which is equally likely to be selected. Group members are not informed of which urn is selected; however, they each receive a piece of private information. Independently for each member, the computer draws a ball from the selected urn and informs each member about the color of their particular ball. Private draws are done with replacement, keeping the composition of balls in the urn constant. The game then proceeds over several decision rounds, with the urn selected by the computer and each member’s private draw staying the same. In the 1st round, after learning the color of one’s own privately drawn ball, each participant is asked to indicate the urn that he thinks is more likely to have been used.²⁸ Next, as the 2nd round starts, each participant is informed of the choices made in the previous round by the agents to whom he is connected, after which he is asked to indicate again the urn that he thinks is more likely to have been used. This process is repeated for 6 rounds in 5-agent networks and for 8 rounds in 7-agent networks.²⁹

²⁷This is explained to participants in the instructions.

²⁸Participants need to choose one and only one urn, i.e. they cannot express probabilistic beliefs over all existing urns.

²⁹Choi, Gale and Kariv (2005) implement 3-agent networks, binary choice set, and repeated choices over 6 rounds. Chandrasekhar, Larreguy and Xandri (2012) implement 7-agent networks, binary choice set, and repeated choices over a random number of rounds (on average 6). Grimm and Mengel (2014) implement 7-agent networks, binary choice set, and repeated choices over 20 rounds. As Table C.2 shows, in our experiment participants update their

Each group plays in a series of independent periods, and each period corresponds to an independent game, i.e. the initial random selection of an urn is independent across periods. Participants in 5-agent networks played 18 periods, while participants in 7-agent networks played 14 periods.³⁰ As the experiment proceeds, each participant can observe on the screen a table which reports, for each previous game (period), whether his choice in each decision round was correct.³¹ The set of periods is divided in two halves: the 2-urn game is played in one half and the 4-urn game in the other half. The order is determined randomly by the computer at the beginning of the session and is communicated to the participants.

While the matching of participants into groups does not change across periods, connections among group members (i.e. the number of connections and who is connected to whom) do. Each participant is informed about his own connections to other agents at the beginning of each period, but not about the connections among other agents. Participants, therefore, do not know the entire network structure, only the links connecting them to their own neighbors, a key feature of both our experimental design and our theoretical model.³² Many real-world social networks (e.g. Facebook) reflect this feature. When each individual has a large number of connections and the network structure is highly complex, it is unrealistic to assume that individuals consider the structure of the entire network in their decision-making.³³

Table C.1 reports the connections implemented for each group member.³⁴ In 5-agent networks, network structures included a complete network, three star networks, and five linked-circle networks, with each subject thus having between 1 and 4 neighbors.³⁵ In 7-agent networks, network structures

choices few times and mostly in early rounds. This evidence justifies the relatively low number of rounds implemented in our design.

³⁰The numbers of periods and rounds per-period were chosen in order to allow the participants of larger networks to experience more decision rounds per period, as information may take longer to spread in larger networks, while keeping the total number of decision rounds and the duration of the experiment similar across all sessions.

³¹The table simply displays ‘correct’ or ‘incorrect’, without reference to the color of the correct urn.

³²Choi, Gale and Kariv (2005) and Chandrasekhar, Larreguy and Xandri (2012) implement common knowledge of the network structure. Grimm and Mengel (2014) include treatments with complete, incomplete and no information of network structure. The DeGroot model is found to outperform the Bayesian model at explaining individual behavior, whether agents have complete information about the network structure or not. Chandrasekhar, Larreguy, and Xandri (2012) find that the uniform-weighting DeGroot model performs the best, and that both the uniform- and the degree-weighting DeGroot model outperform the Bayesian model. Grimm and Mengel (2014) find that the uniform-weighting DeGroot model outperforms the Bayesian model. Hence, the assumption of common knowledge of network structure does not seem to be crucial.

³³The typical (median) teenage Facebook user has 300 friends, according to a recent Pew Research Center study (2013).

³⁴Each network structure reported in Table C.1 was implemented twice, once in a 2-urn game and once in a 4-urn game. Figure C.1 offers a visual representation of the connections. No visualization was provided during the experiment.

³⁵For 5-agent networks, in a complete network, each subject has 4 neighbors; in a star network, one subject has 4 neighbors and other subjects have 1; in a linked-circle network, two subjects have 3 neighbors and other subjects

consisted of six linked-circle networks and one connected complete-components network, with each subject having between 2 and 4 neighbors.³⁶ We chose networks of 5 or 7 subjects to study a richer set of structures compared to what would be possible with less subjects.³⁷ The specific network structures were chosen because they either represent benchmark situations (complete and star network) or allow different nodes in the network to have different characteristics, such as degree and clustering coefficient (linked circle and connected complete components)³⁸. To sum up, within the experiment we vary network size (5 or 7 agents), choice-set size (2 or 4 urns), and network connections.³⁹ Table 1 summarizes the experimental design.

Table 1: Participants

network size	No. subjects	groups	rounds/period	periods	obs.
5	100	20	6	18	10800
7	98	14	8	14	10976
all	198	34			21776

Participants were paid for their performance in every period, but only for one randomly selected round in each period.⁴⁰ At the end of the experiment, the computer randomly selected one round for each period. For each period, participants received CHF 2 (in sessions with 5-subject networks) or CHF 2.5 (in sessions with 7-subject networks) if their decision in the selected round was correct.⁴¹

have 2.

³⁶These networks exhibit small-world characteristics, having a relatively small diameter, a short average path length, and high clustering (compared to an independent random network). In 7-agent networks, network structures included: four networks in which four subjects have 3 and other subjects have 2 neighbors; two networks in which two subjects have 3 and other subjects have 2 neighbors; one network in which one subject has 4, four subjects have 3, and two subjects have 2 neighbors.

³⁷For example, consider the 3-agent networks in Choi, Gale and Kariv (2005).

³⁸The degree of a node is the number of neighbors that the node has. The clustering coefficient of a node is the fraction of pairs of a node's neighbors who are neighbors to each other. If a node has only one neighbor, the clustering coefficient is set to 0. For the 5-agent networks we implement, in the complete network all subjects have a clustering coefficient equal to 1, in the star networks all subjects have a clustering coefficient equal to 0, in the linked circle networks the clustering coefficient equals 0, 1/3 or 1. Also for the 7-agent networks we implement, the clustering coefficient ranges between 0, 1/3 and 1.

³⁹The design is between-subject for the network size and within-subject for the choice-set size and the network connections. The sample of 198 participants is divided into two network size treatments, with approximately 100 subjects in each cell of a between-subjects design. Within each cell, choice-set size and network connections are varied in random order. Then, we have 198 observations in a within design with order effects controlled for, and two between comparisons with approximately 100 observation in each cell. For a discussion on between- and within-subject design, see Charness, Gneezy and Kuhn (2012).

⁴⁰Each session included an initial practice period, which did not count for the determination of earnings nor was included in the analysis.

⁴¹The difference in payment per correct-decision (CHF 2 versus CHF 2.5) was gauged to attain similar expected earnings for subjects irrespective of the session they participated in, in order to comply with the ETH Decision

Otherwise they received nothing. Participants’ earnings ranged between CHF 22 and CHF 42.5, with an average of CHF 34 (including a CHF 10 show-up fee).⁴²

5 Experimental evidence: Quasi-Bayesian updating

In this section we analyze the experimental data in the light of the theoretical framework presented in Section 3. We begin by describing agents’ behavior in the first round. As one might expect, in the first round agents choose an action based on their private signal. Table 2 reports the distribution across individuals of the frequency with which first-round choices coincide with private signals. The mean is 0.97 in 5-agent networks and 0.95 in 7-agent networks, and the median is 1 for both network sizes.

Table 2: Frequency of 1st-round choice equal to private signal. Distribution across individuals.

	obs	mean	median	std
5-agent networks				
all games	100	0.97	1	0.10
2-urn games	100	0.97	1	0.11
4-urn games	100	0.97	1	0.12
7-agent networks				
all games	98	0.95	1	0.13
2-urn games	98	0.94	1	0.14
4-urn games	98	0.96	1	0.13

We then inspect agents’ behavior in rounds following the first, when agents have the opportunity to revise their choice after observing their neighbors’ choices. Table 3 reports the percentage of observations (across rounds $t = 2, 3, \dots$) in which choice behavior is consistent with Quasi-Bayesian updating. At least 90 percent of observations across network size, network types, and choice sets, are consistent with Quasi-Bayesian updating.

Figure 1 shows the empirical distribution of the individual rates at which participants’ choice behavior is consistent with Quasi-Bayesian updating. For each participant, the individual score of Quasi-Bayesian updating is defined as the fraction of choices that satisfy Definition 2. Median individual rates of consistency are extremely high both in 5-agent networks (0.96 and 0.96 in 2- and 4-urn games, respectively) and in 7-agent networks (0.96 and 0.92 in 2- and 4-urn games, respectively).⁴³

Science Lab rules.


⁴²In sessions with 5-subject networks, the range was CHF 22-42 (average CHF 35). In sessions with 7-subject networks, the range was CHF 22.5-42.5 (average CHF 33).

⁴³Mean individual scores are extremely high both in 5-agent networks (0.93 and 0.91 in 2- and 4-urn games, respectively) and in 7-agent networks (0.91 and 0.89 in 2- and 4-urn games, respectively). Kolmogorov-Smirnov test finds no significant difference in the distribution of individual scores in 2-urn games (5-agent networks versus 7-agent


Table 3: Percentage of observations in which choice behavior is consistent with Quasi-Bayesian updating.

	5-agent networks			7-agent networks	
	star	linked circle	complete	small world	connected complete
2-urn	0.95	0.92	0.91	0.91	0.90
4-urn	0.94	0.90	0.90	0.89	0.91


Figure 1: Empirical distribution of participants' individual consistency with Quasi-Bayesian updating. A kernel density estimate is also reported, using an Epanechnikov kernel function with optimal half-width.




(a) 2-urn games, 5-agent networks



(b) 2-urn games, 7-agent networks



(c) 4-urn games, 5-agent networks



(d) 4-urn games, 7-agent networks

networks $p = 0.408$) but finds a significant difference in 4-urn games (5-agent networks versus 7-agent networks $p = 0.016$).

6 Long-run properties of Quasi-Bayesian updating in social networks

The two main questions addressed by the literature on learning in social networks concern the conditions on the environment and the network structure under which consensus (i.e. asymptotic agreement in actions) and information aggregation (i.e. optimality of long-run actions conditional on the pooled private information of all agents) occur.⁴⁴ We address both questions within the framework of Quasi-Bayesian updating presented in Section 3. To do so, let us formally define consensus and information aggregation. Let $\mathbf{a}^* \in A^V$ denote a vector such that $\mathbf{a}_i^* = a^*$ for all $i \in V$ and some action a^* .

Definition 3 *An updating system \mathbf{f} yields **consensus** if for every $(\mathbf{a}^1, G) \in A^V \times \mathcal{G}$ there exists an action a^* and a time t^* such that $\mathbf{a}^t = \mathbf{a}^*$ for all $t \geq t^*$.*

Thus an updating system yields consensus, if consensus is reached in finite time in every strongly connected network, and persists from that time onward. Our definition requires consensus to occur in all strongly connected networks for any initial action profile. Therefore, we are concerned with the class of updating functions that yield consensus independently of the actual structure of the strongly connected network. While a strong requirement, it mirrors the approach to (asymptotic) consensus of the literature on non-Bayesian and Bayesian learning in networks, where consensus refers also to consensus in all strongly connected networks.

Next we define our notion of information aggregation. It is useful to note that a deterministic updating system yields a unique process of actions $\{\mathbf{a}^t\}_{t \in \mathbb{N}}$ for each initial action vector \mathbf{a}^1 . Therefore, the best possible information aggregation outcome is one in which all agents, from some time onward, select an action that is optimal conditional on the first round action vector. In settings with a rich signal space this notion of information aggregation is typically weaker than one which requires the long-run actions of all agents to be optimal conditional on the realized signals of all agents.

Definition 4 *A updating system \mathbf{f} yields **information aggregation** if for every $(\mathbf{a}^1, G) \in A^V \times \mathcal{G}$ there exists a time t^* such that for all agents $i \in V$ and times $t \geq t^*$ we have*

$$a_i^t \in \arg \max_{a \in A} E[u(a, \omega) | \mathbf{a}^1].$$

⁴⁴For an analysis of consensus see DeMarzo, Vayanos and Zwiebel (2003) and Mueller-Frank (2014) for non-Bayesian models, and Gale and Kariv (2003), Rosenberg, Solan and Vieille (2009), and Mueller-Frank (2013a) for Bayesian models. For an analysis of learning see Golub and Jackson (2010), Mueller-Frank (2013a, 2014), Arieli and Mueller-Frank (2014), and Mossel, Sly and Tamuz (2015).

As in the definition of consensus, we require information aggregation to occur in any strongly connected network. This approach mirrors the one taken in the literature on Bayesian learning in social networks. Mueller-Frank (2013a) and Arieli and Mueller-Frank (2014) provide sufficient conditions such that perfect information aggregation generically occurs in any strongly connected network.

Compared to the environment in which Quasi-Bayesian updating in networks was introduced in Section 3, we impose three additional conditions. First, we restrict attention to the case of binary actions. This is a standard approach in the social learning literature.⁴⁵ The restriction to binary actions allows us to make a first step towards better understanding the long run properties of boundedly rational repeated interaction in environments with finitely many actions. Previous theoretical studies of boundedly rational learning restricted attention to the DeGroot model or extensions thereof, all of which require the action space to be infinite. While many real world applications of repeated interaction in networks are well modeled with an infinite action space, many others are clearly not. For such finite environments there is no understanding in the non-Bayesian learning literature of the long run properties of repeated interaction. Second, we assume that reverse symmetry is satisfied.

Definition 5 *The utility function $u : A \times \Omega \rightarrow \mathbb{R}$ satisfies **reverse symmetry** if for any $\omega, \hat{\omega}$ such that $u(a', \omega) > u(a', \hat{\omega})$ we have $u(a'', \omega) < u(a'', \hat{\omega})$.*

In words, reverse symmetry states that if action a' achieves strictly higher utility in state ω than in state $\hat{\omega}$, then a'' achieves strictly higher utility in state $\hat{\omega}$ than in state ω . For example, reverse symmetry holds in the setting commonly considered in the literature on social learning, which features binary states and binary actions, $A = \Omega = \{0, 1\}$. Here agents achieve a utility of 1 if their action matches the state of the world, and a utility of 0 otherwise. Reverse symmetry is satisfied as

$$\begin{aligned} 1 &= u(1, 1) > u(1, 0) = 0 \\ 0 &= u(0, 1) < u(0, 0) = 1. \end{aligned}$$

Note that for every state $\omega \in \Omega$, the state conditional signal distribution and the strategy of agent i induce a probability distribution over the set of actions $A = \{0, 1\}$. To simplify the analysis we make the following assumption on the signal generating distributions in relation to the utility function. Let S_I be the set of signals such that

$$S_I = \left\{ s \in S : \int u(0, \omega) \Pr(d\omega | s) = \int u(1, \omega) \Pr(d\omega | s) \right\}.$$

⁴⁵For models that restrict attention to binary states and actions, see Smith and Sorensen (2000) and Acemoglu, Dahleh, Lobel and Ozdaglar (2011) for sequential social learning, and Mossel, Sly and Tamuz (2012) for Bayesian learning in social networks.

We assume that the state dependent signal generating distribution F_ω assigns probability zero to S_I , for every state ω .⁴⁶ Third, we impose a monotone likelihood property on the environment.

Definition 6 *The environment satisfies the **monotone likelihood property** if for any $\omega, \hat{\omega}$ and $a \in A$ such that $u(a, \omega) > u(a, \hat{\omega})$ then*

$$\Pr(a|\omega) > \Pr(a|\hat{\omega}).$$

This property states that if a achieves higher utility in state ω than in state $\hat{\omega}$, then the probability of action a being taken in state ω is higher than in state $\hat{\omega}$. In order for the analysis of consensus and information aggregation to be non-trivial, we assume that for each action there exists a positive prior probability set of states such that either action is uniquely optimal under the utility function u . Then, monotone likelihood, together with the direct assumption on the signal generating distributions, implies that with positive probability in $\Omega \times S^v$ either action is uniquely optimal in expected utility terms, conditional on the realized signal.

Note that the environment we consider is a generalization of the standard model considered in the social learning literature which features binary actions and satisfies reverse symmetry and the monotone likelihood property.⁴⁷ In the network literature, Mossel, Sly and Tamuz (2012) analyze repeated interaction of Bayesian agents in social networks with an informational structure and utility function as in the standard social learning model. The main difference between our approach and the above-mentioned papers is that we allow for a general state space rather than restricting attention to the binary case. Theorem 1 provides a characterization of consensus. Recall that \mathbf{a}'' denotes a vector with all entries equal to a'' .

Theorem 1 *Consider an environment that satisfies reverse symmetry and monotone likelihood. Consensus occurs if and only if there exists an action a' such that for all $\mathbf{a} \neq \mathbf{a}'', a'' \neq a'$,*

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

and the Quasi-Bayesian updating functions of all agents is identical and satisfies $\beta(\mathbf{a}, N') = a'$ for all $N' \subseteq N$ and $\mathbf{a}_{N'} \neq \mathbf{a}_{N'}''$.

Theorem 1 provides a necessary and sufficient condition on the environment such that consensus occurs in all strongly connected networks, and for all initial action vectors \mathbf{a}^1 . This condition

⁴⁶Note that if $F_\omega(S_I) = 0$ for some ω , then it follows from equivalence of measures that $F_{\omega'}(S_I) = 0$ for all ω' . This assumption makes the distribution over actions the same for all agents as the strategies assign the same action with probability one.

⁴⁷See for example Bikhchandani, Hirshleifer and Welch (1992), Smith and Sorensen (2000) and Acemoglu, Dahleh, Lobel and Ozdaglar (2011). Lemma 1 in A.4 establishes that the social learning model satisfies monotone likelihood.

requires the Quasi-Bayesian updating function to assign one action, say a' , for all possible observation vectors other than consensus in the other action a'' . For the condition to be satisfied the environment needs to be highly asymmetric and as such the condition is very strong.⁴⁸ Theorem 2 provides a characterization of information aggregation.

Theorem 2 *Consider an environment that satisfies reverse symmetry and monotone likelihood. Information aggregation occurs if and only if there exists an action a' such that for all $\mathbf{a} \neq \mathbf{a}''$, $a'' \neq a'$,*

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

Very similar to Theorem 1, our second theorem provides a necessary and sufficient condition for information aggregation. The necessary condition for consensus and information aggregation is identical. Instead, the sufficient condition of Theorem 1 is stronger. The sufficient conditions are identical only in the case where action a' is uniquely optimal conditional on all action vectors but the consensus vector in a'' . In case there is indifference among both actions conditional on a vector where all but one agent select a'' , then, in order to achieve consensus, the Quasi-Bayesian updating functions of all agents need to select action a' in the indifference case.

We emphasize two main differences between our results on consensus and information aggregation for binary action spaces and the literature. First, consensus occurs in all strongly connected networks under weak conditions both in the DeGroot model and in the Bayesian model. For example, in the weighted-average DeGroot model the updating functions of agents need neither be identical nor constant over time for consensus to occur.⁴⁹ In our setting with binary actions and Quasi-Bayesian updating, however, consensus requires a highly asymmetric environment and is, as a result, harder to achieve. Second, under Quasi-Bayesian updating information aggregation might occur in finite networks. This result markedly differs from the existing results on non-Bayesian learning where information aggregation either requires infinite networks (Golub and Jackson 2010) or holds only for specific network structures (DeMarzo, Vayanos and Zwiebel 2003). Finally, under Quasi-Bayesian updating information aggregation and consensus generally coincide, while the same is not true in the DeGroot model in finite networks nor in the Bayesian model with binary actions in finite networks.

To conclude our analysis of Quasi-Bayesian updating in social networks, we contrast our result to the naive learning result established by Golub and Jackson (2010). They consider DeGroot updating in an environment with conditional i.i.d. signals. They show that the eventual consensus opinion of agents converges to the true state as the network size grows to infinity, if the influence

⁴⁸See Appendix A.5 for an example.

⁴⁹See DeMarzo, Vayanos and Zwiebel (2003).

of the most influential agent converges to zero. Theorem 3 establishes that their result does not carry forward to Quasi-Bayesian updating with binary actions.

Theorem 3 *Consider an environment that satisfies the monotone likelihood property and reverse symmetry. Then there exists a finite network size n^* such that consensus and information aggregation fail for all network sizes greater than n^* .*

7 Experimental evidence: consensus and information aggregation

In this section we analyze the experimental data in light of the long-run properties of Quasi-Bayesian updating presented in Section 6: consensus and information aggregation. Table 4 reports the fraction of groups in 5-agent networks and 7-agent networks that, by the last updating round, reach consensus. Table 4 also reports separately (in italics) the fraction of groups that reach consensus with information aggregation. Recall that consensus with information aggregation means that all agents choose an action that is optimal conditional on the first round actions of all agents. The results are aggregated according to the type of network structure (star, linked-circle and complete networks for 5-agent networks, and small-world and connected complete networks for 7-agent networks). We highlight two main results.

		5-agent networks			7-agent networks	
		star	linked circle	complete	small world	connected complete
2-urn	no consensus	0.52	0.41	0.30	0.71	0.50
	consensus	0.48	0.59	0.70	0.29	0.50
	<i>of which with information aggregation</i>	0.97	0.98	1	0.96	1
4-urn	no consensus	0.57	0.53	0.35	0.64	0.71
	consensus	0.43	0.47	0.65	0.36	0.29
	<i>of which with information aggregation</i>	0.96	0.96	1	1	1

Table 4: Fraction of groups reaching consensus by the last updating round. Averages across 20 groups for 5-agent networks (100 participants) and across 14 groups for 7-agent networks (98 participants).

First, consensus often fails to arise. The fraction of groups not reaching consensus ranges between 0.30 and 0.57 in 5-agent networks, and between 0.5 and 0.71 in 7-agent networks.⁵⁰ Second,

⁵⁰In 5-agent networks, consensus occurs more often in complete networks and less so in linked-circle and star networks. In 7-agent networks, however, there is no evidence of differences across the employed network structures.

when consensus does occur, it occurs almost always with information aggregation, which is in line with our theoretical result. Among the groups reaching consensus, the fraction of those achieving information aggregation ranges between 0.96 and 1 both in 5-agent and 7-agent networks. Thus, the empirical evidence is consistent with the results of Theorem 1 and 2: in an environment that satisfies the monotone likelihood property and reverse symmetry, Quasi-Bayesian updating achieves information aggregation if it achieves consensus.

Our results are similar to those in Grimm and Mengel (2014), who find that in 7-agent networks with no information of the network structure the fraction of groups reaching consensus ranges between 0.16 and 0.66, depending on the network structure.⁵¹ Our results instead differ from those in Choi, Gale and Kariv (2005), who find that in 3-agent networks under common knowledge of network structure the fraction of groups reaching consensus is 0.71.⁵² We interpret their result as following from a smaller sized network with common knowledge of network structure, which does not generalize to larger networks with lack of common knowledge of the network structure.

8 Conclusion

In this paper we propose a new general model of boundedly rational observational learning, which is based on the concept of Quasi-Bayesian updating. The Quasi-Bayesian approach addresses the weaknesses of the Bayesian and the DeGroot models. It reduces the complexity of the inferences drawn by decision makers, it is applicable to any environment of observational learning, and it is not arbitrary but is instead rationally founded. We also consider the theoretical long run implications of Quasi-Bayesian updating in a model of repeated interaction in social networks, investigating under which conditions Quasi-Bayesian updating yields consensus and information aggregation. We combine our theoretical model with the analysis of data collected in a laboratory experiment. We find that Quasi-Bayesian updating is strongly supported by the data and that the occurrence of consensus and information aggregation is in line with our theoretical predictions.

We are aware that this paper, by focusing on the analysis of Quasi-Bayesian updating in a

⁵¹Grimm and Mengel (2014) implement a 3x3 design varying the network structure (circle, star and kite) and the amount of information about the network structure available to players (no information, incomplete information, or complete information). In our comparison, we consider the circle and the star network structure treatment and the no information treatment, as these treatments closely reflect our experimental design.

⁵²They implement three network structure treatments (complete, star and circle), and three information treatments, full, high, and low, in which agents receive an informative private signal with probability 1, 2/3, or 1/3, respectively. Our comparison is limited to the full information treatment. We are aware that the comparison is weakened by the difference in experimental design. In our design all participants play in all network structures and with all numbers of urns, but the composition of groups never changes. In Choi, Gale and Kariv (2005) participants instead play in groups whose composition changes randomly in each period, but in network structures and with number of urns which never change.

model of repeated interaction in social networks, does not offer a conclusive answer to whether Quasi-Bayesian updating provides a good description of behavior in other observational learning environments and in more complex settings. We certainly consider an analogous analysis of Quasi-Bayesian updating in a sequential social learning model an interesting topic for further research.

References

- [1] Acemoglu, D., Dahleh, M.A., Lobel, I., and Ozdaglar, A. (2011). Bayesian Learning in Social Networks. *Review of Economic Studies*, 78, 1201-1236.
- [2] Alatas, V., Banerjee, A., Chandrasekhar, A., Hanna, R., Olken, B. A. (2014). Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia. *working paper*.
- [3] Anderson, L. R., and Holt, C.A. (1997). Informational Cascades in the Laboratory. *American Economic Review*, 87, 847-862.
- [4] Arieli, I., and Mueller-Frank, M. (2014). Inferring Beliefs from Actions. *working paper*.
- [5] Bala, V., and Goyal, S. (1998). Learning from Neighbours. *Review of Economic Studies*, 65, 595-621.
- [6] Banerjee, A.V. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 107, 797-817.
Banerjee, A.V., Chandrasekhar, A., Duflo, E., Jackson, M.O. (2013). The diffusion of microfinance, *Science*, vol. 341 (6144).
- [7] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992) A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades. *The Journal of Political Economy*, 100, 992-1026.
- [8] Blondel, V.D., Hendrickx, J.M., Tsitsiklis, J.N. (2009). On Krause’s Multi-Agent Consensus Model with State Dependent Connectivity. *IEEE Transactions on Automatic Control*, 54, 2586-2597.
- [9] Chandrasekhar, A.G., Larreguy, H., and Xandri, J.P. (2012) Testing models of social learning on networks: evidence from a framed field experiment. *working paper*.
- [10] Charness, G., Gneezy, U., and Kuhn, M.A. (2012) Experimental Methods: Between-Subject and Within-Subject Design. *Journal of Economic Behavior and Organization*, 81, 1-8.
- [11] Choi, S., Gale, D., and Kariv, S. (2005). Behavioral Aspects of Learning in Social Networks: An Experimental Study. *Advances in Applied Microeconomics: A Research Annual*, 13, 25-61.

- [12] DeGroot, M.H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69, 118-121.
- [13] DeMarzo, P.M., Vayanos, D., and Zwiebel, J. (2003). Persuasion Bias, Social Influence and Uni-Dimensional Opinions. *Quarterly Journal of Economics*, 118, 909-968.
- [14] Fischbacher, U. (2007) Z-tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10, 171-178.
- [15] Gale, D., and Kariv, S. (2003). Bayesian learning in social networks. *Games and Economic Behavior*, 45, 329-346.
- [16] Golub, B., and Jackson, M.O. (2010). Naive Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2, 112-149.
- [17] Goyal, S. (2014). Networks in Economics: A Perspective on the Literature. *working paper*.
- [18] Greiner, B. (2004) The Online Recruitment System ORSEE 2.0 - a Guide for the Organization of Experiments in Economics. *working paper*.
- [19] Grimm, V., and Mengel, F. (2014). Experiments on Belief Formation in Networks. *working paper*.
- [20] Hung, A.A., and Plott, C.R. (2001). Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions. *American Economic Review*, 91, 1508-1520.
- [21] Kübler, D., and Weizsäcker, G. (2004). Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory. *Review of Economic Studies*, 71, 425-441.
- [22] Lobel, I., Ozdaglar, A., Feijer, D. (2011). Distributed Multi-Agent Optimization with State-Dependent Communication. *Mathematical Programming*, 129, 255-284.
- [23] Lobel, I., and Sadler, E. (2015). Information Diffusion in Networks through Social Learning. *Theoretical Economics*, forthcoming.
- [24] Möbius, M.M., Phan, T., and Szeidl, A. (2010). Treasure Hunt: Social Learning in the Field. *working paper*.
- [25] Mossel, E., Sly, A., and Tamuz, O. (2015). On Agreement and Learning. *working paper*.
- [26] Mossel, E., Sly, A., and Tamuz, O. (2012). Strategic Learning and the Topology of Social Networks. *working paper*.
- [27] Mueller-Frank, M. (2013a). A General Framework for Rational Learning in Social Networks. *Theoretical Economics*, 8, 1-40.

- [28] Mueller-Frank, M. (2013b) Reaching Consensus in Social Networks. *working paper*.
- [29] Mueller-Frank, M. (2014). Does one Bayesian make a Difference?. *Journal of Economic Theory*, 423-452.
- [30] Mueller-Frank, M., and Pai, M. (2015). Social learning with costly search,' *American Economic Journal: Microeconomics*, forthcoming.
- [31] Pew Research Center (2013). *Teens, Social Media, and Privacy*.
- [32] Rosenberg, D., Solan, E., and Vieille, N. (2009). Information externalities and emergence of consensus,' *Games and Economic Behavior*, 66, 979-994.
- [33] Smith, L., and Sorensen, P. (2000). Pathological Outcomes of Observational Learning. *Econometrica*, 68, 371-398.

Appendices

A Proofs

A.1 Proof of Theorem 1

Proof. First we establish three properties of Quasi-Bayesian updating that are implied by reverse symmetry and the monotone likelihood property. Consider any partition $\mathcal{P}^* = \{P_1, P_2, \dots, P_m\}$ of Ω such that

$$\min_{\omega \in P_i} u(a', \omega) > \max_{\omega \in P_{i+1}} u(a', \omega) \quad (1)$$

for all $i = 1, \dots, m-1$. The cardinality m of the partition is an arbitrary natural number $m \geq 3$ bounded above by the cardinality of Ω (if Ω is finite). Let q be any prior probability measure over Ω that is equivalent to the prior p .⁵³

1. As a first step, we show that the expected utility of action a' is strictly greater under the posterior distribution conditional on observing a' than under the prior. We establish that for any $i = 1, \dots, m-1$

$$\Pr(P_{i+1} | a', q) \geq \Pr(P_{i+1} | q) \Rightarrow \Pr(P_i | a', q) > \Pr(P_i | q).$$

and we denote this property as **Property A**. Note that

$$\Pr(P_{i+1} | a', q) \geq \Pr(P_{i+1} | q)$$

is equivalent to

$$\begin{aligned} \frac{\Pr(a' | P_{i+1}, q) \Pr(P_{i+1} | q)}{\Pr(a' | q)} &\geq \Pr(P_{i+1} | q) \\ \Pr(a' | P_{i+1}, q) &\geq \Pr(a' | q). \end{aligned}$$

But monotone likelihood and by property (1) we have $\Pr(a' | P_i, q) > \Pr(a' | P_{i+1}, q)$ establishing the validity of Property A. Note that Property A implies that for any partition \mathcal{P}^* and cell $P_i \neq P_m$ we have

$$\Pr\left(\bigcup_{j=1}^i P_j | a', q\right) > \Pr\left(\bigcup_{j=1}^i P_j | q\right).$$

This implies that the induced distribution on $u(a', \omega)$ conditional on observing a' first order stochastically dominates the induced distribution on $u(a', \omega)$ under the prior distribution q .

⁵³Two priors are equivalent if they are absolutely continuous with respect to each other.

Hence the expected utility of a' under the posterior distribution is higher than under the prior distribution. Analogous reasoning and weak symmetry implies that expected utility of action a'' is lower under the posterior conditioning on a' than under the prior q . Hence, for any prior distribution it holds that if a' is optimal under the prior q , then a' remains optimal under the posterior conditional on observing a' .

2. As a second step, we show that whenever all observed actions are identical and equal to a' , then the conditional expected utility maximizing action is a' . That is, if only a' is observed, then the Quasi-Bayesian updating function selects a' . To see this, note that the first round agent acting on his signal and the observer share the same prior. If a first round agent selects a' , it implies that he observed a signal $s \in S'$ such that the expected utility of a' under the conditional distribution given s is greater or equal to the expected utility of $a'' \neq a'$ under the conditional distribution given s , i.e.

$$\int u(a', \omega) \Pr(d\omega | s) \geq \int u(a'', \omega) \Pr(d\omega | s)$$

for all $s \in S'$. This implies

$$\int_{s \in S'} \left(\int u(a', \omega) \Pr(d\omega | s) \right) (\Pr(ds | S')) \geq \int_{s \in S'} \left(\int u(a'', \omega) \Pr(d\omega | s) \right) \Pr(ds | S').$$

Since

$$E[u(a', \omega) | a'] = \int_{s \in S'} \left(\int u(a', \omega) \Pr(d\omega | s) \right) (\Pr(ds | S'))$$

and

$$E[u(a'', \omega) | a'] = \int_{s \in S'} \left(\int u(a'', \omega) \Pr(d\omega | s) \right) \Pr(ds | S')$$

conditional on observing one action a' , the Quasi-Bayesian revision function assigns action a' . The claim then follows by induction. Suppose that the observation vector \mathbf{a} consists of n observations of a' and assume that the Quasi-Bayesian revision function assigns a' . We need to show that if the observation vector consists of $n + 1$ observations of a' , then the Quasi-Bayesian revision function assigns a' as well. The induction basis is already established. The inductive step follows from step 1 as observing one more action a' increases the expected utility of a' while decreasing the expected utility of a'' .

3. Next we establish a monotonicity property of the Quasi-Bayesian updating function. Suppose that f_i is Quasi-Bayesian. We establish the following monotonicity property: if $f_i(\mathbf{a}, G) = a'$ holds, then $f_i(\mathbf{a}', G') = a'$ holds for any (\mathbf{a}', G') such that (i) the number of by agent i observed a' actions is at least as large as under (\mathbf{a}, G) , and (ii) the number of by agent i observed a'' actions is at most as large as under (\mathbf{a}, G) . To see this, consider a graph G and an action vector \mathbf{a} such that $f_i(\mathbf{a}, G) = a'$. In particular, assume that the observed vector $\mathbf{a}_{N_i(G) \cup i}$ features k' observations of a' and k'' observations of a'' , the number of the respective observed a' and a''

actions is denoted by (k', k'') . By assumption we have that (k', k'') induces a Quasi-Bayesian choice of a' . For monotonicity we have to show that for $l', l'' \geq 0$ an observation vector with $(k' + l', k'' - l'')$ still induces a Quasi-Bayesian choice of a' . We do this in two minor steps.

- (a) Step 1 above provides the case $l'' = 0$ and $l' > 0$.
- (b) Suppose now that both l' and l'' are greater than 0. Consider an observation vector such that $(k' + l', k'' - l'')$ and the corresponding expected utilities of a' and a'' . Note that by step 2 the expected utility of a' for an observation vector $(k' + l', k'' - l'' + l)$ is decreasing in l . Similarly, the expected utility of a'' for an observation vector $(k' + l', k'' - l'' + l)$ is increasing in l . Hence, optimality of a' under $(k' + l', k'')$ implies optimality of a' under $(k' + l', k'' - l'')$.

We can now establish the claim of the theorem. We prove sufficiency for consensus first. Consider a network G and an initial action vector \mathbf{a}^1 . If \mathbf{a}^1 is a consensus vector, then by step 2 above the updated second period action of all agents remains the same. Hence consensus remains to hold in every subsequent period. Suppose instead that \mathbf{a}^1 is not a consensus vector. Denote by i^* an agent that selected a^* in the first round, $\mathbf{a}_{i^*}^1 = a^*$. Monotonicity as established in step 3 together with the assumption on a^* , implies that a^* is conditional expected utility maximizing for any observed vector $\mathbf{a}_{N_i(G) \cup i}$ with at least one entry being equal to a^* . Hence i^* selects a^* in every round $t \in \mathbb{N}$. Consider the set of agents $N_{\rightarrow i^*}^1(G)$ whose longest path to i^* is equal to one,

$$N_{\rightarrow i^*}^1(G) = \{j \in V : d_G(j, i^*) = 1\}.$$

All agents at distance one from i^* select action a^* from round $t = 2$ onward, i.e. we have $\mathbf{a}_j^t = a^*$ for all $j \in N_{\rightarrow i^*}^1(G)$ and $t \geq 2$. By the same reasoning we have $\mathbf{a}_j^t = a^*$ for all $j \in N_{\rightarrow i^*}^k(G)$ and $t \geq k + 1$. Since the graph G is strongly connected we have $\mathbf{a}_j^t = a^*$ for all $j \in V$ and $t \geq d_i^* + 1$ we d_i^* equals the longest shortest path connecting any agent j to i^* .

Finally, we establish necessity of the condition on a^* for consensus in three steps.

1. First we show that consensus, reverse symmetry and the monotone likelihood property imply that a two dimensional action vector $\mathbf{a}_{N_i(G) \cup i} = (a', a'')$ induces a unique optimal action conditional on (a', a'') , i.e.

$$E[u(a', \omega) | \mathbf{a} = (a', a'')] \neq E[u(a'', \omega) | \mathbf{a} = (a', a'')].$$

Suppose not. Then by step 1 above we have

$$\begin{aligned} E[u(a', \omega) | \mathbf{a} = (a', a', a'')] &> E[u(a'', \omega) | \mathbf{a} = (a', a', a'')] \\ E[u(a'', \omega) | \mathbf{a} = (a', a'', a'')] &> E[u(a', \omega) | \mathbf{a} = (a', a'', a'')]. \end{aligned}$$

Now consider an undirected line network and order the agents along the line. Let \mathbf{a} be such that the initial action of agents from 1 to k is a' while all remaining agents select a'' , and where $2 \leq k \leq v - 2$. But then agent k updates his action to a' while agent $k + 1$ updates his action to a'' . All other agents observe only identical actions and hence remain with their first period action. Hence the second period action vector is identical to the first and consensus fails. In the following, let a' be the action that is preferred conditional on observing (a', a'') .

2. Now suppose that the condition on a^* fails. This implies that there exists some non-consensus action vector $\hat{\mathbf{a}} \in \{0, 1\}^v$ such that $f_i(\hat{\mathbf{a}}, G) = a''$ and where $N_i(G) = V \setminus \{i\}$. Let k' be the number of a' actions under $\hat{\mathbf{a}}$. Consider a network G' where $v - k'$ agents form a complete subgraph, that is they all directly link to each other, while each of the remaining k' agents has one undirected edge towards agent \tilde{i} who belongs to the complete subgraph. Hence agent \tilde{i} observes every agent and is observed by every agent. Let $(\hat{\mathbf{a}}, G')$ be such that all agents in the complete subgraph select action a'' and the remaining agents select a' . By step 1, all agents that do not form part of the complete subgraph remain at action a' . Agent \tilde{i} , remains at action a'' as $f_{\tilde{i}}(\hat{\mathbf{a}}, G) = a''$. All remaining agents belonging to the complete subgraph observe only action a'' and hence remain at a'' . Therefore, consensus fails. ■

A.2 Proof of Theorem 2

Proof. Sufficiency follows along the same lines as in the proof of Theorem 1. We prove necessity of the condition as follows. First assume that for all realizations $\mathbf{a} \in A^v$ there exists a unique conditional expected utility (given \mathbf{a}) maximizing action. This directly implies that information aggregation requires consensus. Hence Theorem 1 applies and there needs to exist an action a' such that for all $\mathbf{a} \neq \mathbf{a}'$, $a'' \neq a'$,

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

Next suppose that for some realizations $\hat{\mathbf{a}} \in A^v$ both actions are maximizing the conditional expected utility. We show information aggregation implies that any such vector $\hat{\mathbf{a}}$ has the property that all but one agent agree on one action, i.e. there exist an action that is taken by exactly one agent while all others take the remaining action. First note that similar arguments as in the proof of Theorem 1 imply that conditional on a two-dimensional action vector $\mathbf{a}_{N_i(G) \cup i} = (a', a'')$ there exists a unique conditional expected utility maximizing action which we denote by a' . Suppose there exists a vector $\hat{\mathbf{a}} \in A^v$ such that conditional on $\hat{\mathbf{a}} \in A^v$ both actions are maximizing the conditional expected utility given $\hat{\mathbf{a}}$ and there exist at least two observations of either action in $\hat{\mathbf{a}}$, i.e. using the notation from the proof of Theorem 1 we have (k', k'') such that $k', k'' \geq 2$. Monotonicity as established in the proof of Theorem 1, then shows that conditional on any vector $\tilde{\mathbf{a}} \in A^v$ such that $(k' + 1, k'' - 1)$ action a' is the unique expected utility maximizing action and similarly, that conditional on any vector $\tilde{\mathbf{a}} \in A^v$ such that $(k' - 1, k'' + 1)$ action a'' is the unique expected utility maximizing action.

Now consider the vector $\tilde{\mathbf{a}} \in A^v$ such that $(k' - 1, k'' + 1)$ with $k', k'' \geq 2$. By the arguments above we know that a'' is the unique expected utility maximizing action conditional on $\tilde{\mathbf{a}}$. Information aggregation requires that for an initial action vector $\tilde{\mathbf{a}}$ in the first period, all agents converge to a'' in finite time and in any strongly connected network G . Consider the following network \tilde{G} where $k'' + 1$ agents form a complete subgraph. One of the agents belonging to the complete subgraph, agent \tilde{i} , also has an undirected link to the remaining $k' - 1$ agents each of which is uniquely connected with agent \tilde{i} . Now suppose that in the first period action vector $\tilde{\mathbf{a}}$ all agents in the complete subgraph select action a'' , while all remaining agents select action a' . By the reasoning above, every agent remains at his initial action in all periods, hence contradicting information aggregation.

Monotonicity as established in Theorem 1 then implies that a' is an expected utility maximizing action for all $\mathbf{a} \neq \mathbf{a}'$, $a'' \neq a'$ concluding the proof. ■

A.3 Proof of Theorem 3

Proof. Consider a partition $\{P_1, P_1^C\}$ of Ω such that

$$u(a', \omega') + \delta > u(a'', \omega')$$

for all $\omega' \in P_1$ and some $\delta > 0$, and

$$u(a', \omega') > u(a', \omega'')$$

for all $\omega' \in P_1$ and $\omega'' \in P_1^C$ which by the monotone likelihood property implies

$$\Pr(a' | \omega') > \Pr(a' | \omega'').$$

By the non-triviality assumption there exists such a closed set P_1 with prior probability strictly between 0 and 1. Let

$$\alpha = \inf_{\omega \in P_1} \Pr(a' | \omega)$$

and note that by construction of P_1 and the monotone likelihood property we have

$$(*) \Pr(a' | \omega'') < \alpha$$

for all $\omega'' \in P_1^C$. Further, let \mathbf{a}'_n denote an action vector with all but the first entry being equal to a' . Let p_n be the posterior probability distribution over Ω conditional on \mathbf{a}'_n . By the proof of Theorem 1, we have

$$p_{n+1}(P_1) > p_n(P_1)$$

for all n which implies that $p_n(P_1)$ converges to a limit $l \in (0, 1]$. The theorem is now established in two steps.

1. First, we show that the limit l equals 1. Using Bayes rule, we have that

$$p_n(P_1) = \frac{1}{1 + \frac{\int_{P_1^c} \Pr(a'|\omega) dp_{n-1}}{\int_{P_1} \Pr(a'|\omega) dp_{n-1}}}.$$

As $\Pr(a'|\omega) < \alpha$ for all $\omega \in P_1^c$, we have

$$\int_{P_1^c} \Pr(a'|\omega) dp_n < \alpha \int_{P_1^c} dp_n$$

and

$$\int_{P_1} \Pr(a'|\omega) dp_n \geq \alpha \int_{P_1} dp_n$$

for all n . Further we have

$$\lim_{n \rightarrow \infty} \int_{P_1} \Pr(a'|\omega) dp_n \geq \lim_{n \rightarrow \infty} \alpha \int_{P_1} dp_n,$$

and

$$\lim_{n \rightarrow \infty} \int_{P_1^c} \Pr(a'|\omega) dp_n = \lim_{n \rightarrow \infty} \alpha \int_{P_1^c} dp_n$$

if and only if $l = 1$. Next note that

$$l = \frac{1}{1 + \frac{\lim_{n \rightarrow \infty} \int_{P_1^c} \Pr(a'|\omega) dp_n}{\lim_{n \rightarrow \infty} \int_{P_1} \Pr(a'|\omega) dp_n}} \geq \frac{1}{1 + \frac{\alpha(1-l)}{\alpha l}}$$

and holding with strict inequality iff $l < 1$. First, consider the case $l < 1$, here we have

$$l > \frac{1}{1 + \frac{(1-l)}{l}}$$

which is equivalent to

$$1 > 1$$

establishing a contradiction. This concludes the proof of $l = 1$.

2. By step 1, we have $\lim_{n \rightarrow \infty} p_n(P_1) = 1$. By construction, we have

$$u(a', \omega') + \delta > u(a'', \omega')$$

for all $\omega' \in P_1$. Therefore,

$$\lim_{n \rightarrow \infty} \int_{\Omega} u(a', \omega) dp_n > \lim_{n \rightarrow \infty} \int_{\Omega} u(a'', \omega) dp_n$$

which in turn implies that there exists an $n^* \in \mathbb{N}$ such that for all $n > n^*$

$$\int_{\Omega} u(a', \omega) dp_n > \int_{\Omega} u(a'', \omega) dp_n.$$

Without loss of generality, suppose that

$$\int_{\Omega} u(a', \omega) dp_2 < \int_{\Omega} u(a'', \omega) dp_2$$

Hence the necessary and sufficient condition for information aggregation of Theorem 2 fails for all $n > n^*$ concluding the proof. ■

A.4 Social Learning and the Monotone Likelihood Property

Lemma 1 *The environment of the social learning literature satisfies the monotone likelihood property.*

Proof. Recall that the social learning literature considers a setting with binary states and binary actions, $A = \Omega = \{0, 1\}$ and where agents achieve a utility of 1 if their action matches the realized state, and a utility of 0 otherwise. Denote by S_1 the set of signals for which the posterior probability of state 1 conditional on the signal is greater than half. Hence, S_1 is precisely the set of signals for which action 1 is uniquely optimal. The monotone likelihood property is satisfied if

$$\Pr(S_1 | \omega = 1) > \Pr(S_1 | \omega = 0)$$

which under a uniform common prior is equivalent to

$$F_1(S_1) > F_0(S_1).$$

Consider the posterior probability of state 1 conditional on the set S_1 . As for each $s \in S_1$ the posterior probability of state 1 is greater than half, we have that

$$\Pr(\omega = 1 | S_1) > \frac{1}{2}$$

which under a uniform common prior is equivalent to

$$F_1(S_1) > F_0(S_1)$$

establishing the monotone likelihood property. ■

A.5 Example for Consensus and Information Aggregation

Consider the following example with binary states, binary actions and binary signals, $\Omega = A = S = \{0, 1\}$. Both states are a priori equally likely. Agents achieve a utility of 1 if his action matches the realized state and 0 otherwise. This implies that an action is uniquely optimal if and only if the conditional probability of its corresponding state is greater than half. Conditional on the state of the world, signals are distributed as follows

$$\begin{aligned} \Pr(s = 1 | \omega = 1) &= p_{11} = \frac{1}{4} \\ \Pr(s = 0 | \omega = 0) &= p_{00} = \frac{9}{10}. \end{aligned}$$

Note that conditional on signal $s = 1$ the posterior probability of state $\omega = 1$ is equal to

$$\frac{p_{11}}{p_{11} + (1 - p_{00})} = \frac{\frac{1}{4}}{\frac{1}{4} + \frac{1}{10}} = \frac{40}{56} > \frac{1}{2}$$

and hence observing the signal $s = 1$ induces an optimal action equal to 1. Similarly, the posterior probability of state $\omega = 0$ conditional on $s = 0$ is equal to

$$\frac{p_{00}}{p_{00} + (1 - p_{11})} = \frac{\frac{9}{10}}{\frac{9}{10} + \frac{3}{4}} = \frac{36}{66} > \frac{1}{2}$$

and hence a signal of zero induces an optimal action of 0. Suppose that the network contains three agents. Hence there needs to be an action that is optimal for all but consensus in the other action. Note that $a = 1$ is optimal conditional on k signal's equal to $s = 0$ and one signal $s = 1$ if

$$\frac{p_{11} (1 - p_{11})^k}{p_{11} (1 - p_{11})^k + (p_{00})^k (1 - p_{00})} > \frac{1}{2}$$

which is equivalent to

$$\frac{5}{2} > \left(\frac{6}{5}\right)^k.$$

The inequality is satisfied for $k = 0, 1, 2$ and hence is the necessary and sufficient condition for consensus and information aggregation.

B Instructions

Instructions (sessions with 5-subject networks)

This is an experiment on decision-making. If you follow these instructions and make careful decisions, you may earn a considerable amount of money. Please do not talk with anyone during the experiment.

Participants As the experiment begins, the computer randomly assigns to each participant in the room one of the following 5 labels: A, B, C, D, or E. An equal number of participants is assigned with each label. Then the computer randomly forms 5-person groups. Within a group, each member has a different label. *The label assigned to each participant and the members of each group don't change throughout the experiment.*

During the experiment, you are informed of the choices made by some or all members of your group. If you are informed of the choice of another member, then he is in turn informed of your choice. In other words, you and the other member are connected. A table displays the choices of members connected to you. The choices of members *not* connected to you are displayed as N/A (not available).

Task You play a guessing game together with the members of your group. The game is played with 2 or 4 urns containing colored balls. *Figure B.1* depicts the urns and their composition in the 2-urn game and in the 4-urn game.

In the *2-urn game* there are a White Urn and a Black Urn, each containing 3 colored balls:

- the White Urn has 2 white and 1 black ball.
- the Black Urn has 2 black and 1 white ball.

In the *4-urn game* there are a Red Urn, a Yellow Urn, a Green Urn and a Blue Urn, each containing 7 colored balls:

- the Red Urn has 4 red, 1 yellow, 1 green and 1 blue ball
- the Yellow Urn has 4 yellow, 1 red, 1 green and 1 blue ball
- the Green Urn has 4 green, 1 red, 1 yellow and 1 blue ball
- the Blue Urn has 4 blue, 1 red, 1 yellow and 1 green ball.

You are always informed of which game is played. The games differ only in the number and composition of the urns. They follow the same rules.

The computer starts by randomly selecting an urn. Every urn is equally likely. You are not informed of which urn is selected. The computer then draws a ball from the selected urn and you are informed of the color of the ball. The draw is your private information and is not shared with anyone else. A draw is made independently for each group member. After being drawn, a ball is returned to the urn, keeping the composition of balls in the urn constant.

The game is divided in 6 decision turns. In the 1st turn, after learning about your own draw, you are asked to *indicate the urn that you think is more likely to have been used*. Next, as the 2nd turn starts, you are informed of the choices

made in the previous turn by the group members connected to you. You are then asked again to *indicate the urn that you think is more likely to have been used*. This process is repeated for 6 turns. ***Note that the urn selected by the computer and each member's private draw don't change across turns.***

Repetition of the task You play 18 rounds like the one described above. A table displays whether your guess in each turn of each previous round was correct. In 9 consecutive rounds you play the 2-urn game and in other 9 consecutive rounds you play the 4-urn game. Whether you play the 2-urn game before or after the 4-urn game is determined randomly at the beginning of the experiment. This information is displayed on the screen.

While the members of your group don't change, how many and which members are connected to you may change from one round to the next. In other words, connections among group members may change. Even when your connections don't change, the connections among other members may change.

Earnings At the end of the experiment, the computer randomly selects one turn for each round. For every round, you earn CHF2 if your guess in the selected turn was correct (i.e. if the urn you indicated was the urn selected by the computer), otherwise you earn nothing. It's in your interest to do your best in every turn of every round.

The experiment starts with a practice round, which doesn't count for the determination of earnings.

Instructions (sessions with 7-subject networks)

This is an experiment on decision-making. If you follow these instructions and make careful decisions, you may earn a considerable amount of money. Please do not talk with anyone during the experiment.

Participants As the experiment begins, the computer randomly assigns to each participant in the room one of the following 7 labels: A, B, C, D, E, F or G. An equal number of participants is assigned with each label. Then the computer randomly forms 7-person groups. Within a group, each member has a different label. ***The label assigned to each participant and the members of each group don't change throughout the experiment.***

During the experiment, you are informed of the choices made by some or all members of your group. If you are informed of the choice of another member, then he is in turn informed of your choice. In other words, you and the other member are connected. A table displays the choices of members connected to you. The choices of members *not* connected to you are displayed as N/A (not available).

Task You play a guessing game together with the members of your group. The game is played with 2 or 4 urns containing colored balls. *Figure B.1* depicts the urns and their composition in the 2-urn game and in the 4-urn game.

In the *2-urn game* there are a White Urn and a Black Urn, each containing 3 colored balls:

- the White Urn has 2 white and 1 black ball.
- the Black Urn has 2 black and 1 white ball.

In the *4-urn game* there are a Red Urn, a Yellow Urn, a Green Urn and a Blue Urn, each containing 7 colored balls:

- the Red Urn has 4 red, 1 yellow, 1 green and 1 blue ball

- the Yellow Urn has 4 yellow, 1 red, 1 green and 1 blue ball
- the Green Urn has 4 green, 1 red, 1 yellow and 1 blue ball
- the Blue Urn has 4 blue, 1 red, 1 yellow and 1 green ball.

You are always informed of which game is played. The games differ only in the number and composition of the urns. They follow the same rules.

The computer starts by randomly selecting an urn. Every urn is equally likely. You are not informed of which urn is selected. The computer then draws a ball from the selected urn and you are informed of the color of the ball. The draw is your private information and is not shared with anyone else. A draw is made independently for each group member. After being drawn, a ball is returned to the urn, keeping the composition of balls in the urn constant.

The game is divided in 8 decision turns. In the 1st turn, after learning about your own draw, you are asked to *indicate the urn that you think is more likely to have been used*. Next, as the 2nd turn starts, you are informed of the choices made in the previous turn by the group members connected to you. You are then asked again to *indicate the urn that you think is more likely to have been used*. This process is repeated for 8 turns. ***Note that the urn selected by the computer and each member's private draw don't change across turns.***


Repetition of the task You play 14 rounds like the one described above. A table displays whether your guess in each turn of each previous round was correct. In 7 consecutive rounds you play the 2-urn game and in other 7 consecutive rounds you play the 4-urn game. Whether you play the 2-urn game before or after the 4-urn game is determined randomly at the beginning of the experiment. This information is displayed on the screen.

While the members of your group don't change, how many and which members are connected to you may change from one round to the next. In other words, connections among group members may change. Even when your connections don't change, the connections among other members may change.

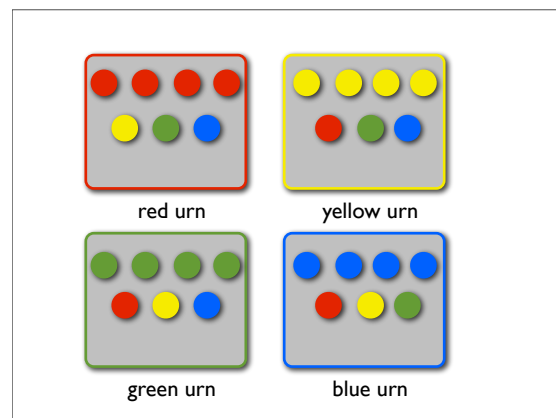
Earnings At the end of the experiment, the computer randomly selects one turn for each round. For every round, you earn CHF2.50 if your guess in the selected turn was correct (i.e. if the urn you indicated was the urn selected by the computer), otherwise you earn nothing. It's in your interest to do your best in every turn of every round.

The experiment starts with a practice round, which doesn't count for the determination of earnings.

Figure B.1: Urn composition



(a) 2-urn game



(b) 4-urn game

C Tables and Figures

Table C.1: Network structures: who's neighbor to whom?

(a) 5-agent networks


network structure	neighbors for each subject in the network				
	A	B	C	D	E
complete	B,C,D,E	A,C,D,E	A,B,D,E	A,B,C,E	A,B,C,D
linked circle (A-C)	B,C,E	A,C	A,B,D	C,E	A,D
linked circle (A-D)	B,D,E	A,C	B,D	A,C,E	A,D
linked circle (B-D)	B,E	A,C,D	B,D	B,C,E	A,D
linked circle (B-E)	B,E	A,C,E	B,D	C,E	A,B,D
linked circle (C-E)	B,E	A,C	B,D,E	C,E	A,C,D
star (A center)	B,C,D,E	A	A	A	A
star (B center)	B	A,C,D, E	B	B	B
star (C center)	C	C	A,B,D,E	C	C

(b) 7-agent networks

network structure	neighbors for each subject in the network						
	A	B	C	D	E	F	G
linked circle (A-C,D-F)	B,C,G	A,C	A,B,D	C,E,F	D,F	D,E,G	A,F
linked circle (B-D,E-G)	B,G	A,C,D	B,D	B,C,E	D,F,G	E,G	A,E,F
linked circle (C-E,F-A)	B,F,G	A,C	B,D,E	C,E	C,D,F	A,E,G	A,F
linked circle (A-D,B-G)	B,D,G	A,C,G	B,D	A,C,E	D,F	E,G	A,B,F
linked circle (A-D)	B,D,G	A,C	B,D	A,C,E	D,F	E,G	A,F
linked circle (A-D,A-E,B-G)	B,D,E,G	A,C,G	B,D	A,C,E	A,D,F	E,G	A,B,F
conn. complete comp.	B,C,D	A,C	A,B	A,E	D,F,G	E,G	E,F

Figure C.1: Network structures

(a) 5-agent networks



(b) 7-agent networks




Table C.2: Choice revisions. Number of observations pooling data across participants and rounds (i.e. network structures and choice sets). For 5-agent networks, 1800 observations = 100 participants \times 18 rounds. For 7-agent networks, 1372 observations = 98 participants \times 14 rounds.

(a) 5-agent networks

round of last choice revision	No. choice revisions						all
	0	1	2	3	4	5	
	1092	0	0	0	0	0	1092
2		216	0	0	0	0	216
3		59	48	0	0	0	107
4		46	37	12	0	0	95
5		25	49	22	13	0	109
6		29	46	46	38	22	181
all	1092	375	180	80	51	22	1800

(b) 7-agent networks

round of last choice revision	No. choice revisions								all
	0	1	2	3	4	5	6	7	
	777	0	0	0	0	0	0	0	777
2		157	0	0	0	0	0	0	157
3		15	16	0	0	0	0	0	31
4		20	23	9	0	0	0	0	52
5		19	27	8	8	0	0	0	62
6		16	12	11	4	2	0	0	45
7		10	24	17	8	7	3	0	69
8		14	48	35	31	19	13	19	179
all	777	251	150	80	51	28	16	19	1372