

TESTING MODELS OF SOCIAL LEARNING ON NETWORKS: EVIDENCE FROM A FRAMED FIELD EXPERIMENT

ARUN G. CHANDRASEKHAR[†], HORACIO LARREGUY[‡], AND JUAN PABLO XANDRI^{*}
(PRELIMINARY AND INCOMPLETE)

ABSTRACT. Theory has focused on two leading models of social learning on networks: Bayesian and DeGroot rules of thumb learning. These models can yield greatly divergent behavior; individuals employing rules of thumb often double-count information and may not exhibit convergent behavior in the long run. By conducting a unique lab experiment in rural Karnataka, India, set up to exactly differentiate between these two models, we test which model best describes social learning processes on networks. We study experiments in which seven individuals are placed into a network, each with full knowledge of its structure. The participants attempt to learn the underlying (binary) state of the world. Individuals receive independent, identically distributed signals about the state in the first period only; thereafter, individuals make guesses about the underlying state of the world and these guesses are transmitted to their neighbors at the beginning of the following round. We consider various environments including incomplete information Bayesian models and provide evidence that individuals are best described by DeGroot models wherein they either take simple majority of opinions in their neighborhood.

KEYWORDS: Networks, Bayesian learning, DeGroot learning, Lab experiments

JEL CLASSIFICATION CODES: D82, D83, D85, C92, C93

1. INTRODUCTION

Understanding the mechanisms of social learning is very important, particularly in developing nations. Due to the lack of formal institutions, missing markets, and information aggregating mechanisms, agents in developing economies often rely on social connections for information and opportunities. Social learning is often characterized by word-of-mouth communication. Given that social learning is a fundamental component of many economic processes, before employing models of social learning to make policy recommendations we

Date: March 2012.

We are grateful to Daron Acemoglu, Abhijit Banerjee, Esther Duflo, Matthew O. Jackson, Anna Mikusheva, Ben Olken, and Rob Townsend who provided support, guidance, and extremely helpful suggestions. Essential feedback was provided by Juan Dubra, Ben Golub, Dhruv Parthasarathy, Xiao Yu May Wang, and Luis Zermeno. This project was possible thanks to the financial support of the Russell Sage Behavioral Economics Grant. Chandrasekhar is grateful for support from the National Science Foundation GFRP. Larreguy thanks the Bank of Spain and Caja Madrid for financial support. The Poverty Action Lab and The Centre for Microfinance at the Institute for Financial Management and Research provided valuable logistical assistance. Gowri Nagraj and our field team were vital to the project's progress. We also thank Gayathri Chandrasekhar for assistance with translation. Alejandra Menchaca and Archana Prakash provided support and patience throughout the project. All errors are our own.

[†]Massachusetts Institute of Technology. Email: arunc@mit.edu.

[‡]Massachusetts Institute of Technology. Email: larreguy@mit.edu.

^{*}Massachusetts Institute of Technology. Email: jxandri@mit.edu..

must first understand which models best describe features of empirical social learning. The mechanics of the learning process are of policy interest. If the features of social learning are better described by certain models, those models should be the environment in which the relevant economic outcomes are studied. The two leading classes of social learning models employed in the network literature are Bayesian learning and DeGroot rules of thumb models. The models often exhibit distinct behavior. Individuals employing rules of thumb often double-count information and may not reach consensus in the long run.

In this paper, we will address the question of whether Bayesian learning or DeGroot rules of thumb models do a better job of describing empirical learning processes on networks. To study this question we conduct a unique lab experiment in the field across 19 villages in rural Karnataka, India. We ran our experiments in villages so that we could study the relevant population of interest, namely those who could be potentially targeted by policy that depends greatly on social learning. Our approach was to design simple networks that provide statistical power to distinguish between the different learning models. We then conducted a lab experiment in the field using these networks to address the proposed question.

We created networks of seven individuals and gave each individual a map of the entire graph so that the full informational structure was comprehended. The underlying state of the world was either one or zero with equal probability. At $t = 0$ each individual received an independent identically distributed (iid) signal about the underlying state of the world and were informed that signals were correct with probability $5/7$. After receiving the signal each individual privately made a guess about the state of the world. These guesses were communicated to each individual’s network neighbors at the start of the first period, $t = 1$.

Thereafter, in any given period, each individual knew the guesses of all of her network neighbors from all past periods. Using this information, she made a guess about the state of the world, which in turn was communicated to each of her network neighbors at the beginning of the following period. Every individual was paid for her guess in a randomly chosen round from the set of rounds that she played that day over the course of all the experiments. Consequently, participants had strong incentives to make their best guess in each round.¹

The work most closely related to ours is Möbius et al. (2011), who study how information decays as it spreads through a network. They test between DeGroot models and a streams model that they develop in which individuals “tag” information by describing where it comes from. Their experiment uses Facebook network data in conjunction with a field experiment and finds evidence in favor of the streams model. In our experiment, we shut

¹While it could be the case that players were extremely sophisticated and engaged in experimentation in early rounds, anecdotal evidence from participants suggests that this is not the case. In addition, the theoretical and experimental literature uses this assumption.

down this ability for individuals to “tag” information to be able to compare the Bayesian model to DeGroot alternatives.

We are able to analyze the data at two levels: the network level and the individual level. Network level analysis considers the entire network and sequence of actions as a single observation. That is, theory predicts a path of actions under a model of social learning, for each individual in each period given a network. At the network level, we address a question about how well social learning behaves; the observational unit in this approach is the social network itself. Meanwhile individual level analysis considers the action of an individual, given a history, as the observational unit.

Our core results are as follows. First, at the network level, we find evidence that a DeGroot rule of thumb model better explains the data than the Bayesian learning model. This is not to say, however, that the social learning process does not resemble the data generated by Bayesian learning. In fact, the Bayesian learning model explains 62% of the actions taken by individuals while the best DeGroot rule of thumb explains over 76% of the actions taken by individuals.²

Second, at the individual level, we find that a DeGroot rule of thumb model of learning performs significantly better than Bayesian learning in explaining the actions of an individual given a history of play. In fact this model explains nearly 87% of the actions taken by individuals given a history.

Third, we make a methodological contribution by identifying key problems in estimating models of Bayesian learning on networks with trembles or quantal response equilibrium (QRE), the standard technique in the literature. We demonstrate that networks that are small enough to avoid computational constraints are not large enough to tease out the differences between DeGroot and Bayesian learning with trembles. Meanwhile those that are large enough to separate the models become computationally infeasible to study using trembles or QRE.

We also establish several supplementary results which may also be of independent methodological interest. First, we identify a key asymmetry in the social learning on networks literature. Usually this literature has focused on asymptotic learning, whether individuals end up learning the state of the world as time goes to infinity. Models of Bayesian learning have focused on action models; namely, individuals observe the actions taken by their neighbors and update their beliefs accordingly. Under some regularity conditions, this model of learning has nice asymptotic properties and learning occurs well. Meanwhile, the literature emphasizing DeGroot learning, e.g. [Golub and Jackson \(2010\)](#), has focused on a communication model wherein each individual communicates her belief to

²When we say a model explains $x\%$ of the actions, we are interested in $x := \frac{y-50}{50}$, where y is the percent of actions predicted correctly. This is the right normalization since we could always explain half the actions by flipping a coin.

her neighbors every period. This literature, too, has found that asymptotic learning happens with minimal regularity conditions. However, we point out that such models allow for myopic individuals (those following rules of thumb) to transmit too much information relative to the Bayesian learning models. By considering the fair comparison, an action model of DeGroot learning, we add to the literature the first example of networks which satisfy the properties to have asymptotic learning under both Bayesian (action) and DeGroot (communication) learning models, but fail to have asymptotic learning with DeGroot (action) learning models.

Second, we establish a unique algorithm to simulate Bayesian learning on networks which is computationally tight in the sense that asymptotically there can be no faster algorithm. Namely, the algorithm is $O(T)$ where T is the number of rounds played. This algorithm is actually asymptotically efficient given our setup.³

Third, we provide a discussion of why such questions of social learning must be addressed in a structural manner and, in particular, using a lab experiment. Researchers may be tempted to use the intuitions of Bayesian learning and DeGroot learning to test for correlations in regression analysis of social learning data. Two natural examples are (a) regressions that check whether individuals overweight the same information if they receive it through multiple channels and (b) regressions that explore whether individuals treat old information that cycles back to them as if it was new, additional information. Under these frames, the null hypothesis is the Bayesian model, since in these specifications the Bayesian models suggest that the relevant parameters ought to be zero. A rejection of zero in the coefficients of these regressions provides evidence in the direction of the DeGroot rules of thumb. To test the effectiveness of regression-based analysis, we simulate learning data from a Bayesian model using the aforementioned algorithm as well as from DeGroot models. We show that such regression-based approaches do not work, namely, the data generated even by Bayesian learning models do not conform to the intuition motivating the regressions. We maintain that, in turn, the researcher ought to proceed by a structural analysis. Given the computational constraints for structural estimation of learning models in large networks, this suggests that separating models of social learning are best addressed in a lab setting.

There is little empirical evidence comparing Bayesian learning with rules of thumb learning in non-experimental contexts. Without experimental data it is difficult to control priors of agents in the network and the signal quality. Moreover, even in field experiments separating between Bayesian and DeGroot models may be difficult. First, structural approaches are computationally infeasible even with moderately sized networks (10 nodes), as it will become clear below in our discussion of computational complexity. Second, reduced form tests may not suffice for separating between these models. Third, empirical network data

³An algorithm is $O(T)$ if the number of computations as a function of T , $f(T)$, is such that $\frac{f(T)}{T} \rightarrow M$ for some constant M . In particular, this is true if $f(T) = MT$, as it is in our algorithm

may not be precisely measured, affecting the conclusions of a researcher who is trying to select between these models of learning. There may be problems with estimating a structural model on a sampled network and the survey-obtained social network may not be precisely the communication channels used in practice, both of which would induce biases (Chandrasekhar and Lewis, 2010). Since network-based estimation (which is inherently structural even when using reduced form regressions) is often sensitive to misspecification of the network, it is difficult to cleanly identify which model best describes the data in a non-laboratory context. Fourth, we are unable to know exactly what information is being transmitted in empirical data. Without knowing whether the information transmitted in this context is beliefs, actions or something else all together, one may mistakenly select the wrong model because of not properly specifying the information that is communicated.

Meanwhile, we believe that for our purposes, conducting a lab experiment outside the field of interest is insufficient because we desire to describe the social learning process for our population of interest. We are precisely interested in studying the social behavior of rural populations in a developing country as this is the relevant population in the aforementioned literature.

Related Literature. Acemoglu et al. (2010) and Jackson (2008) provide excellent reviews of the social learning on networks literature. The literature is partitioned by whether the learning is Bayesian or myopic (following some rule of thumb). On top of this, the literature layers a myriad of questions such as whether individuals learn from the communication of exact signals (or beliefs or payoffs of other agents) or by observing others' actions, whether the information arrives once or enters over time, whether the interaction is simultaneous or sequential, etc.

Bayesian Learning. The benchmark models of Bayesian learning come from Banerjee (1992) and Bikhchandani et al. (1992). They examine sequential decisions by Bayesian agents who observe past actions. These papers point out that consensus are formed and thereafter agents end up making the same decision, which may in fact be the wrong decision. Smith and Sorensen (2000) and Celen and Kariv (2005) look at Bayesian-rational sequential decision making and explore the conditions under which asymptotic learning is attainable. Acemoglu et al. (2010) extend the framework to consider a social network environment in which individuals have stochastic neighborhoods. Their main result is that asymptotic learning occurs even with bounded beliefs for stochastic topologies such that there is an infinitely growing subset of agents who are probabilistically “well informed” (i.e. with some probability observe the entire history of actions) with respect to whom the rest of the agents have expanding observations. Gale and Kariv (2003) study an environment that is closest to ours. They only focus on Bayesian learning and extend the learning model to a finite social network with multiple periods. At time t each agent makes a decision given her information set, which includes the history of actions of each of their neighbors in the

network. Via the martingale convergence theorem, they point out that connected networks with Bayesian agents will yield uniform actions in finite time with probability one.

DeGroot Learning. Ellison and Fudenberg (1993, 1995) are benchmarks for the rule of thumb learning models. Bala and Goyal (1998) investigate an extension of the two-armed bandit problem to the network context with boundedly rational agents and demonstrate that there will be payoff equalization and uniformity of action in the limit via an imitation principle.

In a network context, DeGroot (1974) provides the most influential non-Bayesian framework. Agents observe signals just once and communicate with each other and update their beliefs via a weighted and possibly directed trust matrix. Golub and Jackson (2010) characterize the asymptotic learning for a sequence of growing networks. They argue that crowds are wise, provided that there are not agents that are too influential. DeMarzo et al. (2003) also consider a DeGroot style model and show that as agents fail to account for the repetition of information propagating through the network, persuasion bias may be a serious concern affecting the long run process of social opinion formation. Moreover, they show that even multidimensional beliefs converge to a single line prior to obtaining a consensus, thereby demonstrating how a multidimensional learning process can be characterized by a uni-dimensional convergence. Chatterjee and Seneta (1977), Berger (1981), Friedkin and Johnsen (1997), and Krackhardt (1987) are among other papers that examine variations on the DeGroot models.

The main contribution of this paper is to bridge the experimental literature with common myopic models of learning, motivated by the DeGroot model tradition, in a framework that allows us to study the Bayesian and DeGroot-like features of empirical social learning processes. We fill the gap by conducting experiments designed to precisely aid us in model selection procedures and establish regularities of the learning process on networks.

Learning Experiments. The literature on social learning experiments begins with Anderson and Holt (1997), Hung and Plott (2001), and Kubler and Weizsacker (2004). Explicit network structure are considered in a series of papers by Gale and Kariv (2003), Choi et al. (2005, 2009), and Çelen, Kariv, and Schotter (2010).

Choi et al. (2005, 2009) make a seminal contribution to the empirical literature of social learning by testing the predictions derived by Gale and Kariv (2003) in a laboratory experiment. They are able to show that features of the Bayesian social learning model fit the data well for networks of three individuals. However, they do not allow for statistical power under the DeGroot alternatives. In extremely simple networks, such as the ones studied in their paper, there are few (if any) differences in the predicted individual learning behavior by the Bayesian and the rule of thumb learning models.

As mentioned above, Möbius et al. (2011) have conducted a related a field experiment using the Facebook network of Harvard undergraduates to distinguish between streams

models versus DeGroot models of social learning. This paper is complementary to ours. We address whether individuals observing only actions of their neighbors behave in a Bayesian or DeGroot manner and they are interested in whether myopic individuals, who can communicate more information by tagging, appear to behave in a streams manner or a DeGroot manner.

Structure of the Paper. The rest of the paper is organized as follows. Section 2 develops the theoretical framework. Section 3 contains the experimental setup. Section 4 describes the structural estimation procedure and the main results of the estimation. Section 5 presents the discussion of the difficulties of reduced form approaches. Section 6 concludes.

2. FRAMEWORK

2.1. Notation. Let $G = (V, E)$ be a graph with a set V of vertices and E of edges and put $n = |V|$ as the number of vertices. We denote by $A = A(G)$ the adjacency matrix of G and assume that the network is an undirected, unweighted graph, with $A_{ij} = \mathbf{1}\{ij \in E\}$. Let $N_i = \{j \in V : ij \in E\} \cup \{i\}$ the set of neighbours of i (i.e. the individuals who are connected to i) plus himself. Individuals in the network are attempting to learn about the underlying state of the world, $\theta \in \Theta = \{0, 1\}$. Time is discrete with an infinite horizon, so $t \in \mathbb{N}$.

At $t = 0$, and only at $t = 0$, agents receive iid signals $s_i|\theta$, with $P(s_i = \theta|\theta) = p$ and $P(s_i = 1 - \theta|\theta) = 1 - p$. The signal correctly reflects the state of the world with probability p . We denote by W the set of all possible combinations of signals among agents, which we will refer to as “worlds”. Therefore $s \in S$ is an element $s = (s_1, \dots, s_n)$ with $s_i \in \{0, 1\}$. Note that $|W| = 2^n$. We will use $d_i = \sum_j A_{ij}$ to refer to the vector of degrees for $i \in \{1, \dots, n\}$ and ξ for the eigenvector corresponding to the maximal eigenvalue of A .

2.2. Bayesian Learning. In our analysis we consider a model of Bayesian learning with incomplete information. Individuals will have common priors over the relevant state spaces (described below) and update according to Bayes’ rule in each period. We formalize the model in Appendix B. Each agent is drawn from a population which has π share Bayesian agents and $1 - \pi$ share DeGroot agents and this fact is common knowledge, as is the structure of the entire network. However, there is incomplete information about the types of the other agents in the network, and the Bayesian individuals will attempt to learn about the types of the other agents in the network along the path while attempting to learn about the underlying state of the world. The incomplete information setup is a useful step beyond the fully Bayesian environment, restricting $\pi = 1$. For instance, if an individual believes that her neighbor does not act in a Bayesian manner, she will process the information about observed decisions accordingly; as outside observers, the econometricians might think that she is not acting as a Bayesian. This is a very common problem when testing Bayesian

learning, because we need to make very strong assumptions about common knowledge. A model in which there is incomplete information about how other players behave attempts to address this issue while only minimally adding parameters to be estimated in an already complicated system.

2.3. DeGroot Learning. We begin with a classical model of rule of thumb learning on networks and discuss three specific and natural parametrizations. Jackson (2008) contains an excellent review of DeGroot learning models. In our experiment, we consider DeGroot action models as opposed to communication models. In action models individuals observe the historical actions of their network neighbors, while in communication models individuals will be able to communicate their beliefs to their neighbors.

We are interested in action models for several reasons. First, the models of Bayesian learning on networks are action models, so it is the appropriate comparison. Second, it is extremely difficult to get reliable, measurable, and believable data of beliefs in a communication model for a lab experiment conducted in the field in rural villages. Third, as it is difficult to control and map into data exactly what is (or isn't) communicated by various agents in a more general communication model, we are able to focus on the mechanics of the learning process by restricting communication to observable actions. Fourth, this also fits with the motivating literature wherein individuals may observe the actions, such as technology or microfinance adoption decisions, of their neighbors.

Let $T = T(A)$ be a weighted matrix which parametrizes the weight person i gives to the action of person j . We study three natural parametrizations of the DeGroot model. The first is *uniform weighting* wherein each individual weights each of her neighbors exactly the same. The weight matrix $T^u(A)$ is given by

$$T_{ij}^u = \frac{A_{ij}}{d_i + 1} \text{ and } T_{ii}^u = \frac{1}{d_i + 1}$$

meaning that each individual puts $(d_i + 1)^{-1}$ weight on each of her d_i neighbors as well as on herself.

The second model we consider is *degree weighting*. Each individual weights her neighbors by their relative popularity, given by degree. $T^d(A)$ is given by

$$T_{ij}^d = \frac{d_j}{\sum_{j \in N_i} d_j + d_i} \text{ and } T_{ii}^d = \frac{d_i}{\sum_{j \in N_i} d_j + d_i}.$$

The third model is *eigenvector weighting*. An individual places weight on her neighbor proportional to the neighbor's relative importance, given by eigenvector centrality. $T^e(A)$ is given by

$$T_{ij}^e = \frac{\xi_j}{\sum_{j \in N_i} \xi_j + \xi_i} \text{ and } T_{ii}^e = \frac{\xi_i}{\sum_{j \in N_i} \xi_j + \xi_i}$$

where ξ is the eigenvector corresponding to the maximal eigenvalue of A . This is motivated by the idea that an individual may put greater weight on more information-central neighbors, which eigenvector centrality captures.

While a very natural parametrization of learning, the DeGroot model misses strategic and inferential features of learning. Behavior is as follows. At time $t = 0$, individuals receive signals $s = (s_1, s_2, \dots, s_n)$, and accordingly, actions $a_{i,0} = \mathbf{1}\{s_i = 1\}$ are taken. Let $\mathbf{a}_0 = (a_{1,0}, a_{2,0}, \dots, a_{n,0})$. At the beginning of $t = 1$, individual i observe $a_t^{(j)}$ for all $j \in N_i$ and aggregates information to form beliefs to $b_1 = T\mathbf{a}_0$. In turn, actions are chosen $\mathbf{a}_1 = \mathbf{1}\{b_1 > 1/2\}$. Now consider time $t = k$ with action profile \mathbf{a}_k . Then beliefs for stage $k + 1$ are formed $b_{k+1} = T\mathbf{a}_k$ and accordingly actions are chosen $a_{k+1} = \mathbf{1}\{b_{k+1} > 1/2\}$. In turn, we have if the limit exists,

$$\begin{aligned} a_\infty &= \lim_{k \rightarrow \infty} \mathbf{1}\{Ta_{k+1} > 1/2\} \\ &= \lim_{k \rightarrow \infty} \mathbf{1}\{T \cdot \mathbf{1}\{Ta_k > 1/2\} > 1/2\}, \quad a_k = \mathbf{1}\{Ta_{k-1} > 1/2\}. \end{aligned}$$

While we cannot easily analyze this using the theory of linear operators (due to nested indicator functions), we will discuss its implications in section 2.5.

2.4. A Fair Comparison. We argue that there is an important gap in the literature of learning on networks. Models of Bayesian learning in the literature have largely been action based; individuals observe the actions taken by their neighbors and accordingly update their beliefs and in this manner asymptotic consensus is achieved under certain regularity conditions (Gale and Kariv). In a sequential learning environment, provided no finite set of individuals is asymptotically excessively influential, asymptotic learning occurs (Acemoglu et al.). However, the literature on DeGroot learning, e.g. Golub and Jackson (2010), has focused on models of communication to obtain asymptotic learning results under similar regularity conditions. So long as there is no finite set of agents that are too central in a particular sense, the network is wise in the sense that asymptotic learning occurs. Notice that the wisdom of the DeGroot learning hinges on the fact that an extensive amount of information is passed along such a model relative to the action model of Bayesian learning. For a parallel, in Bayesian learning if we introduced a communication model, then the filtering problem would be simpler since an agent would know her neighbors' posteriors exactly.

We argue that there is an asymmetry in the literature; the right abstraction to think about social learning ought to be parallel across the competing Bayesian and DeGroot models. By considering the fair comparison, an action model of DeGroot learning, we add to the literature a key example of networks which satisfy the properties to have asymptotic learning under both Bayesian (action) and DeGroot (communication) learning models, but fail to have asymptotic learning with DeGroot (action) learning models.

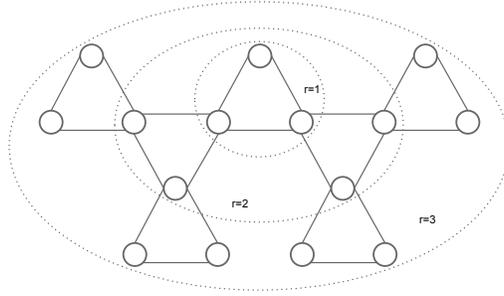


FIGURE 1. A concentric social quilt.

2.5. An Illustrative Example: Concentric Social Quilts. We present a simple setup which yields asymptotic learning under communication DeGroot models and consensus under action Bayesian models, but fails asymptotic learning and violate consensus with action DeGroot models. Namely, a number of nodes will become “stuck” in an information trap in a local neighborhood of the network. This demonstrates a wedge between DeGroot and Bayesian learning and therefore motivates our learning experiment.

The motivation for the graph structure comes from Jackson et al. (2010) who study network architecture that arise as equilibria in favor exchange games. They show that these networks will be social quilts; a social quilt is a patchwork of substructures (e.g., triangles) pasted together in specific ways. We take a very simple example of this style of a network. From the applied perspective, the intuition is that if graphs are constructed as equilibria of risk-sharing or favor-exchange games, then they may have such quilt-like substructures. While these quilt-like structures enable network members to maintain favor-exchange relationships in equilibrium through local punishment of misbehavior, the same networks are also the surface on which information passes among members. We caution that it may be the case that information does not transmit efficiently through social quilts.

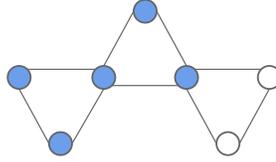
To illustrate this we define a concentric social quilt (CSQ) as a graph that consists of triangles quilted together around a central triangle such that every triangle (in the interior of the structure) is connected to exactly three triangles in the following way. Consider a sequence of CSQs which can be constructed following a recursive process as the number of nodes goes to infinity. We index this sequence by $r \in \mathbb{N}$.

- (1) Take CSQ_{r-1} and let T_r be the set of terminal nodes of CSQ_{r-1} .
- (2) To each terminal node $i \in T_r$, attach a triangle with two new nodes added.

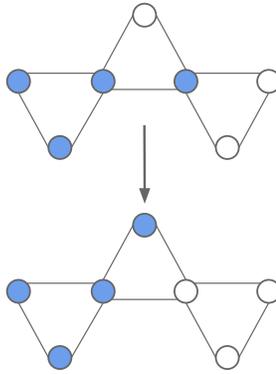
Figure 1 shows such a network and the model is developed in detail in Appendix D.

DEFINITION 2.1. We say that node $i \in V_r$ is stuck if there exists a $t_i \in \mathbb{N}$ such that for all $t \geq t_i$, $a_{i,t} = 1 - \theta$.

A node is stuck if the node for all but finitely many periods takes the same (wrong) action. Figure 2 provides two examples of nodes that get stuck despite the majority of nodes in the network receiving the right signal.



Panel A:



Panel B:

FIGURE 2. In Panel A two nodes are stuck for all periods $t \in \mathbb{N}$, even though 5 of the 7 nodes have received the true signal. In Panel B in the first period 4 nodes receive the true signal, and after one node switches, 3 are asymptotically stuck.

Panel A of Figure 2 illustrates the problem. Assume that for some subtree of the CQR , which connects to the rest of the network through the top-most node, we have the initial signal endowment shown. Any information from the rest of the graph will come via top-most node, which we will call parent node. To get a lower bound on the number of nodes that get stuck in the wrong action, we can simply assume that the parent node of the subtree always chooses the right action for all rounds. However, even in this case the nodes in the lower right triangle act in the same (wrong) manner for all but finitely many periods. As the sequence of networks grow, $r \rightarrow \infty$, there will be a non-vanishing fraction of subtrees with this initial configuration of signals. These subtrees will have at least $3/7$ nodes which become stuck. This example has demonstrated the following result.

PROPOSITION 2.1 (Wedge in Learning Models). *For a sequence of concentric social quilts with iid signals with probability p , with probability approaching one*

- (1) *under the DeGroot communication model with uniform weighting the network is wise,*⁴
- (2) *and under the DeGroot action model with uniform weighting a non-vanishing fraction of nodes get stuck.*

Proof. All proofs are contained in Appendix E. □

That asymptotic learning occurs with the DeGroot communication model follows from Corollary 1 of Golub and Jackson (2010), and the result for the DeGroot action model is apparent from the previous example. To illustrate the severity of Proposition 2.1, in Table 1 we show lower bounds on the expected fraction of nodes that are stuck. Even with high quality signals ($p = 0.7$) at least 16% of nodes become stuck and do not asymptotically learn. In particular, recall that the benchmark for mistakes is 50%, since a node can always randomly guess. Therefore, relative to the expected fraction of nodes that should have learned, at least 25% actually get stuck with the wrong information.

In addition to motivating the study of DeGroot action models in our experiment, this example is of independent interest as it illustrates a wedge in the learning on networks literature between the usual DeGroot and Bayesian models. It also raises the question about whether certain network structures are better for social learning, given that asymptotic learning may not occur due to this stuck property. We conjecture that expander graphs will generate asymptotic learning even with action DeGroot models and address this more thoroughly in ongoing work.

3. EXPERIMENT

3.1. Setting. Our experiment was conducted at 19 villages in Karnataka, India. The villages range from 1.5 to 3.5 hours' drive from Bangalore. A village setting was chosen because social learning through networks such as by word-of-mouth communication is of the utmost importance in rural environments; information about new technology (Conley and Udry, 2010), microfinance (Banerjee et al., 2010), among other things propagates through the social network.

3.2. Overall Game Structure. In each village, individuals played the social learning game three times, each time with a different network structure. The three networks (see Figures 3) were played with a random order in each village. Every network consisted of seven individuals and each participant was shown the entire network structure as well as her own location in the network.

At the beginning of each game, every individual was shown two identical bags, one with three yellow balls and one blue ball and the other which has three blue balls and one yellow

⁴Let $\mu = \theta p + (1 - p)(1 - \theta)$ and T_n is a sequence of convergent row-normalized matrices. As defined in Golub and Jackson (2010), the sequence is wise if $\text{plim}_{n \rightarrow \infty} \sup_{i \leq n} |\lim_{t \rightarrow \infty} T_n^t s_n - \mu| = 0$. In our context wisdom corresponds to asymptotic learning since in the limit a share of nodes that have belief μ goes to one and therefore the nodes can distinguish $\mu > 0$ or $\mu < 0$, as p is known.

ball. One of the two bags was chosen at random to represent the state of the world and the goal of the game was that the participant had to independently guess whether the blue bag or the yellow bag had been selected. Since there was an equal probability that either bag could be chosen, we induced priors of $1/2$. As the selected bag contained three balls reflecting the state of the world, individuals anticipated receiving independent signals that were correct with probability $5/7$.

After receiving their signals in round zero, all individuals simultaneously and independently made their best guesses about the underlying state of the world (which bag had been selected). The game continued to the next round randomly and on average lasted 6 rounds. If the game continued to the second round, at the beginning of the second round each individual was shown the round one guesses of the other individuals in her neighborhood, N_i . Agents updated their beliefs about the state of the world and then again made their best guesses about the state of the world. The game again continued to the following round randomly. This process repeated until the game came to a close. Notice that after the time zero set of signals, no more signals were drawn during the course of the game. Individuals would only observe the historical decisions of their neighbors and update their own beliefs accordingly.

Individuals were paid for a randomly chosen round from a randomly chosen game and therefore faced non-trivial incentives to submit a guess which reflected their belief about the underlying state of the world.

3.3. Network Choice. We selected networks specifically so that we could separate between various DeGroot and Bayesian models considered in the paper. The previous experimental literature on Bayesian learning on networks (Choi et al. (2005, 2009)) make use of several interesting and parsimonious three-person networks. However, we are unable to borrow these networks for our study as they were not designed for the purpose of separating between DeGroot and Bayesian learning. In fact, the networks utilized in Choi et al. (2005, 2009) lack power to pit Bayesian learning against the DeGroot alternatives posited above. Panel A of Table 2 shows the fraction of observations that differ across complete information Bayesian learning and the DeGroot alternatives for each of the three networks used in Choi et al. (2005) and Choi et al. (2009). In two of the networks, there are no differences between the equilibrium paths of Bayesian learning and each of the DeGroot alternatives and in the third network the differences are on the order of 15% of the observations.

Given our goal of separating between Bayesian and DeGroot alternatives, we move to an environment with seven agents as opposed to three agents, so that we obtain more power to distinguish between these models while still maintaining computational tractability.⁵

We considered all connected, undirected networks with seven nodes. Next, we established a model selection criterion function. This criterion function was a divergence function which

⁵Moving to seven agents, for instance, would be exponentially more difficult for our structural estimation.

measures the share of node-time observations for which the Bayesian model (with $\pi = 1$) and a DeGroot model pick different actions,

$$D(G) := \frac{1}{nT} \sum_{w \in W} \sum_{t=1}^T \sum_{i=1}^n \left| a_{i,t}^B(w | G) - a_{i,t}^m(w | G) \right| \cdot P(w | \theta = 1),$$

where $a_{i,t}^B(w|G)$ is the action predicted under the Bayesian model and $a_{i,t}^m(w|G)$ is the action predicted under DeGroot with m -weighting, with m one of uniform, degree, or eigenvector. Moreover, to be able to convey these networks in a very simple manner in the rural setting, we balanced the divergence of the graph with how simply the network could be drawn on a sheet of paper.

The procedure yields the following networks. Network 1 was chosen to maximize power when testing the Bayesian model against uniform and degree weighting. Network 2 provided power in separating all the models. Network 3 was chosen to maximize power when testing the Bayesian model against eigenvector weighting. Note that in network 3, both the uniform and degree weighting models predict the same behavior as Bayesian learning along the equilibrium path.

4. TESTING THE THEORY

In order to test how well a model m fits the data in village r , we will use a divergence function, defined as

$$D(m, r) := \frac{1}{n(T_r - 1)} \sum_{i=1}^n \sum_{t=2}^{T_r} D_{i,t,r}^m$$

where $D_{i,t,r}^m = |a_{i,t,r}^{obs} - a_{i,t,r}^m|$. This computes the share of actions taken by players that are not predicted by the model m .⁶ To examine how poorly model m predicts behavior over the entirety of the data set, we define the *divergence function* as

$$\mathbf{D}(m) := \frac{1}{R} \sum_{r=1}^R \frac{1}{(T_r - 1)n} \sum_{i=1}^n \sum_{t=2}^{T_r} D_{i,t,r}^m.$$

Model selection will be based on the minimization of this divergence metric, which can also be thought of as the fraction of actions not explained by model m .

The divergence is the deviation of the observed data from the theory. We may define the action prescribed by theory in one of two ways: network level, which looks at the social learning process as the unit of observation, and individual level, in which the unit of observation is the individual.

When studying *network level divergence*, we consider the entire learning process as a single observation. Theory predicts a path of actions under the true model for each individual in each period given a network and a set of initial signals. This equilibrium path

⁶Since all models and all empirical data have a fixed first stage (given by the signal endowment), the first round should not enter into a divergence metric. In turn, we restrict attention to $t \geq 2$.

that model m predicts is gives the theoretical action, $a_{i,t,v}^m$. When using this approach, we try to assess how the social learning process as a whole is explained by a model. This method maintains that the predicted action under m is not path-dependent and is fully determined by the network structure and the set of initial signals.

When we consider the individual level divergence, the observational unit is the individual. The action prescribed by theory is conditional on the information set available to i at $t - 1$: $a_{i,t,v}^m$ is the action predicted for agent i at time t in treatment r given information set $I_{i,r,t}$.

4.1. Learning at the Network Level. We begin by treating the observational unit as the social network itself and take the whole path of actions $\{a_{i,r}\}_{r \leq R}$ as a single observation.

In our incomplete information Bayesian learning model, we consider three natural candidates for π (the probability that an agent is Bayesian as opposed to DeGroot). First, we can consider the parameter that minimizes the network level divergence: the best-fitting value of π , given by $\hat{\pi} = \operatorname{argmin} \mathbf{D}(m, \pi)$. Second, we can consider the limit as $\pi \rightarrow 1$. This is the limiting case in which all agents are almost certain that everyone is Bayesian. Third, we can consider the limit as $\pi \rightarrow 0$.

We find that the optimal $\hat{\pi}$ is a corner solution. In our dataset, the network level divergence is invariant to π .⁷ In particular, while the beliefs change with π , the threshold for choosing an alternative action is never crossed. In turn, the optimal value corresponds to $\pi \rightarrow 0$ or $\pi \rightarrow 1$ as a function of which ever has the lower divergence. Thus, we are relegated to a situation identical to that of the complete information environment, wherein we can consider the two cases of all agents being Bayesian and all agents being DeGroot.

Table 3 presents the network level divergence for each of the three DeGroot models as well as the complete information Bayesian model. Across all the networks, uniform weighting fails to explain 12% of the data, degree weighting fails to explain 14% of the data, complete information Bayesian learning fails to explain 19% of the data, and eigenvector centrality fails to explain 27% of the data. This suggests that the degree and uniform DeGroot models as well as the Bayesian learning models each explain between 60 to 80% of the observations.

Uniform weighting dominates each of the other models in terms of fit. Meanwhile, eigenvector centrality weighting performs uniformly worse. The data shows that these models outperform the eigenvector weighting model considerably: they explain approximately 84% of all agent-round observations, while the latter explains less than 67% of these observations.

As the optimal $\hat{\pi}$ are corner solutions, we need only compare the complete information Bayesian learning model to each of the DeGroot alternatives. To be able to statistically test the difference of fits across these different models, we conduct a non-nested hypothesis test in a GMM setting (e.g. Rivers and Vuong, 2002). Specifically, we perform a nonparametric

⁷Details are provided in a supplementary appendix, available upon request from the authors.

bootstrap at the village-game level wherein we draw, with replacement, a 22 village-game blocks of observations, and compute the network level divergence. We then create the appropriate test statistic, which is a normalized difference of the divergence functions from the two competing models.

Our key hypothesis of interest is a one-sided test with the null of Bayesian learning against the alternative of the DeGroot model. Table 4 contains the results of the inference procedure. First, when we look across all topologies, we find evidence to reject the Bayesian model in favor of the uniform weighting alternative and the degree weighting alternative.

Finally, we note that eigenvector centrality is dominated. When looking across all networks, it is summarily rejected in favor of any alternative model. Looking at the bootstrap for a given topology, we find that only in the case of network 1 against Bayesian learning and network 3 against uniform weighting does eigenvector weighting remain statistically indistinguishable.

Ultimately, the bootstrap provides strong evidence that the null of Bayesian learning is rejected in favor of the degree weighted DeGroot alternative, suggestive evidence that Bayesian learning underperforms relative to the uniform weighted DeGroot alternative, and support for the claim that eigenvector weighting explains the data quite poorly.

4.2. Learning at the Individual Level. Having looked at the network level divergence, we turn our attention to individual level divergence. While this does not purely address the mechanics of the social learning process as a whole, it does allow us to look at individual learning patterns. Understanding the mechanics of the individual behavior may help us microfound the social learning process.⁸

4.2.1. DeGroot Models. We begin by calculating the individual level divergence for the DeGroot models.⁹ Panel A of Table 5 contains the results of the exercise. First, uniform weighting systematically outperforms eigenvector weighting (by a large margin) and degree weighting (by a smaller margin). It is worth noting how well the DeGroot models perform in terms of predicted individual behavior. Across all three networks, the uniform weighting model explains approximately 88.9% of all individual observations. Degree and eigenvector centrality weighting models predict 86% and 75% of all individual observations, respectively.

Furthermore, we look at the individual level divergence village by village. Panel B of Table 5 shows the number of villages for which each DeGroot model was the best fitting. For network 1, uniform weighting is the best fitting (17 villages), for network 2 both uniform and degree weighting are tied (11 villages each), and for network 3 uniform weighting is by far the best fitting model (21 villages).

⁸It may be the case that agents themselves do not each behave in according to a particular model while the aggregate social group may best be described by such a model

⁹When an agent faces a tie, they stay with their previous action. We considered a random tie-breaking alternative as well, which does not change the results much.

4.2.2. *Bayesian Learning with Incomplete Information.* We now turn our attention to the Bayesian learning model. Unlike the myopic models, when considering the empirical divergence and the subsequent predicted action $a_{i,t,v}^m$, we need to consider the whole path of observed actions for all agents.

A potential problem arises: since our model of Bayesian learning implies that actions taken by individuals are deterministic functions of the underlying environment, this implies that the support of the set of potential paths that individuals could have observed is rather limited. Therefore, there is a possibility that empirically, Bayesian agents may arrive to an information set that has zero probability of occurrence. This is problematic for identification, since the Bayesian learning model is mute when agents have to condition their inference on zero probability events; any observed action from then on would be admissible for a Bayesian learning agent. Table 6 shows that zero probability information sets are hit quite frequently.

With any DeGroot alternative, 100% of the treatments in networks 1 and 2 have at least one agent hitting a zero-probability information set. Moreover, at least 62% of players in these networks have hit a zero-probability information set at some point. Though one may be tempted to interpret the lack of support itself as evidence against a Bayesian model, this is a delicate issue requiring a more careful treatment.

To compute the divergence across all observations, we need to make the support of possible paths extend over our entire data set. The usual way to deal with this problem is to introduce disturbances. In the following subsection we explore the possibility of estimating a trembling hand or quantal response equilibrium (QRE) style version of Bayesian learning in which we introduce the possibility of making mistakes by all agents. In such a model, individuals can make mistakes with some probabilities, and Bayesian agents, knowing the distribution of these disturbances, integrate over this possibility when updating.

4.2.3. *Bayesian Learning with Disturbances and Complexity Problems.* To simplify exposition, we restrict attention to the case of a complete information Bayesian model where each agent is Bayesian. In this environment, each agent makes a mistake with probability ε and chooses the opposite action that a Bayesian agent would chose. This guarantees full support; any agent can take any action given any history with positive probability.¹⁰ Moreover, one could potentially describe what a Bayesian learning model prescribes by taking the limit $\varepsilon \rightarrow 0$.

Introducing disturbances comes at great computational cost in an environment where agents learn on networks. The only sufficient statistic for the information set that each agent sees is the information set itself, as there is no deterministic function between signal

¹⁰Haile et al. (2008) show that QRE imposes no falsifiable restrictions and can rationalize any distribution of behavior in normal form games. Relating this intuition to our context, one may be able to pick a distribution of ε such that it rationalizes the incomplete information Bayesian model as describing the data well.

endowments and information sets. This means that through time, the relevant state space (the histories each agents could have seen) grows exponentially. We show that this makes the problem intractable for any practical purpose.

First, we note that the algorithm that we use to simulate the Bayesian learning model without trembles is computationally “tight” in the sense that, asymptotically, there is no faster algorithm.¹¹ Because any algorithm would have to take order T steps to print output for each of the T periods, an algorithm that is $O(T)$ is asymptotically tight.

PROPOSITION 4.1. *The algorithm for computing Bayesian learning with no disturbances is $O(T)$. Moreover, it is asymptotically tight; i.e. any algorithm implementing Bayesian learning running time must be at least $O(T)$*

Specifically, the algorithm is $\Theta(n4^nT)$.¹² Notice that if n was growing this algorithm would be exponential time. Second, we show that the extension of this algorithm to an environment with disturbances is computationally intractable.

PROPOSITION 4.2. *Implementing the Bayesian learning algorithm with disturbances has computational time complexity of $\Theta(4^{nT})$.*

To see how computationally intractable is this algorithm, take as an example our experimental design. Assume that the original code takes one second to run. With $n = 6$ and $T = 5$ the computational time is on the order of 6.9175×10^{18} seconds, which is 8.0064×10^{13} days or 2.1935×10^{11} years. To get some perspective, let us compare the number of calculations with this very simplistic algorithm using the Choi et al. (2005, 2009) environment in which $n = 3$. In this setup, the expected time would be 1.2288×10^5 seconds or 1.42 days which is entirely reasonable for structural econometrics.

In the above exercise, we used the most natural algorithm and one that was efficient for the case without disturbances; an objection may be made that there could perhaps be a more efficient algorithm. The decision problem we are interested in is determining whether an agent i in time period t given a history always picks the same action under a proposed algorithm as under the Bayesian model with trembles. We conjecture that the problem is NP-hard, which we are investigating in ongoing work (Chandrasekhar, Cornejo, and Xandri, 2011). This means that the computational problem is at least as hard as NP-complete problems.¹³ Whether there may or may not be polynomial time solutions for NP-hard problems is open; if $P \neq NP$, then none would exist. The computer science literature studying Bayesian learning networks shows that obtaining the probabilities is NP-hard (Cooper, 1990) in any given network of events. In this context the networks

¹¹Our environment consists of finite graphs where n does not grow in T .

¹²Recall that we say $f_1(n) \in \Theta(f_2(n))$ if f_1 is asymptotically bounded above and below by f_2 , upto a constant scale. Formally, if $\exists c_1, c_2 > 0, \underline{n}$ such that $\forall n > \underline{n}, c_1 \cdot |f_2(n)| < |f_1(n)| < c_2 \cdot |f_2(n)|$.

¹³A problem is said to be NP-complete if (a) it is NP which is to say that a given solution can be verified in polynomial time and (b) it is NP-hard so that any NP problem can be converted to this problem in polynomial time.

are networks of events. Translating our framework into this setup involves constructing a network of belief states for each individual in the network and each time period, so a node in the Bayesian learning network would be a pair (i, t) , so the size of it would be NT . Our ongoing work seeks to extend their argument to our decision problem which involves checking that the action taking by each person in each time period is identical when comparing a proposed algorithm with the true Bayesian learning model. The intuition is that the learning network is growing linearly in the number of periods and individuals and therefore for any algorithm there can be some action sequence such that to be able to decide whether individual i at time t , given the history, needs to decide whether to guess 0 or 1, one needs all the probabilities. Based on Cooper (1990), which applies to a broader class of networks (and therefore will have weakly worse complexity), we conjecture that the argument for our sub-class of networks will also be NP-hard.

4.2.4. *Results.* We have argued that estimating the Bayesian model with trembles has computational complexity constraints. In turn, we now turn to studying which model best fits the data, taking these constraints into account. We are interested in studying which model best fits the data. We look at the deviation of each agent’s action, given the history that the agent has observed at that time, from the predicted action by the model for that agent given the history. The formalities are developed in Appendix B.

To compute the individual level divergence for the model with incomplete information, we need an estimate of π . Again we use the filter described in Appendix C to compute the probability that an agent i in village v is a Bayesian learner. Recall that we have three natural candidates for π : the best-fitting value, given by $\hat{\pi} = \arg \min \mathbf{D}(m, \pi)$, the limit as $\pi \rightarrow 1$, and the limit as $\pi \rightarrow 0$.

Since guaranteeing full support in this model by reintroducing trembles induces computational problems, we make the following arguments regarding the relative performance of the Bayesian model. First, the fact that we repeatedly observe agents facing zero probability events, even when there is positive probability that agents may be behaving in another manner, may be taken as prima facie evidence supporting the idea that this model of Bayesian learning with incomplete information on networks fails to explain the experimental data.

Second, one could make the objection that the considered incomplete information Bayesian model is not sufficiently rich to capture the characteristics of the data and that, perhaps, one needs a more nuanced model. This could indeed be the case, but as demonstrated in Proposition 4.2, it would be computationally infeasible to estimate a model generating full support.

Third, it might be the case that we have the right model of incomplete information Bayesian model but we lack a theory of what individuals do once they hit zero probability events. This implies that we assume the existence of a correct set of beliefs when encountering zero probability events that rationalizes individuals’ actions. If this is the case we may

take two different approaches. First, we could be agnostic about the correct off equilibrium beliefs. Second, we could consider the case of a richer Bayesian model that rationalizes the actions taken after an agent hits a zero probability event and precisely matches the supposed off equilibrium behavior. Such a model, of course, has the degree-of-freedom problem.

We begin with the first approach, by being agnostic about the off-equilibrium behavior and instead restrict attention to observations for which agents were in an information set that had a positive probability of occurrence. This is the only feasible comparison we can do given our assumption and agnosticism about the off-equilibrium beliefs. In this subset of observations, we can calculate the individual level divergence, since Bayes' rules applies and the Bayesian learning models gives us a concrete prediction. Of course, we have not eliminated observations at random, but rather we have eliminated those that were not in accordance to the Bayesian learning model equilibrium (i.e. those that happened with zero probability). This is an admittedly ad hoc approach, requiring the assumption that the DeGroot model does not perform sufficiently worse off-equilibrium (where the Bayesian model in principle could rationalize anything). We will return to this issue below.

Under such an assumption, if it turns out that even in this subset of observations, Bayesian performs worse than the alternative myopic models considered, then this would be further evidence that, at the individual level, the Bayesian learning model would not seem to fit the experimental data well. Based on this idea, we present the calculation of the individual level divergence measure for observations that were in the support of the model. Table 7 contains the results of this exercise. It shows that across all networks, as well as for each network, both degree and uniform weighting DeGroot models have a lower divergence than the Bayesian model. Specifically, it suggests that a model with a limiting $\pi \rightarrow 0$ dominates the corresponding model with $\pi \rightarrow 1$.

As in the case of the network level analysis, the individual level divergence does not vary with π , as long as $\pi \in (0, 1)$. That is, by varying the ex-ante probability of any given agent being Bayesian, and as long as it is strictly different from 0 or 1 (as we have assumed throughout), a Bayesian agent would choose the same action on identical information sets. Upon inspection, this is not because there is no actual effect on the inference of each agent on the true state of the world θ , but instead because the change in inference is not sufficiently strong that it changes a guess from 0 to 1 (or vice versa). Of course, there are equilibrium considerations here. Even if, in principle, varying the parameter π affects the decision, if no one chooses different actions for different π 's, then there cannot be any difference in the inference made by agents on the actions of others.

To be able to perform inference on the null hypothesis of the Bayesian learning model against the alternative DeGroot models, again we perform a nonparametric bootstrap. Our main hypothesis of interest is a one-sided test of the Bayesian learning null against the

DeGroot alternatives. Table 8 shows us that both the complete and incomplete information Bayesian models are strongly rejected against the alternatives of uniform and degree weighting DeGroot models. Moreover, as in the case of the network level analysis, we find that eigenvector centrality is uniformly rejected against all alternative models. Third, and perhaps more puzzling, at the individual level we find that uniform weighting is a better description of the behavior than degree weighting; the result holds as significant at a 10% one-sided test. While the evidence is somewhat weak, it still overturns the finding in the network level case wherein degree weighting performed better than uniform weighting. This suggests that while the entire process may be best described by degree weighting, conditional on an information set, an individual's action may be best predicted by uniform weighting. That being said, the divergence for each of these models is quite low, suggesting that preferential weighting is a strong fit as well.

We now return to the second approach. If we assume that indeed we have the right model of incomplete information Bayesian model but we are simply missing the correct off equilibrium behavior, we could consider the case that a richer Bayesian model could be the one that rationalizes the actions taken after an agent hits a zero probability event and precisely matches the supposed off equilibrium behavior. Notice that even if, for short T , the Bayesian model might be under-performing, with probability 1 this will be the opposite in the long run, inverting the present results. This follows because, if we consider the Bayesian model as rationalizing anything off-equilibrium, once we are off equilibrium, as $t \rightarrow \infty$, Bayesian would never be penalized while DeGroot will be penalized infinitely often (assuming the behavior does not identically match the DeGroot model for all but finitely many periods).

To summarize this section's results, first we have presented evidence that the considered model of Bayesian learning result arrives at zero probability information sets extremely often. This can be taken as evidence against these particular models. Second, we provide computational theory that shows that models with trembles, which would smooth out the zero probability information set problem, are of little practical use to structurally evaluate empirical data. In turn, methodologically, structural approaches must restrict themselves to models which allow for zero probability information sets. Third, we take a pass at the data by ignoring the off-equilibrium information sets. This evidence suggests that, when restricting the analysis to positive probability information sets, the divergence minimizing models have in the limit no Bayesian agents. Finally, we point out that this approach, while ad hoc, may be inappropriate for a model of incomplete information behavior wherein the off-equilibrium behavior is well-matched. However, assuming the researcher is interested in incomplete information models (because of the computational infeasibility of QRE models), the argument in favor of the ad hoc approach is rejected only if the researcher believes in an untestable model which performs well off equilibrium (since we know on equilibrium it performs poorly). But such a model is unlikely to be the empirically relevant object.

A natural exercise may be to project the intuitions of these models into a regression framework. Such a reduced form approach would allow us to avoid computational constraints and see whether the patterns in regression provide suggestive evidence of DeGroot learning as opposed to Bayesian learning. We discuss the prospect of this in the following section.

5. REDUCED FORM

In this section, we discuss two reduced form approaches to study the data. Our motivation for this exercise is twofold. First, given the computational limits of the structural approach, we are interested in seeing whether reduced form patterns of Bayesian learning (as opposed to DeGroot learning) may be obtained from the data. Second, since larger networks, such as those found in empirical data sets, do not lend themselves to structural approaches for computational reasons, it is worth looking into the effectiveness of reduced form approaches to address these questions.

The central intuition we focus on has to do with double counting information. Under any of the aforementioned Bayesian models, Bayesian agents should not double-count information. However, DeGroot agents do double-count information, and it is on this intuition that we build the exercise.

We provide two examples of regressions which researchers may run. The first set of regressions explores whether individuals overweight the same information if they receive it through multiple channels. The second set of regressions explore whether individuals treat old information that cycles back to them as if it is new, additional information. The null in these regressions is Bayesian model, since one would assume that the relevant parameters ought to be zero. Thus, a rejection of a zero may provide evidence in the direction of the DeGroot rules of thumb. The empirical data shows that both these reduced form analyses seem to provide support in favor of the DeGroot alternatives. However, because we are able to simulate out the data under the null, we show that these intuitions are wrong. Specifically, when we simulate social learning data under the Bayesian null, the coefficients are not as one may have expected.

5.1. Multiplicity. We define a variable which is a dummy for whether individual i makes a guess of 1 in the final period T , $y_i := 1 \{a_{i,T} = 1\}$. As before, d_i is the degree of individual i and N_i is the set of (direct) neighbors $N_i = \{j \in V : ij \in E\}$. Note that $d_i = |N_i|$. Moreover, N_{2i} is the set of second-neighbors of person i ; that is, $j \in N_{2i}$ means that there is at least one path of length two between i and j , but no path of length one. Finally, we define N_{2i}^l to be the set of second neighbors to whom she has exactly l paths.

The first regression we run is of the form

$$(5.1) \quad y_i = \beta_0 + \beta_1 s_i + \beta_2 \mathbb{E}_{N_i} [s_j | j \in N_i] + \sum_l \beta_{3l} \mathbb{E}_{N_{2i}^l} [s_j | j \in N_{2i}^l] + \epsilon_i.$$

This is a regression of whether or not individual i ultimately makes a guess of 1 on whether the individual's signal is 1 (s_i) the share of ones ($\mathbb{E}_{N_i}[s_j|j \in N_i]$) in individual i 's neighborhood, and the share of ones given to each subset of second neighbors to whom i has exactly l paths ($\mathbb{E}_{N_{2i}^l}[s_j|j \in N_{2i}^l]$).

The interpretation is as follows. β_2 measures the impact of her neighborhood receiving a greater share of ones on an individual's guess. We expect $\beta_2 > 0$. Moreover, β_{3l} measures the impact of the subset of her second-neighborhood with multiplicity l . The intuition is that as the signals of individuals with greater multiplicity ought not be double-counted under a Bayesian frame, $\beta_{3l+1} > \beta_{3l}$ would be evidence of overweighting redundant information that has arrived via multiple channels, while $\beta_{3l+1} = \beta_{3l}$ would provide evidence in favor of the Bayesian hypothesis.

Table 9 provides the simulation and empirical results. Looking at the empirical results, as expected, an individual's own signal being one increases the probability of the final guess being one by 38pp, in a linear probability model framework. Moreover, if the share of individuals in one's neighborhood with signals of one increase by 10%, an agent has a 4.6pp increase in the probability of her final guess being one. While this seems to be consistent with the intuition that agents engage in double-counting, the simulation exercise shows that these patterns cannot be interpreted in that manner.

Given the learning model, the network structure, and signal endowment, we simulated out the learning path and then ran the relevant regressions. We present results when simulating the learning process from the complete information Bayesian model (every agent is Bayesian) as well as each of the three DeGroot alternatives. The table shows that the Bayesian null does not have coefficients that are near identical across multiplicities 1, 2 and 3. Furthermore, the increasing correlation with indirect friends of higher multiplicities is also not uniformly found across the DeGroot models. Ultimately, the regressions suggest that the linear projection of this learning process is complex and may depend crucially on the network structure and set of initial signals.

5.2. Historical Information. Another reduced form that one may look at is whether individuals re-incorporate historical information that they have previously observed. Consider an individual at period 3. They have observed both their own signals and the signals of their direct neighbors (insofar as the first period guesses of their neighbors will be identical to their signals). In period three, therefore, a Bayesian individual's guess should not re-incorporate this information. Instead, it should only update using information about second-neighbors and the like, about whom they have yet to receive information.

To examine this formally, we perform the following regression. We regress the period three guess of individual i on her own signal (s_i) and the average signal of her neighborhood ($\mathbb{E}_{N_i}[s_j|j \in N_i]$) which she would have seen in period three. We also include as regressors the average signal of second neighbors ($\mathbb{E}_{N_{2i}}[s_k|k \in N_{2i}]$) which should be new information in period three. Lastly, we include the average signal of direct neighbors whose signals

can cycle back via a path of length two back to individual i . Of course, we also include the agent herself in this set. (Formally, we use $\mathbb{E}_{C_i}[s_j | j \in C_i]$, where $C_i = \{j \in V - \{i\} : A_{ij}^2 A_{ij} > 0\} \cup \{i\}$.) The regression is as follows.

$$(5.2) \quad y_i = \alpha_0 + \alpha_1 s_i + \alpha_2 \mathbb{E}_{N_i}[s_j] + \alpha_3 \mathbb{E}_{N_{2i}}[s_k | k \in N_{2i}] + \alpha_4 \mathbb{E}_{C_i}[s_j | j \in C_i] + \epsilon_i.$$

We test the hypothesis of whether $\alpha_4 = 0$, which is our Bayesian null. Notice that $\alpha_4 > 0$ provides evidence that individuals reincorporate information that they already knew as it cycles through the network.

Table 10 presents the simulation and empirical results. The empirical results show that $\alpha_4 > 0$ and $\alpha_3 > 0$. While this seems to provide suggestive evidence for the intuition that DeGroot weighting reincorporates old information, the simulation results provide evidence that for our environment $\alpha_3 < 0$ for some of these models and $\alpha_4 > 0$ even for the Bayesian model.

5.3. Reflection on Reduced Forms. Taken together, Tables 9 and 10 have shown that natural reduced form approaches to test between these models may be misguided without first checking whether the patterns by the learning processes actually match the intuitions. We are able to study the reduced form projections of the Bayesian model using our simulation algorithm and find evidence that, when projected onto a regression for these networks with this environment, the Bayesian data suggests that the coefficients can deviate greatly from our intuitions. This, we argue, provides a strong motivation for the structural approach to studying the models.

6. CONCLUSIONS

In this paper we have investigated whether social learning patterns on small networks resemble those which would emerge if agents behaved in a Bayesian manner or if they are better modeled with DeGroot alternatives which are myopic and more simplistic models. To do so, we developed a simple experiment on networks that were designed to distinguish between these models, large enough to give us power on this dimension, but small enough to ensure that simulating a Bayesian learning on networks model was not computationally intractable. Given our experimental data, we were able to study the social learning process as a whole by taking the network as the unit of observation and study the behavior of an individual, which addresses whether an agent acts in a Bayesian manner.

At the network level we find evidence that the degree weighting DeGroot model best explains the data. The Bayesian learning null is rejected in favor of this alternative model. However, we maintain that Bayesian learning did an adequate job of describing the data process, explaining (beyond pure random guessing) 62% of the actions taken as compared to 76% by the DeGroot alternative.

At the individual level we find that uniform weighting DeGroot performs the best and both uniform and degree weighting DeGroot models outperform the Bayesian model. However, we show that the Bayesian model encounters the problem that many individuals come across zero probability information set. First, this provides suggestive evidence of the lack of fit of this incomplete information Bayesian model. Second, we demonstrate that introducing disturbances to smooth out the distribution cannot be a solution in this environment. The computational complexity of the problem is damaging to the very approach of applying QRE or trembles to the Bayesian learning on networks environment. As such, we recommend that researchers focus on computationally tractable models which will be easier to falsify.

We also show that reduced form approaches may be problematic. We provide two natural examples of regressions which build on intuitions separating DeGroot and Bayesian learning patterns. Equipped with our Bayesian learning algorithm, we simulate learning data from the Bayesian model as well as from DeGroot models and show that the reduced form regression outcomes do not conform to the intuitions.

Ultimately, the findings suggest that agents and the learning process as a whole may better be thought of as coming from DeGroot action models where individuals myopically weight their neighbors' actions when updating their own beliefs. This may imply that social learning processes empirically may be sub-optimal, with information often getting stuck in pockets of the network. Having constructed an example of a network which satisfies asymptotic learning for DeGroot communication models, but where asymptotic learning fails for DeGroot action models, we argue that in action-learning environments DeGroot processes may be more damaging to the wisdom of society than previously anticipated.

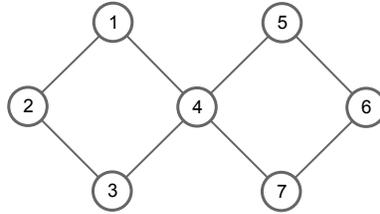
REFERENCES

- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2010): "Bayesian Learning in Social Networks," .
- ANDERSON, L. AND C. HOLT (1997): "Information cascades in the laboratory," *The American Economic Review*, 87, 847–862.
- BALA, V. AND S. GOYAL (1998): "Learning from neighbours," *Review of Economic Studies*, 65, 595–621.
- BANERJEE, A. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 797–817.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. JACKSON (2010): "Social networks and microfinance," MIT working paper.
- BERGER, R. (1981): "A necessary and sufficient condition for reaching a consensus using DeGroot's method," *Journal of the American Statistical Association*, 76, 415–418.

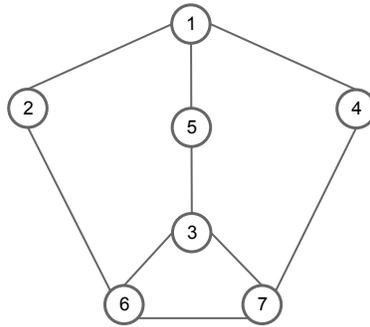
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of political Economy*, 100.
- CELEN, B. AND S. KARIV (2005): “An experimental test of observational learning under imperfect information,” *Economic Theory*, 26, 677–699.
- CHANDRASEKHAR, A. AND R. LEWIS (2010): “Econometrics of sampled networks,” MIT working paper.
- CHATTERJEE, S. AND E. SENETA (1977): “Towards consensus: Some convergence theorems on repeated averaging,” *Journal of Applied Probability*, 89–97.
- CHOI, S., D. GALE, AND S. KARIV (2005): “Behavioral aspects of learning in social networks: An experimental study,” *Advances in Applied Microeconomics: A Research Annual*, 13, 25–61.
- (2009): “Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data,” *WP, May*.
- CONLEY, T. AND C. UDRY (2010): “Learning about a new technology: Pineapple in Ghana,” *The American Economic Review*, 100, 35–69.
- DEGROOT, M. (1974): “Reaching a consensus,” *Journal of the American Statistical Association*, 69, 118–121.
- DEMARZO, P., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions*,” *Quarterly journal of economics*, 118, 909–968.
- ELLISON, G. AND D. FUDENBERG (1993): “Rules of thumb for social learning,” *Journal of Political Economy*, 101, 612–643.
- (1995): “Word-of-mouth communication and social learning,” *The Quarterly Journal of Economics*, 93–125.
- FRIEDKIN, N. AND E. JOHNSEN (1997): “Social positions in influence networks,” *Social Networks*, 19, 209–222.
- GALE, D. AND S. KARIV (2003): “Bayesian learning in social networks,” *Games and Economic Behavior*, 45, 329–346.
- GOLUB, B. AND M. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–149.
- HAILE, P., A. HORTAÇSU, AND G. KOSENOK (2008): “On the empirical content of quantal response equilibrium,” *The American Economic Review*, 98, 180–200.
- HUNG, A. AND C. PLOTT (2001): “Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions,” *American Economic Review*, 91, 1508–1520.
- JACKSON, M. (2008): *Social and economic networks*, Princeton University Press.
- JACKSON, M., T. BARRAQUER, AND X. TAN (2010): “Social Capital and Social Quilts: Network Patterns of Favor Exchange,” .
- KRACKHARDT, D. (1987): “Cognitive social structures,” *Social Networks*, 9, 109–134.

- KUBLER, D. AND G. WEIZSACKER (2004): “Limited depth of reasoning and failure of cascade formation in the laboratory,” *Review of Economic Studies*, 71, 425–441.
- MÖBIUS, M., T. PHAN, AND A. SZEIDL (2011): “Treasure Hunt,” *Mimeo*.
- SMITH, L. AND P. SORENSEN (2000): “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398.

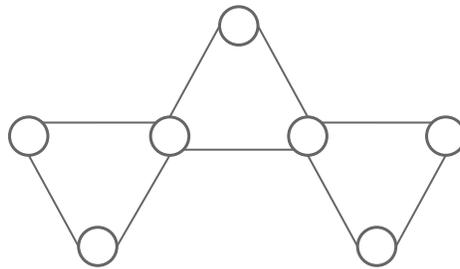
FIGURES AND TABLES



Panel A: Network 1



Panel B: Network 2



Panel C: Network 3

FIGURE 3. Network structures chosen for the experiment.

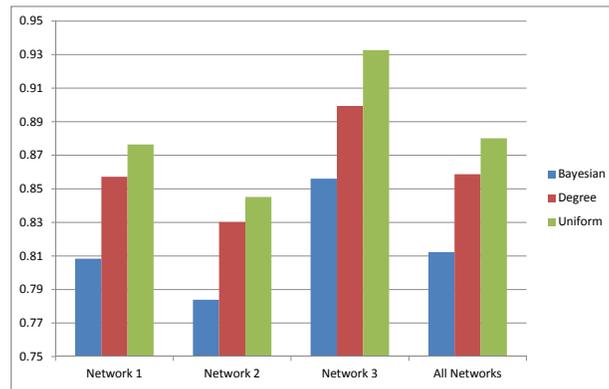


FIGURE 4. Network structures chosen for the experiment.

Table 1: Asymptotic Lack of Learning

<i>Signal Prob.</i>	<i>Fraction Stuck</i>		<i>Extent of Learning</i>	
	(1)	(2)	(3)	(4)
0.51	41.82%	55.06%	7.18%	14.64%
0.55	35.69%	48.80%	9.31%	20.69%
0.6	28.40%	41.40%	11.60%	29.00%
0.65	21.74%	33.80%	13.26%	37.89%
0.7	15.85%	27.14%	14.15%	47.17%
0.75	10.86%	21.30%	14.14%	56.56%
0.8	6.82%	16.42%	13.18%	65.90%
0.9	1.62%	9.59%	8.38%	83.80%

Note: Column (1) presents a lower bound on the fraction of nodes who get stuck and column (2) presents an upper bound. Column (3) presents an upper bound on the fraction of nodes that are right in the limit less p , the fraction of agents who receive the right signal in the beginning. Column (4) presents an upper bound on the fraction of agents who began with the wrong initial signal who ended up with the right guess in the limit.

Table 2: Fraction of Observations that Differ with the Bayesian Model

<i>Panel A: Networks from Choi et al. (2005, 2009)</i>						
<i>Network</i>	<i>Total Divergence</i>			<i>Divergence in Final Period</i>		
	<i>Uniform</i>	<i>Degree</i>	<i>Eigenvector</i>	<i>Uniform</i>	<i>Degree</i>	<i>Eigenvector</i>
1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
3	9.37%	21.87%	8.98%	12.67%	18.67%	7.67%

<i>Panel B: Networks Selected in This Paper</i>						
<i>Network</i>	<i>Total Divergence</i>			<i>Divergence in Final Period</i>		
	<i>Uniform</i>	<i>Degree</i>	<i>Eigenvector</i>	<i>Uniform</i>	<i>Degree</i>	<i>Eigenvector</i>
1	6.89%	10.64%	6.89%	9.13%	13.89%	9.13%
2	9.91%	10.00%	10.48%	15.08%	13.66%	12.10%
3	8.75%	9.67%	9.67%	12.24%	13.18%	13.18%

Notes: Fraction of observations for which the complete information Bayesian model differs with the DeGroot alternative. In Panel A, network 1 is the “complete network”, network 2 is the “directed network”, and network 3 is the “incomplete network” of Choi et al. (2005, 2009).

Table 3: Network Level Divergence

<i>Network</i>	<i>Total Obs</i>	<i>Bayesian</i>	<i>Uniform</i>	<i>Degree</i>	<i>Eigenvector</i>
All Networks	9,205	0.1878	0.1198	0.1413	0.2703
1	3,045	0.1917	0.1236	0.1428	0.2229
2	3,031	0.2161	0.1548	0.1698	0.3026
3	3,129	0.1440	0.0673	0.1006	0.2909

Notes: Network level divergence for the complete information Bayesian model, uniform DeGroot weighting, degree DeGroot weighting, and eigenvector DeGroot weighting.

Table 4: Significance Tests for Network Level Divergence

H_0	H_a	Informative Treatments	Informative Observations	All Networks	Network 1	Network 2	Network 3
Bayesian	Degree	71.5385	0.251	0.3787*** 0.0001	0.2010*** 0.0041	0.2900*** 0.0019	0.1654** 0.0133
Bayesian	Uniform	65.7692	0.1955	0.5565*** 0.0001	0.2923*** 0.0001	0.3767*** 0.0001	0.2814*** 0.0001
Bayesian	Eigenvector	80.7692	0.3333	-0.7724*** 0.9999	-0.1579* 0.9397	-0.5671*** 0.9999	-0.6673*** 0.9999
Degree	Uniform	29.2308	0.2092	0.2866*** 0.0006	0.1556* 0.0615	0.1140* 0.0696	0.2041*** 0.0001
Degree	Eigenvector	69.6154	0.2847	-1.2291*** 0.9999	-0.4144*** 0.9999	-0.8441*** 0.9999	-0.9329*** 0.9999
Uniform	Eigenvector	70.3846	0.319	-1.3691*** 0.9999	-0.5151*** 0.9999	-0.8849*** 0.9999	-0.8849*** 0.9999

Notes: The test statistic is the normalized difference in the divergence functions of the null and the alternative model, and below we show the probability that the test statistic is less than 0, estimated via bootstrap with replacement. Denote *** as significant at the one-sided 1% level, ** at the one-sided 5% level, * at the one-sided 10% level.

Table 5a: Individual Divergence for DeGroot Models*Panel A: Divergence Across Networks*

<i>Network</i>	<i>Observations</i>	<i>Divergence</i>		
		<i>Uniform</i>	<i>Degree</i>	<i>Eigenvector</i>
<i>All</i>	<i>9,205</i>	<i>0.0648</i>	<i>0.135</i>	<i>0.1083</i>
<i>1</i>	<i>3,045</i>	<i>0.0699</i>	<i>0.1513</i>	<i>0.0961</i>
<i>2</i>	<i>3,031</i>	<i>0.0788</i>	<i>0.1544</i>	<i>0.1562</i>
<i>3</i>	<i>3,129</i>	<i>0.0386</i>	<i>0.0866</i>	<i>0.0598</i>

Panel B: Share of Best-Fitting Sessions

<i>H₀</i>	<i>H₁</i>	<i>Informative</i>	<i>H0 beats H1</i>	<i>Best model</i>
<i>Uniform</i>	<i>Degree</i>	<i>0.65</i>	<i>0.8225</i>	<i>Uniform</i>
<i>Uniform</i>	<i>Eigenvector</i>	<i>0.4538</i>	<i>0.7712</i>	<i>Uniform</i>
<i>Degree</i>	<i>Eigenvector</i>	<i>0.3577</i>	<i>0.172</i>	<i>Eigenvector</i>

Notes: Panel A shows the individual level divergence across the three networks for each of the three DeGroot models.

In Panel B, the first column indicates the share of informative sessions over all sessions, and the second column indicates the share of informative sessions where the H0 model beats the H1 model.

Table 5b: Lack of Bayesian Individual Behavior

H_i	Observations	Share Bayesian
<i>Degree</i>	46	17.39%
<i>Uniform</i>	98	17.35%
<i>Eigenvector</i>	98	17.35%

Note: "Observations" are the number of cases where there are discrepancies between the parent node action (which from $t > 3$ Bayesian prescribes peripheral nodes should follow) and the action that the H_1 model prescribes peripheral nodes. "Share Bayesian" is the share of observations where peripheral nodes indeed follow the parent node action.

Table 6: Zero Probability Information Sets Reached

Panel A: Complete Information Model

Network	% Individuals	% Treatments	% Observations
1	36.36%	54.55%	26.44%
2	54.09%	100.00%	53.00%
3	74.24%	100.00%	46.36%

Panel B: Incomplete Information Model

Degree Weighting Alternative

Network	% Individuals	% Treatments	% Observations
1	36.36%	54.55%	26.44%
2	80.30%	100.00%	51.12%
3	65.15%	100.00%	37.93%

Uniform Weighting Alternative

Network	% Individuals	% Treatments	% Observations
1	36.36%	54.55%	26.44%
2	70.45%	100.00%	43.07%
3	62.88%	100.00%	41.00%

Eigenvector Weighting Alternative

Network	% Individuals	% Treatments	% Observations
1	34.09%	54.55%	22.99%
2	80.30%	100.00%	50.00%
3	66.67%	100.00%	38.89%

Notes: Panel A presents results for the complete information Bayesian model. Panel B presents results for the incomplete information Bayesian model against DeGroot alternatives. % Individuals refers to the fraction of individuals who reach a zero probability information set. % Treatments refers to the fraction of treatments (network x village) that reaches a zero probability information set. % Observations refers to the fraction of individual x time units that reach a zero probability information set.

Table 7: Individual Level Divergence on Information Sets in Support

<i>Model</i>	<i>Network</i>	<i>No. Obs.</i>	$\pi \rightarrow 0$	$\pi \rightarrow 1$
Uniform	1	304	0.337	0.499
	2	337	0.316	0.469
Degree	1	261	0.410	0.540
	2	324	0.337	0.536

Note: Players number of agents that eventually hit a zero probability information set. Treatments denotes the number of experiments for which at least one agent hits a zero probability information set. Observations denotes the number of triples (individual, village, treatment) that is in a zero probability information set.

Table 8: Significance Tests for Individual Level Divergence

H_0	H_a	All Networks	Network 1	Network 2	Network 3
Bayesian	Degree	0.0173	-0.0318	0.0193	0.0405*
		0.3438	0.794	0.357	0.0768
Bayesian	Uniform	0.3652***	0.1933***	0.2709***	0.0762***
		0.0001	0.0001	0.0001	0.0001
Bayesian	Eigenvector	0.1799***	0.1509***	0.011	0.1081***
		0.0001	0.0001	0.4267	0.0001
Degree	Uniform	0.3928***	0.2608***	0.293***	0.1068***
		0.0001	0.0001	0.0001	0.0001
Degree	Eigenvector	0.2295***	0.2054***	-0.0368***	0.1055***
		0.0001	0.0001	0.9999	0.0001
Uniform	Eigenvector	-0.2958***	-0.1325***	-0.3023***	-0.0488***
		0.9999	0.9992	0.9999	0.9999

Notes: The test statistic is the normalized difference in the divergence functions of the null and the alternative, and below we show the probability that the test statistic is less than 0, estimated via bootstrap with replacement. Denote *** as significant at the one-sided 1% level, ** at the one-sided 5% level, * at the one-sided 10% level.

Table 9: Weight on indirect neighbors according to the number of multiple direct neighbors

	Data		Bayesian		Degree		Uniform	
	All (1)	Restricted (2)	All (3)	Restricted (4)	All (5)	Restricted (6)	All (7)	Restricted (8)
Signal	.4974*** [.0218]	.5351*** [.0279]	.3965*** [.017]	.4429*** [.0223]	.5725*** [.0169]	.6204*** [.0211]	.6017*** [.0159]	.677*** [.0201]
Direct	.7735*** [.0311]	.7485*** [.0438]	.775*** [.0323]	.8048*** [.0384]	.9125*** [.0335]	.8913*** [.0425]	.8954*** [.0381]	.9078*** [.0475]
One Way	.3157*** [.0405]	.2609*** [.0542]	.7056*** [.0363]	.8189*** [.0486]	0.0461 [.035]	-0.0614 [.0421]	.1306*** [.0368]	.077* [.0447]
Two Ways	.2177*** [.0355]	.1751*** [.0424]	.3582*** [.0337]	.4096*** [.0338]	.0797* [.0421]	0.024 [.0449]	.1985*** [.0459]	.1841*** [.0519]
TwoWays > OneWay								
T-statistic	-1.97	-1.557	-8.653	-8.136	0.7675	1.925	1.481	2.215
Probability	0.9741	0.9386	1	1	0.2224	0.0287	0.071	0.0146
Observations	1750	1271	1750	1271	1750	1271	1750	1271
R -squared	0.4724	0.373	0.5322	0.4632	0.588	0.5453	0.5974	0.5571

Note: Outcome variable is action in round 3. "Direct" is the average signal of direct neighbors, "One Way" is the average signal of indirect neighbors only through one direct neighbor, and "Two Ways" is the average signal of indirect neighbors through two direct neighbors.

Robust standard errors, clustered at the village by game level, in brackets. * p<.1, ** p<.05, *** p<.01 Odd columns are regressions with all data. Even columns are regressions restricting to treatments that are informative for the comparisons Bayesian - Degree and Bayesian - Uniform.

Table 10: Weight on indirect neighbors according to whether they provide new information

	Data		Bayesian		Degree		Uniform	
	All (1)	Restricted (2)	All (3)	Restricted (4)	All (5)	Restricted (6)	All (7)	Restricted (8)
Signal	.4284*** [.0347]	.4573*** [.0425]	.4729*** [.0286]	.5264*** [.0356]	.4507*** [.032]	.482*** [.0359]	.5694*** [.0307]	.6468*** [.0357]
Direct	.7508*** [.0363]	.7011*** [.0471]	.8095*** [.0285]	.8007*** [.0346]	.8436*** [.0402]	.8167*** [.0489]	.7576*** [.0391]	.7275*** [.0492]
Indirect New	.2455*** [.0331]	.1803*** [.0377]	.3768*** [.0261]	.3489*** [.0317]	.1558*** [.0333]	.0883** [.0384]	.2548*** [.0283]	.2063*** [.0338]
Indirect Repeated	.1715*** [.0416]	.1669*** [.048]	.1029*** [.0354]	.0801* [.0427]	.2463*** [.0398]	.2619*** [.0447]	.1724*** [.0417]	.1554*** [.0492]
Observations	1587	1250	1587	1250	1587	1250	1587	1250
R -squared	0.4628	0.3958	0.4953	0.4135	0.5819	0.5475	0.5945	0.5687

Note: Outcome variable is action in round 3. "Direct" is the average signal of direct neighbors, "Indirect New" is the average signal of indirect neighbors that provide new information, and "Indirect Repeated" is the average signal of indirect neighbors that do not provide new information.

Robust standard errors, clustered at the village by game level, in brackets. * $p < .1$, ** $p < .05$, *** $p < .01$ Odd columns are regressions with all data. Even columns are regressions restricting to treatments that are informative for the comparisons Bayesian - Degree and Bayesian - Uniform.

Table 11: Weight on neighbors according to distance.

	Data		Bayesian		Degree		Uniform	
	All	Restricted	All	Restricted	All	Restricted	All	Restricted
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Signal	.5131*** [.0235]	.5535*** [.0303]	.4118*** [.0219]	.4457*** [.0266]	.5907*** [.0209]	.6599*** [.0259]	.6489*** [.0197]	.7337*** [.022]
Direct	.7902*** [.0345]	.7847*** [.0537]	.9039*** [.0367]	.957*** [.0486]	.9347*** [.0369]	.9666*** [.0508]	.9118*** [.0407]	.9379*** [.0544]
Distance 2	.3261*** [.0398]	.2872*** [.0573]	.6484*** [.0394]	.7207*** [.0518]	.1645*** [.045]	.1392** [.0548]	.2369*** [.0428]	.2264*** [.0523]
Distance 3	.1042** [.043]	0.0764 [.0587]	.3517*** [.0428]	.3947*** [.0533]	.1125*** [.0357]	0.0673 [.045]	.1609*** [.0335]	.1315*** [.0425]
Distance 4	0.0583 [.0656]	0.0266 [.0777]	.15** [.0594]	.1778** [.0733]	0.0295 [.0581]	-0.0278 [.0668]	0.0498 [.05]	0.0081 [.0562]
Test Distance 2 > Distance 3								
T-statistic	5.021	4.289	6.807	7.036	0.9166	1.132	1.319	1.453
Probability	1.20E-06	2.30E-05	4.60E-10	2.00E-10	0.1808	0.1304	0.0951	0.0749
Test Distance 3 > Distance 4								
T-statistic	0.8282	0.8034	3.893	3.686	1.433	1.43	1.809	1.758
Probability	0.2048	0.2119	9.20E-05	2.00E-04	0.0776	0.0782	0.0368	0.0411
Observations	1330	1001	1330	1001	1330	1001	1330	1001
R -squared	0.4822	0.3903	0.5706	0.5078	0.6279	0.5833	0.6442	0.6133

Note: Outcome variable is action in last round. "Direct" is the average signal of direct neighbors, "Distance #" is the average signal of indirect neighbors with distance #.

Robust standard errors, clustered at the village by game level, in brackets. * p<.1, ** p<.05, *** p<.01 Odd columns are regressions with all data. Even columns are regressions restricting to treatments that are informative for the comparisons Bayesian - Degree and Bayesian - Uniform.

APPENDIX A. COMPLETE INFORMATION BAYESIAN ALGORITHM

A.1. Setup. To begin with, suppose that all agents learn about the state of the world using Bayes' rule. Assume that this fact is common knowledge for all agents. We will extend this framework below, but it is the most pedagogical model to help establish notation. It is crucial to keep in mind that the signal endowment w is a sufficient statistic over the actions that agents take, since θ is never observed. In turn, the inference that agents need to do from the other agents' play is concerned with the signal endowment. To proceed we must establish some notation.

We define $\mu_0^\theta(w)$ as the probability of signal endowment w when the true state of the world is θ . Based on the previous assumptions, we have that:

(A.1)

$$\mu_0^{\theta=1}(w) := \Pr(w \mid \theta = 1) = p^{\#\{i:w_i=1\}} (1-p)^{N-\#\{i:w_i=1\}} = p^{\left(\sum_{i=1}^n w_i\right)} (1-p)^{n(1-\sum_{i=1}^n w_i)}$$

(A.2)

$$\mu_0^{\theta=0}(w) := \Pr(w \mid \theta = 0) = p^{\left(\sum_{i=1}^{i=N} \frac{w_i}{n}\right)} (1-p)^{\left(N-\sum_{i=1}^{i=N} \frac{w_i}{n}\right)}$$

Following the same reasoning, define $\mu_{i,t}^\theta(w, \bar{w})$ as the belief probability distribution that agent i has at period t of the game over states $w \in W$, given that the true signal endowment is \bar{w} . Observe that for different signal endowments, the information sets that each agent will observe are clearly going to be different. In turn, the belief over signal endowments that each agent has at each state depends on what was the true signal endowment. Of course, this measure will be constant over signal endowments that generate the same information sets.

Also, define $p_{i,t}(\bar{w})$ as the belief probability that agent i has at stage t of the game puts on the event $\theta = 1$, if the information set he observes comes from the true signal endowment being \bar{w} . Moreover, put $\mathbf{a}_{i,t}(\bar{w}) \in \{0, 1\}$ as the action that agent i takes at stage t of the game, if the information set reached comes from the true signal endowment \bar{w} . The objects that will be iteratively found will be $\mu_{i,t}^\theta(w, \bar{w})$, $\mathbf{a}_{i,t}(\bar{w})$ and $p_{i,t}^\theta(\bar{w})$.

To start the iterations, we need to find these objects at $t = 1$. These are given by

(A.3)

$$p_{i,1}^\theta(\bar{w}) = \begin{cases} p & \text{if } \bar{w}_i = s_i = 1 \\ 1-p & \text{if } \bar{w}_i = s_i = 0 \end{cases}$$

(A.4)

$$\mathbf{a}_{i,1}(\bar{w}) = s_i$$

(A.5)

$$\mu_{i,1}^{\theta=1}(w, \bar{w}) = \begin{cases} 0 & \forall w \in W : w_i \neq s_i \\ \frac{\mu_0^{\theta=1}(w)}{\Pr(s_i|\theta=1)} & \text{if } w_i = s_i \end{cases}$$

(A.6)

$$\mu_{i,1}^{\theta=0}(w, \bar{w}) = \begin{cases} 0 & \forall w \in W : w_i \neq s_i \\ \frac{\mu_0^{\theta=0}(w)}{\Pr(s_i|\theta=0)} & \text{if } w_i = s_i \end{cases}$$

To model what is exactly what each agent observes, in order to see how beliefs are updated over the course of the game, we need to introduce some network notation. In particular, $N(i) := N_i \cup \{i\}$, the set of neighbors of agent i , including i herself. Next, define

$$(A.7) \quad \mathbf{a}_t^{(i)}(\bar{w}) := \left(\underbrace{\mathbf{a}_{j_1, t-1}(\bar{w}), \mathbf{a}_{j_2, t-1}(\bar{w}), \dots, \mathbf{a}_{j_{d(i)}, t-1}(\bar{w})}_{\text{actions taken by neighbors}}, \underbrace{\mathbf{a}_{i, t-1}(\bar{w})}_{\text{own past action}} \right)$$

to be the action profile that agent i sees at the beginning of state t , when the true state of the world is \bar{w} . If we just write $a_t^{(i)}$ (without bold letters) we refer to a particular observed action profile.

A.2. Time $t + 1$ iteration. At time t , we have

$$\begin{aligned} \mu_{i,t}^{(\theta=1)}(w, \bar{w}) &= \Pr(w \mid \theta = 1, I_{i,t}(\bar{w})) \\ \mu_{i,t}^{(\theta=0)}(w, \bar{w}) &= \Pr(w \mid \theta = 0, I_{i,t}(\bar{w})) \\ p_{i,t}^{(\theta=1)}(\bar{w}) &= \Pr(\theta = 1 \mid I_{i,t}(\bar{w})) \\ p_{i,t}^{(\theta=0)}(\bar{w}) &= \Pr(\theta = 0 \mid I_{i,t}(\bar{w})) \end{aligned}$$

where $I_{i,t}(\bar{w})$ is the information set agent i is at stage t of the game, given that the true type endowment is \bar{w} . Of course, agent i does not know \bar{w} , but only the information set $I_{i,t}$. Suppose now the agents receive new information, namely, $w \in I_{t+1}(\bar{w}) \subseteq I_t(\bar{w})$. Then

$$(A.8) \quad \mu_{i,t+1}^{(\theta=1)}(w, \bar{w}) := \Pr(w \mid \theta = 1, I_{i,t+1}(\bar{w})) = \begin{cases} 0 & \text{if } w \notin I_{i,t+1}(\bar{w}) \\ \frac{\mu_{i,t}^{(\theta=1)}(w, \bar{w})}{\sum_{w' \in I_{t+1}} \mu_{i,t}^{(\theta=1)}(w', \bar{w})} & \text{if } w \in I_{i,t+1}(\bar{w}) \end{cases} .$$

Likewise,

$$(A.9) \quad \mu_{i,t+1}^{(\theta=0)}(w, \bar{w}) := \Pr(w \mid \theta = 0, I_{i,t+1}(\bar{w})) = \begin{cases} 0 & \text{if } w \notin I_{i,t+1}(\bar{w}) \\ \frac{\Pr(w \mid \theta=0, I_t(\bar{w}))}{\sum_{w' \in I_{t+1}} \Pr(w' \mid \theta=0, I_t(\bar{w}))} & \text{if } w \in I_{i,t+1}(\bar{w}) \end{cases} .$$

Based on the new probability distribution over signal endowments, we can get the probability over θ as

$$(A.10) \quad \begin{aligned} p_{i,t+1}^{(\theta=1)}(\bar{w}) &= \Pr(\theta = 1 \mid I_{i,t+1}(\bar{w})) = \Pr(\theta = 1 \mid I_{i,t}(\bar{w}) \cap I_{i,t+1}(\bar{w})) \\ &= \frac{p_{i,t}^{(\theta=1)}(\bar{w}) \sum_{w \in I_{i,t+1}(\bar{w})} \mu_{i,t}^{(\theta=1)}(w, \bar{w})}{p_{i,t}^{(\theta=1)}(\bar{w}) \sum_{w \in I_{i,t+1}} \mu_{i,t}^{(\theta=1)}(w, \bar{w}) + \left(1 - p_{i,t}^{(\theta=1)}(\bar{w})\right) \sum_{w \in I_{i,t+1}} \mu_{i,t}^{(\theta=0)}(w, \bar{w})} . \end{aligned}$$

Therefore, we need to compute the relevant information sets. Let $\mathbf{a}_{i,t}(w)$ be the action that agent i takes at time t if configuration of signals is w . Then we can consider the set

of worlds that have positive probability at time t , given by

$$(A.11) \quad W_{t+1}^{(i)}(\mathbf{a}_t^{(i)}) = \{w \in W : \mathbf{a}_{j,t}(w) = a_{j,t} \text{ for all } j \in N(i)\}$$

and hence,

$$(A.12) \quad I_{i,t+1}(\bar{w}) = W_{t+1}^{(i)}(\mathbf{a}_t^{(i)}(\bar{w}))$$

In deriving the information set, we were able to eliminate the path of actions observed by agent i so far by realizing that the actions taken by an agent are deterministic functions of what they observe. Thus, once we have conditioned on the signal endowment w , the actions observed $a_t^{(i)}(w)$ are completely determined; conditioning on them does not change the probability calculation. However, in environments where we allow for random actions, as opposed to deterministic actions, this fact is no longer true. While it is a perfectly sensible and reasonable model, it complicates things computationally in an astronomical manner. We will return to this problem below.

A.3. Actions. The algorithm described so far gives us how a Bayesian agent i would update her beliefs if she knows $\mathbf{a}_{j,t}(\bar{w})$ for all j and has prior beliefs on signal endowments and states of the world given by

$$(A.13) \quad \left\{ \mu_{i,t}^{(\theta=1)}(w, \bar{w}), \mu_{i,t}^{(\theta=0)}(w, \bar{w}), p_{i,t}^{(\theta=1)}(\bar{w}) \right\}.$$

If agent i is Bayesian, then the decision at $t+1$ is given by (when there are no ties)

$$(A.14) \quad \mathbf{a}_{i,t+1}(\bar{w}) = \begin{cases} 1 & \text{if } p_{i,t+1}^{(\theta=1)}(\bar{w}) > \frac{1}{2} \\ 0 & \text{if } p_{i,t+1}^{(\theta=1)}(\bar{w}) < \frac{1}{2} \\ a \in \{0, 1\} & \text{if } p_{i,t+1}^{(\theta=1)}(\bar{w}) = \frac{1}{2} \end{cases}$$

Note that when $p_{i,t+1}^{(\theta=1)}(\bar{w}) = \frac{1}{2}$, we need to use some tie breaking rule. We will use the “past action” rule. That is, when faced with a tie, an individual will play the action she played in the previous round, $p_{i,t+1}^{(\theta=1)}(\bar{w}) = \frac{1}{2} \implies \mathbf{a}_{i,t+1}(\bar{w}) = \mathbf{a}_{i,t}(\bar{w})$. Of course, one could think of many other tie breaking rules, including random tie breaking rule, where the agent plays each action with the same probability. However, as we will see, this such a model will be computationally intractable in our framework.

Observe that the above framework extends to situations where some agents play ad hoc decision rules via DeGroot learning. Suppose that each agent i may be of some type $\eta \in H := \{\eta_1, \eta_2, \dots, \eta_K\}$. For example, take the type space to be

$$(A.15) \quad H = \{\text{Bayesian, Eigenvector, Naive, Preferential}\}$$

so each agent may be either a Bayesian agent or a DeGroot agent who constructs simple linear indexes from the past actions taken by neighbors. In particular, suppose that agent

i has type $\eta_i = \text{Preferential}$. In world \bar{w} and time $t + 1$, she observes actions $\mathbf{a}_t^{(i)}$. Based on this, she defines the following index:

$$(A.16) \quad \text{Preferential}_{i,t}(\bar{w}) := \sum_{j \in N(i)} \mathbf{a}_{j,t}(\bar{w}) T_{ij}^{\text{Pref}}$$

Therefore, the corresponding action rule is

$$(A.17) \quad \mathbf{a}_{i,t+1}^P(\bar{w}) = \begin{cases} 1 & \text{if } \text{Preferential}_{i,t}(\bar{w}) > \frac{1}{2} \\ 0 & \text{if } \text{Preferential}_{i,t}(\bar{w}) < \frac{1}{2} \\ a \in \{0, 1\} & \text{if } \text{Preferential}_{i,t}(\bar{w}) = \frac{1}{2} \end{cases}.$$

Similarly, we can construct $\mathbf{a}_{i,t+1}^N(\bar{w})$, $\mathbf{a}_{i,t+1}^E(\bar{w})$ using T^{Naive} and T^{Eig} , respectively.

As long as *agents' types are common knowledge*, the algorithm described so far can handle heterogeneity in agents' types without changing the nature of the Bayesian updating.

A.4. Algorithm To Find Action Rules.

Step: 1: Initiate the algorithm with $\mu_{i,1}^{\theta=1}(w, \bar{w})$ and $\mu_1^{\theta=0}(w, \bar{w})$ given by (A.1) and (A.2), action rule $\mathbf{a}_{i,1}(\bar{w})$ as in (A.4) and $p_{i,1}(\bar{w})$ as in (A.3).

Step: 2: At period $t + 1$, start with $\{\mu_{i,t}^{\theta=1}(w, \bar{w}), \mu_{i,t}^{\theta=0}(w, \bar{w}), p_{i,t}(\bar{w}), \mathbf{a}_{i,t}(\bar{w})\}$. Derive $\mu_{i,t+1}^{\theta=1}(w)$ and $\mu_{i,t+1}^{\theta=0}(w)$ using using equations (A.9) and (A.8).

Obtain $p_{i,t+1}(\bar{w})$ from (A.10), and then derive the action that each agent takes, $\mathbf{a}_{i,t+1}(\bar{w})$ depending on the agent's type.

Step: 3: Repeat Step 2 until $t = T$.

APPENDIX B. INCOMPLETE INFORMATION BAYESIAN ALGORITHM

We only include the modification from the above. The only difference now is how to start the algorithm: see that now

$$(B.1) \quad \mu_0^{\theta=1}(w) = \mu_0^{\theta=1}(s, \eta) = \left(p \sum_{i=1}^n s_i (1-p)^{n-\sum_{i=1}^n s_i} \right) \prod_{i=1}^n \pi_i^{\eta_i} (1-\pi_i)^{1-\eta_i}$$

$$(B.2) \quad \mu_0^{\theta=0}(w) = \mu_0^{\theta=0}(s, \eta) = \left((1-p) \sum_{i=1}^n s_i p^{n-\sum_{i=1}^n s_i} \right) \prod_{i=1}^n \pi_i^{\eta_i} (1-\pi_i)^{1-\eta_i}$$

See that the assumption of independent types is immaterial to the description of the algorithm: we could substitute in principle the term $\prod_{i=1}^n \pi_i^{\eta_i} (1-\pi_i)^{1-\eta_i}$ for some function $F(\eta) = \Pr(\eta)$.

After each agent sees their "signal" (s_i, η_i) , we can calculate the derived measures over worlds as

$$\begin{aligned} \mu_{i,1}^{\theta=1}(w, \bar{w}) &= \Pr((s, \eta), (\bar{s}_i, \bar{\eta}_i) \mid (s_i, \eta_i), \theta = 1) \\ &= \frac{\Pr(w \mid \theta = 1) \Pr((s_i, \eta_i) \mid w, \theta = 1)}{\sum_{z \in W} \Pr(z) \Pr((s_i, \eta_i) \mid z, \theta = 1)} = \begin{cases} 0 & \text{if } s_i \neq \bar{s}_i, \text{ or } \eta_i' \neq \bar{\eta}_i \\ \frac{\mu_0^{\theta=1}(w)}{\sum_{z \in W} \mu_0^{\theta=1}(z) \Pr(s_i, \eta_i \mid z, \theta=1)} & \text{otherwise} \end{cases} \end{aligned}$$

and similarly for $\mu_{i,1}^{\theta=0}(w, \bar{w})$. For the other needed objects for the algorithm, we use:

$$\begin{aligned} \mathbf{a}_{i,1}(\bar{w}) &= \begin{cases} 1 & \text{if } \bar{s}_i = 1 \\ 0 & \text{otherwise} \end{cases} \\ p_{i,1}(\bar{w}) &= \begin{cases} p & \text{if } \bar{s}_i = 1 \\ 1 - p & \text{if } \bar{s}_i = 0 \end{cases} \end{aligned}$$

For the rest of the algorithm, the action rule will depend on the type. Let $\mathbf{a}_{i,t}^B(\bar{w})$ denote the probability with which agent i at stage t plays $a = 1$ in world \bar{w} if acts as a Bayesian, and $\mathbf{a}_{i,t}^M(\bar{w})$ be the analogous if the agent act as a M -weighter (and then, does not depend on which state of the world she's in, which only affect the learning that she has done so far). Then, the action profile at time t is

$$\mathbf{a}_{i,t}(\bar{w}) = \begin{cases} \mathbf{a}_{i,t}^B(\bar{w}) & \text{if } \bar{\eta}_i = 1 \\ \mathbf{a}_{i,t}^M(\bar{w}) & \text{if } \bar{\eta}_i = 0 \end{cases}$$

APPENDIX C. FILTERING

Define

$$F_0(\eta) = \Pr(\eta \mid I_0, \hat{\pi})$$

with I_0 the information set of the statistician at $t = 0$. Since we are controlling the experiment, we know that $\{s, \theta\} = I_0$. Call s^* and θ^* the chosen values by the experimentalist. For example, if type endowments are independent of both θ and s and the location on the network

$$(C.1) \quad F_0(\eta) = \prod_{i=1}^n \pi_i^{\eta_i} (1 - \pi_i)^{1 - \eta_i}$$

Now, suppose that we have calculated at to time $t - 1$:

$$F_{t-1}(\eta) = \Pr(\eta \mid I_{t-1}, \hat{\pi})$$

Define

$$(C.2) \quad \mathcal{A}_{i,t-1}^*(\eta) = \Pr(a_{i,t} = 1 \mid I_{t-1}, (s^*, \eta), \hat{\pi}) = \begin{cases} 1 - \varepsilon & \text{if } A_{i,t-1}(\eta, s^*) = 1 \\ \varepsilon & \text{if } A_{i,t-1}(\eta, s^*) = 0 \end{cases}$$

that means, is the probability distribution that the statistician has over actions if she knew the true type endowment. The probability ε is to ensure that as statisticians, we put positive weight on every history, and later we take $\varepsilon \rightarrow 0$. This will not be a problem empirically, since all histories have positive probability empirically.

See that $\mathcal{A}_{i,1}^*(\eta) = \mathbf{1}(s_i = s_i^*)$ for all $\eta \in H$ for any (reasonable) model that for any agent, if they see only their signal, they choose their own signal. Let \mathbf{a}_t^* be the $N \times 1$ action vector observed by the experimenter at time t .

This function allows us to get the conditional probabilities over actions of all agents as:

$$(C.3) \quad \Pr(\mathbf{a}_t \mid I_{t-1}, (s^*, \eta), \hat{\pi}) = \left(\prod_{j:a_j=1} \Pr(a_{j,t} = 1 \mid I_{t-1}, (s^*, \eta), \hat{\pi}) \right) \left(\prod_{j:a_j=0} (1 - \Pr(a_{j,t-1} = 1 \mid I_{t-1}, (s^*, \eta)), \hat{\pi}) \right) \\ \left(\prod_{j:a_j=1} \mathcal{A}_{i,t-1}^*(\eta) \right) \left(\prod_{j:a_j=0} (1 - \mathcal{A}_{i,t-1}^*(\eta)) \right)$$

Using in (C.3) that agents randomize over actions independently. Then, after observing action vector $\mathbf{a}_t^* \in \{0, 1\}^n$, the statistician updates her beliefs over types as

$$(C.4) \quad F_t(\eta) = \Pr(\eta \mid I_{t-1}, \mathbf{a}_t^*) = \frac{\Pr(\eta \mid I_{t-1}, \hat{\pi}) \Pr(\mathbf{a}_t^* \mid I_{t-1}, \eta, \hat{\pi})}{\sum_{\tilde{\eta}} \Pr(\tilde{\eta} \mid I_{t-1}) \Pr(\mathbf{a}_t^* \mid I_{t-1}, \tilde{\eta}, \hat{\pi})} \\ = \frac{F_t(\eta) \left(\prod_{j:a_j^*=1} \mathcal{A}_{i,t-1}^*(\eta) \right) \left(\prod_{j:a_j^*=0} (1 - \mathcal{A}_{i,t-1}^*(\eta)) \right)}{\sum_{\tilde{\eta} \in H} F_t(\tilde{\eta}) \left(\prod_{j:a_j^*=1} \mathcal{A}_{i,t-1}^*(\tilde{\eta}) \right) \left(\prod_{j:a_j^*=0} (1 - \mathcal{A}_{i,t-1}^*(\tilde{\eta})) \right)}$$

To finish up the algorithm, we need to calculate $\mathcal{A}_{i,t}^*(\eta)$, but this comes directly from the algorithm described above.

The algorithm then to get the distribution of η conditional on the whole set of information is:

Step: 1: Initiate algorithm with $F_0(\eta)$ as in (C.1) and an action function \mathcal{A}^* as described above. Moreover, introduce information about s^* (only thing that we actually need)

Step: $t < T$: Taking $\mathcal{A}_{i,t-1}^*(\eta)$ as given, run learning code as in the previous section and calculate $\mathcal{A}_{i,t}^*(\eta)$ as in (C.2)

Step: $t = T$: Once $\mathcal{A}_{i,T-1}^*(\eta)$ is calculated, calculate likelihood over type endowments as

$$(C.5) \quad F_T(\eta) = \frac{F_{T-1}(\eta) \left(\prod_{j:a_{Tj}^*=1} \mathcal{A}_{i,T-1}^*(\eta) \right) \left(\prod_{j:a_{Tj}^*=0} (1 - \mathcal{A}_{i,T-1}^*(\eta)) \right)}{\sum_{\tilde{\eta} \in H} F_{T-1}(\tilde{\eta}) \left(\prod_{j:a_{Tj}^*=1} \mathcal{A}_{i,T-1}^*(\tilde{\eta}) \right) \left(\prod_{j:a_{Tj}^*=0} (1 - \mathcal{A}_{i,T-1}^*(\tilde{\eta})) \right)}$$

Last: Step : Get the probability of being Bayesian of agent i as

$$(C.6) \quad \Pi_{i,t,v}(\pi) \equiv \Pr(\eta_i = 1) = \sum_{\tilde{\eta} \in H: \tilde{\eta}_i = 1} F_T(\tilde{\eta})$$

APPENDIX D. STUCK NODES AND CONCENTRIC SOCIAL QUILTS

D.1. Stuck Nodes. Given an undirected graph $G = (V, E)$ and a subset of nodes $v \subseteq V$ we define $G_v = (v, E_v)$ as the induced subgraph for subset v , where $E_v = \{(ij) \in E : \{i, j\} \subseteq v\}$. Given a subgraph G_v , let $d_i(G_v)$ be the degree of node i in subgraph G_v . Let $a_{it} \in \{0, 1\}$ be the action that node $i \in V$ takes at round $t \in \mathbb{N}$, which we will assume follows the uniform DeGroot action model; i.e. $a_{i,t} = \mathbf{1} \left\{ \frac{1}{d_i+1} \sum_{j \in N_i} a_{j,t-1} > \frac{1}{2} \right\}$. We allow for any tie-breaking rule when $\frac{1}{d_i+1} \sum_{j \in N_i} a_{j,t-1} = \frac{1}{2}$.

LEMMA D.1. *Take a subset of individuals $v \subseteq V$ such that there exists $h \in \mathbb{N}$ with*

$$h \leq d_i(G_v) \leq d_i < 2h + 1 \text{ for all } i \in v$$

If agents behave according to the uniform weighting DeGroot action model, and at some $T \in \mathbb{N}$ we have $a_{i,T} = a \in \{0, 1\}$ for all $i \in v$, then $a_{i,t} = a$ for all $t \geq T$.

Proof. The proof is by induction: without loss of generality, suppose $a_{i,T} = 1$ for all $i \in v$. Of course, for $t = T$ the result is trivially true. Suppose now that $a_{i,t} = 1$ for all $i \in v$ and $t \geq T$, and we need to show that $a_{i,t+1} = 1$ too. Let $I_{i,t+1} = \frac{1}{d_i+1} \sum_{j \in N_i} a_{j,t}$ be the index of uniform weighting. We then now that $I_{i,t} \geq \frac{1}{2}$ for all nodes in v , and it suffices to show that this implies $I_{i,t+1} \geq \frac{1}{2}$. Observe,

$$\begin{aligned} I_{i,t+1} &= \frac{\sum_{j \in N_i} a_{j,t}}{d_i + 1} = \frac{\sum_{j \in v \cap N_i} \underbrace{a_{j,t}}_{=1} + \sum_{j \in N_i - v} a_{j,t}}{d_i + 1} \stackrel{(i)}{\geq} \frac{h + 1 + \sum_{j \in N_i - v} a_{j,t}}{d_i + 1} \\ &\geq \frac{h + 1}{d_i + 1} \stackrel{(ii)}{>} \frac{1}{2}. \end{aligned}$$

We have used in (i) that $d_i(G_v) \geq h$ and $a_{j,t} = 1$ for all $j \in v$. Inequality (ii) comes from the fact that

$$\frac{h + 1}{d_i + 1} > \frac{1}{2} \iff d_i < 2h + 1.$$

Therefore, we have that $I_{i,t+1} > \frac{1}{2}$ for any $i \in v$, implying then that $a_{i,t+1} = 1$, as we wanted to show. \square

This lemma says that whenever we find a subset of nodes v such that each node has more connections to nodes in v than it has outside v , then whenever they reach consensus, they would remain there forever. We present an useful corollary of Lemma D.1, which we will use when studying the family CSQ_r .

COROLLARY D.1 (Regular Subgraphs). *Take a family of nodes $v \in V$ such that there exists $k \in \mathbb{N}$ such that*

- (1) G_v is a k -regular graph
- (2) $d_i < 2k + 1$ for all $i \in v$.

Then, if agents behave according to the uniform weighting DeGroot action model, and at some $T \in \mathbb{N}$ we have $a_{i,T} = a \in \{0, 1\}$ for all $i \in v$, then $a_{i,t} = a$ for all $t \geq T$.

Proof. Simply take $h = k$ and apply Lemma D.1. □

See that any triangle in CSQ_r is a 2-regular subgraph, and that each node in it has $d_i = 4 < 2 * 2 + 1 = 5$, so we apply Corollary D.1 with $k = 2$. So, whenever a triangle achieves consensus, it remains there forever.

D.2. Concentric Social Quilts: Preliminaries. Let's now focus on Concentric Social Quilts. Define $S_r = \{i \in V_r : i \text{ gets stuck}\}$ and let $N_r = \#(V_r)$, the number of nodes in CSQ_r . Our object of interest is the random variable

$$\mathcal{F}_r(r) = \text{Fraction of nodes in } CSQ_r \text{ that gets stuck} \equiv \frac{\#S_r}{N_r}$$

which is a random variable. Our objective is to get an asymptotic bound on \mathcal{F}_r . Since we do not yet know whether \mathcal{F}_r has a limit for almost every realization, we define $\underline{\mathcal{F}}$ and $\overline{\mathcal{F}}$ as

$$(D.1) \quad \underline{\mathcal{F}} = \liminf_{r \rightarrow \infty} \mathcal{F}_r \text{ and } \overline{\mathcal{F}} = \limsup_{r \rightarrow \infty} \mathcal{F}_r$$

which is well defined for all realizations of the sequence $\mathcal{F}(r)$ and so it is a well defined random variable. Namely, we want to get the tightest asymptotic lower and upper bounds for the fraction of stuck nodes. Our objective is to get a number $F \in [0, 1]$ such that $\underline{\mathcal{F}} \geq \underline{F}$ and $\overline{\mathcal{F}} \leq \overline{F}$ almost surely; i.e. $\mathbb{P} \left\{ \underline{F} \leq \underline{\mathcal{F}} \leq \overline{\mathcal{F}} \leq \overline{F} \right\} = 1$.

Any CSQ_r has an initial parent triangle $P = (i_0, i_1, i_2)$. For any node $i \in V_r$ we define the distance to parent $d(i, P) = \min \{d(i, i_0), d(i, i_1), d(i, i_2)\}$ where $d(i, j)$ is the minimum number of links we have to go through to connect node i with node j . Likewise, given a triangle $T = (i_s, i_k, i_j)$, we define the distance between triangle T and the parent triangle P as

$$d(T, P) = \max_{i \in T} \{d(i, P)\}$$

DEFINITION D.1 (Ring). *Given a graph $CSQ_r = (V_r, E_r)$ and $s \in \{1, \dots, r\}$ we define R_s , the level s ring of CSQ_r as the subgraph $R_s = (V_r^s, E_r^s)$ of nodes that lie in triangles with distance to parent $d(T, P) = s - 1$.*

Intuitively, a ring is just the collections of triangles that lie in the s -th level of the concentric social quilt as seen in Figure 5. Note that for all s , R_s is a graph that consists of disconnected triangles.

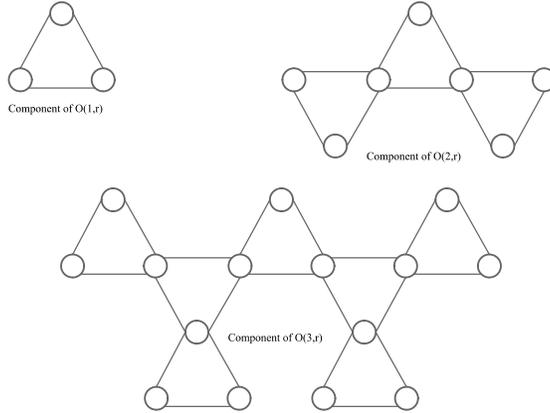


FIGURE 5. Outer Rings.

Define

$$OR_r(k) := \bigcup_{s=0}^{s=k} R_{r-s}$$

as the subgraph formed by the outer rings from $r - k$ to r . This subgraph is also disconnected, with a lot of components, which now are no longer triangles, but rather “trees of triangles” as pictured in Figure 5.

Let $C \subset OR_r(k)$ be a component (a subgraph as shown in Figure 5). The last level of nodes in every component correspond to terminal nodes. The most important property of these components is that **the only connection between each component C and the rest of the graph is the parent node of the component C** , denoted by i_C . This will be the key property of these components, which we will try to explore.

Define

$$\Psi_r(k) = \frac{\#\{OR_r(k) \cap S_r\}}{\#\{OR_r(k)\}}$$

to be the fraction of stuck nodes in $OR_r(k)$

$$\underline{\Psi}(k) := \liminf_{r \rightarrow \infty} \frac{\#\{OR_r(k) \cap S_r\}}{\#\{OR_r(k)\}}$$

and

$$\overline{\Psi}(k) := \limsup_{r \rightarrow \infty} \frac{\#\{OR_r(k) \cap S_r\}}{\#\{OR_r(k)\}}$$

which is also a well defined random variable. These are the tightest asymptotic lower and upper bounds on the fraction of nodes stuck in the last k rings. That is, a lower bound on the fraction of nodes in $OR_r(k)$ that get stuck.

LEMMA D.2. For all $k \in \mathbb{N}$, $\lim_{r \rightarrow \infty} \frac{\#\{OR_r(k)\}}{N_r} = \frac{2^{k+1}-1}{2^{k+1}}$.

Proof. Let L_r = number of terminal triangles in ring r . Of course, we have that $L_r = T_{r-1}$. Because of how CSQ_r grows, we have the following recursion for L_r :

$$L_{r+1} = 2L_r \text{ and } L_2 = 3.$$

It can be easily shown that

$$(D.2) \quad L_r = 3 * 2^{r-2}.$$

We also need to calculate $N_r = \#(V_r)$. Again, because of how CSQ_r is generated, we have the following recursion for N_r :

$$N_{r+1} - N_r = 2L_{r+1}$$

and it can be also easily shown that

$$(D.3) \quad N_r = 3(2^r - 1)$$

Finally, let n_k be the number of nodes in a component $C \subset OR_r(k)$. It is also easy to show that

$$n_k = 2^{k+1} - 1$$

Now, we can state the result. Observe that

$$\begin{aligned} \frac{\#\{O_r(k)\}}{N_r} &= \frac{\overbrace{n_k}^{\text{nodes per component}} \times \overbrace{L_{r-k+1}}^{\text{number of components}}}{N_r} = (2^{k+1} - 1) \frac{3 \times 2^{r-k+1-2}}{3(2^r - 1)} \\ &= \frac{2^{k+1} - 1}{2^{k+1}} \left(\frac{2^r}{2^r - 1} \right) \xrightarrow{r \rightarrow \infty} \frac{2^{k+1} - 1}{2^{k+1}} \end{aligned}$$

as we wanted to show. □

The following proposition is the key to understand how to get bounds on $\underline{\mathcal{F}}$ and $\overline{\mathcal{F}}$ by getting bounds on $\underline{\Psi}(k)$ and $\overline{\Psi}(k)$

PROPOSITION D.1. *Suppose there exist functions $\overline{\psi}, \underline{\psi} : \mathbb{N} \rightarrow [0, 1]$ such that for all k we have*

$$\underline{\psi}(k) \leq \underline{\Psi}(k) \leq \overline{\Psi}(k) \leq \overline{\psi}(k) \text{ almost surely.}$$

Then, for all $k \in \mathbb{N}$ almost surely,

$$(D.4) \quad \underline{\mathcal{F}} \geq \frac{2^{k+1} - 1}{2^{k+1}} \underline{\psi}(k)$$

and

$$(D.5) \quad \overline{\mathcal{F}} \leq 1 - \left(\frac{2^{k+1} - 1}{2^{k+1}} \right) [1 - \overline{\psi}(k)].$$

Proof. Lets focus only on inequality D.4, since D.5 follows the same reasoning. See that

$$\begin{aligned} \mathcal{F}(r) &= \frac{\#\{O_r(k)\}}{N_r} \left(\frac{\#\{O_r(k) \cap S_r\}}{\#\{O_r(k)\}} \right) + \frac{\#\{S_r - O_r(k)\}}{N_r} \\ &\geq \frac{\#\{O_r(k)\}}{N_r} \left(\frac{\#\{O_r(k) \cap S_r\}}{\#\{O_r(k)\}} \right) \end{aligned}$$

so, for all realizations,

$$\begin{aligned} \underline{\mathcal{F}} &= \liminf_{r \rightarrow \infty} \mathcal{F}(r) \geq \left(\lim_{r \rightarrow \infty} \frac{\#\{O_r(k)\}}{N_r} \right) \left(\liminf_{r \rightarrow \infty} \left(\frac{\#\{O_r(k) \cap S_r\}}{\#\{O_r(k)\}} \right) \right) \\ &= \left(\lim_{r \rightarrow \infty} \frac{\#\{O_r(k)\}}{N_r} \right) \underline{\Psi}(k) = \frac{2^{k+1} - 1}{2^{k+1}} \underline{\Psi}(k). \end{aligned}$$

This, together with the fact that $\underline{\Psi}(k) \geq \underline{\psi}(k)$ almost surely, finishes the proof. \square

Note that this proposition is true for any learning model (Bayesian or DeGroot). The learning model plays a role when calculating the bounds $\underline{\psi}$ and $\bar{\psi}$. See that condition D.4 and D.5 are bounds on $\underline{\mathcal{F}}$ and $\bar{\mathcal{F}}$, which do not depend on k : therefore, these are bounds for all k : the higher k , the tighter the bound we get.

D.3. Bounding stuck nodes in the Uniform Weighting model. To normalize, we will assume that the true state of nature is $\theta = 1$, which implies that as $r \rightarrow \infty$ the fraction of nodes with true signals is $p > \frac{1}{2}$. The idea is pretty simple: take a component $C = (V_C, E_C) \subset OR(k, r)$. As we mentioned before, the only connection between C and the rest of the graph is through the parent node i_C (as seen in Figure 5). Let $W_C = \{0, 1\}^{n_k}$ be the set of signal endowments for nodes in C . We will try to find a lower bound $\underline{\psi}_k(w)$ for each signal endowment realization in C such that, when signal endowment is w , the fraction of stuck nodes in C is larger than $\underline{\psi}_k(w)$ fraction of stuck nodes in C if endowment is $w \geq \underline{\psi}_k(w)$. If we can find such $\underline{\psi}_k(w)$, then we can use a law of large numbers to argue that

$$\underline{\Psi}(k) \geq \underline{\psi}(k) \equiv \mathbb{E}_{w \in W_C} \left\{ \underline{\psi}_k(w) \right\} \text{ almost surely}$$

because the realizations of w in each component C is independent of each other. Likewise, if we can find a function $\bar{\psi}_k(w)$ to bound from above the fraction of stuck nodes, and then

$$\bar{\Psi}(k) \leq \bar{\psi}(k) \equiv \mathbb{E}_{w \in W_C} \left\{ \bar{\psi}_k(w) \right\} \text{ almost surely}$$

Imagine first that the signal endowment of the upper triangle in C is $(0, 0, 0)$. Then, using Lemma D.1 we know that the upper triangle of C will get stuck from period $t = 1$ on, and we can get the expected value of stuck nodes in C from there on. See that the fraction of nodes that get stuck in this component is only a function of the realization of $w \in W_C$, which is independent of the realization of the signal endowment of other components on $OR_r(k)$

When the signal endowment of the upper triangle in C is different from $(0, 0, 0)$, we make use of the other property we knew from C : that the only connection to the rest of

the graph is through i_C , the uppermost node in C . Therefore, **a way of getting a lower bound on the number of nodes that get it wrong, is assuming that from round $t = 2$ on, node i_C knows the truth, and plays $a_{i_C,t} = 1$ for all $t \geq 2$.** Intuitively, we are making the graph to have the biggest effect possible in convincing nodes in C that actually, $\theta = 1$, which can only do by making $a_{i_C,t} = 1$ for all rounds other than $t \geq 2$. Once we have that, we can simulate the learning model on C , and calculate $\underline{\psi}_k(w)$ and $\overline{\psi}_k(w)$ in this way.

There are two ways of calculating $\mathbb{E}_{w \in W_C} \{ \underline{\psi}_k(w) \}$:

- (1) **Doing it explicitly:** This can be done for $k = 2$ and $k = 3$, because $\# \{W_C\} = 128$. The bound when $k = 3$ is the one we present in this paper.
- (2) **Monte-Carlo:** Of course, as k goes bigger, it is computationally unfeasible to calculate the expected value of $\underline{\psi}_k(w)$ explicitly, since

$$\# \{W_C\} = 2^{2^{k+1}-1} = O(\exp(\exp(k)))$$

which grows super-exponentially. However, we can simulate random draws of $w \in W_C$ and get an estimate for $\mathbb{E}_{w \in W_C} \{ \underline{\psi}_k(w) \}$ using law of large numbers.

The above method will also work for different learning models on the CSQ_r family.

APPENDIX E. PROOFS

Proof of Proposition 2.1. The first part follows from Golub and Jackson (2010), since every node has degree 2 or 4, Corollary 1 applies. Namely, $\max_{1 \leq i \leq n} \frac{d_i}{\sum d_i} \rightarrow 0$ along our sequence and therefore the social learning process is wise. The second part follows from Lemma D.2. \square

Proof of Proposition 4.2. Let $f(n, T)$ be the amount of calculations that should be done for a network of size n and played for T rounds.

$$\begin{aligned} f(n, T) &= \sum_{t=1}^T n 2^{nt+1} (1 + 2^{nt}) = 2n \left(\sum_{t=1}^T 2^{nt} + \sum_{t=1}^T 4^{nt} \right) \\ &= 2n \left[\frac{2^{n(T+1)} - 2^n}{2^n - 1} + \frac{4^{n(T+1)} - 4^n}{4^n - 1} \right] = \Theta(n 4^{nT}) \end{aligned}$$

Meanwhile, the amount of calculations for the original model (with no trembles) is simply given by

$$g(n, t) = \sum_{t=1}^{t=T} n(2^{2n} + 2^{n+1}) = nT(4^n + 2^{n+1}) = \Theta(nT4^n)$$

Thus, the complexity ratio between the model with trembles and the model with no trembles is

$$\frac{2n \left[\frac{2^{n(T+1)} - 2^n}{2^n - 1} + \frac{4^{n(T+1)} - 4^n}{4^n - 1} \right]}{nT(4^n + 2^{n+1})} = \Theta \left(\frac{1}{T} 4^{n(T-1)} \right)$$

which completes the proof. \square