

Working Paper No. 571

Testing Models of Social Learning on Networks: Evidence from a Lab Experiment in the Field

Arun G. Chandrasekhar | Horacio Larreguy | Juan Pablo Xandri

May 2016



Stanford

Center for International
Development

John A. and Cynthia Fry Gunn Building | 366 Galvez Street, Stanford, CA 94305-6015

TESTING MODELS OF SOCIAL LEARNING ON NETWORKS: EVIDENCE FROM A LAB EXPERIMENT IN THE FIELD

ARUN G. CHANDRASEKHAR[‡], HORACIO LARREGUY[§], AND JUAN PABLO XANDRI^{*}

ABSTRACT. Agents often use noisy signals from their neighbors to update their beliefs about a state of the world. The effectiveness of social learning relies on the details of how agents aggregate information from others. There are two prominent models of information aggregation in networks: (1) Bayesian learning, where agents use Bayes' rule to assess the state of the world and (2) DeGroot learning, where agents instead consider a weighted average of their neighbors' previous period opinions or actions. Agents who engage in DeGroot learning often double-count information and may not converge in the long run. We conduct a lab experiment in the field with 665 subjects across 19 villages in Karnataka, India, designed to structurally test which model best describes social learning. Seven subjects were placed into networks with common knowledge of the network structure, designed to maximize our experimental power to distinguish between these two models of learning. Subjects attempted to learn the underlying (binary) state of the world, having received independent identically distributed signals in the first period. Thereafter, in each period, subjects made guesses about the state of the world, and these guesses were transmitted to their neighbors at the beginning of the following round. We structurally estimate a model of Bayesian learning, relaxing common knowledge of Bayesian rationality by allowing agents to have incomplete information about other players' types (Bayesian or DeGroot updating types). Our estimates show that, despite the flexibility in modeling learning in these networks, agents are robustly best described by DeGroot-learning models wherein they take a simple majority of previous guesses in their neighborhood.

KEYWORDS: networks, social learning, Bayesian learning, DeGroot learning, experiments

JEL CLASSIFICATION CODES: D83, C92, C91, C93

1. INTRODUCTION

The way in which individuals aggregate information is a critical feature of many economic environments. Information and opinions about new technologies, job opportunities,

Date: First Version: August 2011, This Version: May 2016.

We are grateful to Daron Acemoglu, Abhijit Banerjee, Esther Duflo, Ben Golub, Matthew O. Jackson, Markus Möbius, and Adam Szeidl for extremely helpful discussions. Essential feedback was provided by Kalyan Chatterjee, Juan Dubra, Rema Hanna, Ben Olken, Evan Sadler, Rob Townsend, Xiao Yu Wang, Luis Zermeno and participants at numerous seminars and conferences. We also thank Mounu Prem for excellent research assistance. 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. JPAL and CMF at IFMR provided valuable logistical assistance. Gowri Nagraj and our field team were vital to the project's progress. All errors are our own.

[‡]Stanford University, Department of Economics and NBER.

[§]Harvard University, Department of Government.

^{*}Princeton University, Department of Economics.

products, and political candidates, among other things, are largely transmitted through social networks. However, the information individuals receive from others about the state of the world often contain noise. Moreover, individuals might differ in how they imagine that others interpret and process information they retransmit. A critical aspect of social learning, therefore, concerns how agents handle and aggregate noisy information in order to resolve uncertainty and estimate the state of the world, on which they base their subsequent decisions.

Consider the case where an individual A hears which of two choices each of her two friends B and C intend to make in the coming period. For example, this could be which of two fertilizers her friends intend to use or which of two candidates her friends are likely to vote for. Both B and C have friends of their own, including some shared friends, and have discussed their choices with them. Knowing all of this, when A updates her opinion as to which choice is best, how should she incorporate B 's and C 's views? She may naively treat them as independent clues about the optimal choice, or perhaps she is more sophisticated and worries that there is a common component to their opinions because they have friends in common. Or A may even worry about how B and C themselves make inferences about the world, and thus, arrive at their opinions. In this paper, we conduct a lab experiment in the field in order to carefully study the nature of social learning on a network, focusing on the two canonical models in the learning on networks literature: Bayesian and DeGroot learning (Jackson, 2008; Acemoglu et al., 2011; Choi et al., 2015). Relative to existing literature, our main contribution is that we allow for incomplete information by agents about other agents' types, and thus we do not simply restrict to the race between complete information Bayesian and DeGroot learning, allowing then a fair comparison between the two learning models.

There are two broad paradigms of modeling social learning through networks. The first is Bayesian learning, wherein individuals process information using Bayes' rule (see, e.g., Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), Gale and Kariv (2003), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Mossel and Tamuz (2010), Lobel and Sadler (forthcoming)). An environment often studied considers a situation where each individual initially receives a signal about the state of the world and then subsequently observes her neighbors' guesses, before revising her own belief and offering an updated guess in each period. A core result in this setting is that individuals eventually converge in their beliefs as to which state of the world they are in; this guess is correct in large networks when some regularity conditions hold (Gale and Kariv, 2003; Mossel and Tamuz, 2010). While Bayesian agents in large networks are able to come to the right opinion in this environment, these results rely on sophisticated agents who are able to implicitly discern between repeated and new information that they receive through their neighbors over time.

This cognitive load that the Bayesian-learning model imposes over agents when these learn in social networks led to a second paradigm: DeGroot learning (DeGroot, 1974). In

these models, agents are myopic and, after seeing the behavior or beliefs of their network neighbors, they take a weighted average of these to construct their belief going into the subsequent period. Ellison and Fudenberg (1993; 1995), DeMarzo, Vayanos, and Zwiebel (2003), Eyster and Rabin (2010; 2014), Golub and Jackson (2010), and Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012) among others, have studied related models of this form. When agents communicate beliefs with their neighbors and engage in DeGroot learning they converge to the truth in large networks (Golub and Jackson, 2010). However, this convergence can be inefficient and agents located in networks that exhibit significant homophily – wherein agents tend to be connected more with similar others – converge slowly to the truth (Golub and Jackson, 2012).¹

The contrast between DeGroot and Bayesian learning grows when one moves away from rich communication structures – where agents exchange beliefs with their neighbors – in DeGroot learning models.² In many environments of interest, communication can be far more coarse: individuals might reveal which candidate they voted for but are probably less willing to discuss the reasons (e.g., underlying beliefs about candidate competence) behind such vote choice, farmers might observe which fertilizer their neighboring farmers adopted but struggle to elicit from them what fertilizer quality considerations led them to such fertilizer choice, individuals might observe their neighbors joining a microfinance organization but their neighbors might struggle to communicate the motives behind such take up. Moreover, and perhaps especially in a context such as ours where education levels are low, assuming that fine probabilities and complex beliefs are communicated by agents may be unrealistic. It is unsurprising then that much of the Bayesian learning on networks literature focuses on coarse communication, where agents observe their neighbors’ actions as opposed to their beliefs. The differences between the Bayesian and DeGroot paradigms are particularly pronounced in these environments. While Bayesian learning typically generates convergence to the truth in large societies under certain regularity conditions, DeGroot learning may generate misinformation traps, wherein pockets of agents hang on to an incorrect opinion for infinitely many periods. Thus, understanding which learning paradigm better predicts group behavior is important both for understanding whether information transmission is efficient and for thinking through policy designs that rely on information dissemination.³

¹Eyster and Rabin (2014) also restrict to a continuous action space that fully reveals players’ posterior beliefs but consider larger class of behaviors than DeGroot-style behavior. They show that players that use social learning rules that neglect redundancy can converge to wrong long-run beliefs with positive probability. They conclude that abundant imitation might be socially harmful.

²For a recent review article, see Sadler (2014).

³Consider the case where the state of the world is either 0 or 1. In this binary environment, individuals engaging in DeGroot learning often double-count information, fail to not take into account dependencies, and may not reach consensus in the long run even for extremely large graphs (as we show below). Meanwhile, in such an environment, Bayesian learning mechanisms, under some regularity conditions to avoid ties in posterior probabilities, generically generate consensus in finite graphs and, moreover, in very large graphs

In this paper, we study whether Bayesian or DeGroot learning does a better job of describing learning on networks. We conducted a unique lab experiment in the field in 19 villages in rural Karnataka, India, with 665 subjects. We ran our experiments directly in the villages so that we could study a population of interest to development economists and policy makers – namely, those who could be potentially targeted by policy that depends on social learning (e.g., the introduction of a new technology, health practices, microfinance). We designed simple networks that maximized statistical power to distinguish between the two canonical models in the learning on networks literature: Bayesian and DeGroot learning (Choi et al., 2015). We then conducted a lab experiment in the field, using these networks to structurally study which canonical learning model better fits learning in social networks. The network selection was at the conscious potential cost of not being able to separate between other models of social learning studied in the literature and discussed later, but reflected the need to select small but sufficiently large networks to have power to distinguish between the two models of learning the literature deems as most relevant (Choi et al., 2015).

We created networks of seven individuals, giving each individual a map of the entire graph so that the full complete informational structure was comprehended. The underlying state of the world was either 1 or 0 with equal probability. At $t = 0$ each individual received an independent identically distributed (i.i.d.) signal about the underlying state of the world and was 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, which was communicated to each individual’s network neighbors. Thereafter, each individual made guesses about the underlying state of the world, and these guesses were transmitted to her neighbors at the beginning of the following round. Using this information, each individual made a new 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. Individuals were incentivized to guess correctly, as we discuss in further detail below.⁴

In comparing Bayesian and DeGroot learning, the following subtle issue must be considered. The standard Bayesian-learning model typically encodes two important but distinct features: (1) agents are Bayesian, and therefore apply Bayes’ rule to make inferences about neighbors’ information sets; and (2) agents have common knowledge of the ways in which their neighbors map information sets to actions. The first condition is that agents are Bayesian, and the second condition is that they have correct beliefs about the other agents’ types, i.e., whether they are Bayesian or DeGroot learners. The most extreme version of this is a model is one of complete information where everyone is Bayesian and

the populations’ limit belief coincides with the true state of the world (Gale and Kariv, 2003; Mossel and Tamuz, 2010). That is, if the world was truly 0, all individuals would eventually come to believe this.

⁴While 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 assumes away experimentation (see, e.g., Choi et al. (2012), Mossel et al. (2015)).

this is common knowledge. Our approach is unique in that relaxes both of these features when studying the data, and thus allow for a better comparison of Bayesian and DeGroot learning models. A model of Bayesian agents with common knowledge of Bayes-rationality is, in a sense, a strawman; Bayesian agents who are uncertain of their neighbors' types may behave differently than Bayesian agents who have common knowledge of Bayes-rationality of their network. We structurally estimate a Bayesian learning model with incomplete information, where agents need not know ex-ante how others are learning, and relax common knowledge, as well. The key parameter in our incomplete information Bayesian learning model is π , which is the share of agents in the population who are Bayesian, with the assumption that a $1 - \pi$ share of agents are DeGroot learners in their thinking. We are then interested in what $\pi \in [0, 1]$ best describes the data from the experiment.

To assess how a particular model fits the experimental data, we look at the data from two different perspectives: the network level and the individual level. The network-level analysis considers the entire network and sequence of actions by agents as a single observation. That is, we consider the path of actions predicted by theory under a model of social learning for each individual in each period. This is an important frame for analysis because, while all models are wrong, what we are interested in, as researchers studying social networks, is what the learning process better explains the aggregate learning patterns in our experimental data. In the individual-level analysis, instead of focusing on the the social network itself as the observational unit, we consider the action of an individual given a history of actions. We then study which learning model better predicts individual-level decisions, given a history of observed signal and neighbors' actions.

Our core results are as follows. Considering first the network-level analysis, we find that the incomplete information model that better explains the data is one where the share of Bayesian agents in the population (π) is 0. Thus, the data is best fit by a model where all agents are DeGroot. Specifically, this model explains between 86% and 88% of the the actions taken in the data. This is not to say that social learning is not well approximated by Bayesian learning; a complete information Bayesian learning model explains 82% of the experimental data. However, this fit largely originates from the fact that the predictions of the Bayesian and DeGroot models often agree. Moreover, Bayesian block bootstrap, which accounts for dependence in the data due to individuals playing the same learning game, indicates that the difference in fit between the DeGroot and Bayesian learning models is statistically significant. Second, at the individual level, again we find that incomplete information Bayesian models fit the data better when the proportion of Bayesian agents (π) is 0 and there is common knowledge about this. In fact, the DeGroot-learning models explain between 89% and 94% of the actions taken by agents given a history of actions. Meanwhile, the complete information Bayesian model only explains 74% of those actions. Again, the difference in fit between the DeGroot and Bayesian learning models is statistically significant.

We also establish several supplementary results which may also be of independent interest. First, we explore the DeGroot action model (also called the voter model) and demonstrate that even in simple networks where Bayesian learners who observe their neighbors' actions and DeGroot agents who fully communicate their beliefs would converge to the truth, a non-vanishing and non-trivial share of agents could converge to the wrong beliefs. This is important because while Bayesian learning where actions are observed (or coarse information is transmitted) is typically studied in the literature, DeGroot learning models typically study settings where very fine information is communicated. In our view, the fair analog of Bayesian learning where actions are observed is DeGroot learning where actions are observed. It turns out that in this setting, DeGroot-induced errors in learning can destroy efficient social learning in networks that would otherwise be asymptotically efficient under either Bayesian (action) learning or DeGroot communication learning.

Second, we develop a simple, non-parametric test that directly distinguishes between DeGroot-like and Bayesian behavior and show that a considerable share of time the data is consistent with an all-agent DeGroot model.

Third, we develop a simple algorithm to compute incomplete information Bayesian learning on networks that is the best network-independent algorithm. Specifically, it is computationally tight in the sense that, asymptotically, there can be no faster algorithm that is independent of network structure. This is a challenge that, to our knowledge, previous work had not overcome, precluding a structural analysis of learning on networks like ours. Namely, the algorithm is $O(T)$, where T is the number of rounds played.⁵

Fourth, an approach taken in related work looks at learning models with trembles or, relatedly, quantal response equilibrium (QRE) a la [Choi et al. \(2012\)](#). We demonstrate that networks small enough to avoid computational constraints are not large enough to separate between DeGroot and Bayesian learning with trembles. Meanwhile those that are large enough to separate between those models become computationally infeasible to study using trembles or QRE.

Fifth, we discuss why selecting between DeGroot and Bayesian learning models must be done through structural estimation and in a lab setting. We show that natural examples of reduced form analyses, where the intuitions of Bayesian and DeGroot learning are used to test for predictions in regression analysis of social learning data, may be problematic. Namely, the data generated by Bayesian and DeGroot learning models do not necessarily conform to the intuition motivating the regressions. Thus, there much to be gained from a structural analysis. The computational constraints for structural estimation of learning models in large networks, however, suggests the importance of small networks in lab settings to separate between models of social learning. Moreover, lab settings allow controlling priors of agents in the network and the signal quality, as well as restricting the

⁵An 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.

communication among individuals and the social-network structure. Since structural estimation is often sensitive to misspecification, it is difficult to cleanly identify which model best describes the data in a non-laboratory context.

In terms of related work, [Gale and Kariv \(2003\)](#) study the Bayesian learning environment that we build upon. 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 her neighbors in the network. Via the martingale convergence theorem, they point out that connected networks with Bayesian agents yield uniform actions in finite time with probability one. [Choi et al. \(2005, 2012\)](#) 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. Closer to the goal of our paper, [Choi \(2012\)](#) estimates a cognitive-hierarchy model of social learning on networks using the experimental data from [Choi et al. \(2005\)](#) and finds that Bayesian behavior better fits the experimental data. However, note that, while an important step towards our goal, none of these papers study whether Bayesian or DeGroot learning models better explain the experimental learning behavior. The networks 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 DeGroot-type learning models.⁶

The works most closely related to ours are [Möbius et al. \(2015\)](#), conducted in 2004, who study how information decays as it spreads through networks, [Corazzini et al. \(2012\)](#), [Mueller-Frank and Neri \(2013\)](#) and [Mengel and Grimm \(2014\)](#), who conduct lab experiments similar to ours. [Möbius et al. \(2015\)](#) test between DeGroot models and an alternative model that they develop in which individuals “tag” information by describing its origin (called a “streams” model). Note that they do not study Bayesian learning alternatives per se, because that would be computationally prohibitive in their setting. Their experiment uses network data from Harvard undergraduates in conjunction with a field experiment and finds evidence in favor of the “streams” model model in which individuals “tag” information. In our experiment, we shut down the possibility that individuals “tag” information. This allows us to compare the Bayesian model to DeGroot alternatives since, as described above, looking at Bayesian learning even in a complete information context (let alone incomplete information) is impossible in such a field setting. As noted by [Möbius et al. \(2015\)](#), the conclusions of our work and theirs suggest that in complicated environments where tagging can be difficult, agents may exhibit more DeGroot-like learning behavior.

⁶The literature on social learning experiments begins with [Anderson and Holt \(1997\)](#), [Hung and Plott \(2001\)](#), and [Kubler and Weizsacker \(2004\)](#). Explicit network structures are considered in a series of papers by [Gale and Kariv \(2003\)](#), [Choi et al. \(2005, 2012\)](#), and [Celen et al. \(2010\)](#).

Prior to our work, [Corazzini et al. \(2012\)](#), study one simple network where Bayesian and DeGroot learning models differ in their predictions. They show reduced form experimental evidence consistent with DeGroot learning. The authors further propose a generalized updating rule that nests DeGroot as a special case, and they show that it fits the data well. However, as [\(Choi et al., 2015\)](#) highlight, [Corazzini et al. \(2012\)](#) only investigate one simple network in which there is a difference between the outcomes of the Bayesian and DeGroot learning, making it difficult to generalize their findings.⁷ Moreover, in section 5 we show that selecting between Bayesian and DeGroot learning models through reduced form approaches may be problematic. Subsequent to our work, and independently of us, both [Mueller-Frank and Neri \(2013\)](#) and [Mengel and Grimm \(2014\)](#) conducted lab experiments to look at Bayesian versus DeGroot learning. Like [Corazzini et al. \(2012\)](#), these studies offer some modified models of learning. [Mueller-Frank and Neri \(2013\)](#), for instance, develop a general class of models called Quasi-Bayesian updating and show that long run outcomes in their experiments are consistent with this model. Crucial differences from our work include the fact that [Mueller-Frank and Neri \(2013\)](#) and [Mengel and Grimm \(2014\)](#) don't reveal the information about the entire network to their subjects, making the inference problem more complicated. More importantly, none of these studies allow for their agents to have incomplete information, which can greatly change the benchmark Bayesian's behavior. We are particularly interested relaxing the complete information Bayesian model with common knowledge of Bayesian rationality since, in a sense, it serves as a weak straw man. In theory, with a mix of Bayesian and naive agents, a Bayesian agent understanding the heterogeneity in the population would behave in a manner such that the data would look very different from a world in which all agents are Bayesian. Taken together, there is a vibrant literature interested in whether learning in networks can be thought of as DeGroot or Bayesian.

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. Setup.

2.1.1. *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 $A_{ij} = \mathbf{1}\{ij \in E\}$.

⁷[Brandts et al. \(2015\)](#) extends the work by [Corazzini et al. \(2012\)](#) showing further experimental evidence consistent with DeGroot learning considering other networks. However, the simple networks they consider do not allow for statistical power to distinguish between Bayesian and DeGroot alternatives.

Individuals in the network are attempting to learn about the underlying state of the world, $\theta \in \{0, 1\}$. 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 . In every subsequent period, the agent takes action $a_{i,t} \in \{0, 1\}$, which is her guess of the underlying state of the world. Figure 1 provides a graphical illustration of the timeline.

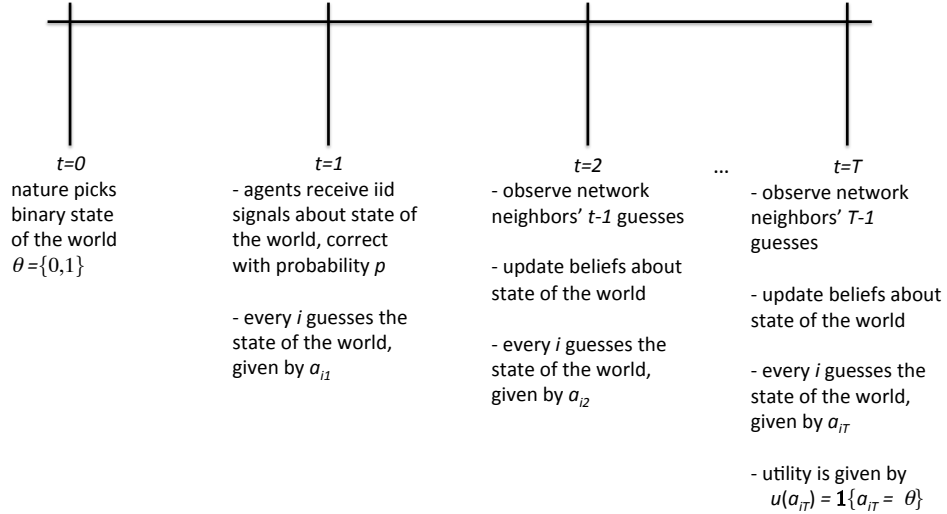


FIGURE 1. Timeline

Through their interactions, agents try to learn about the initial signal configuration $s = (s_1, \dots, s_n)$, with $s_i \in \{0, 1\}$.⁸ Note that the set of all such configurations, $S := \{0, 1\}^n$, has 2^n elements. Sometimes we refer to signal endowments as *states of the world* $\omega = s$, in the language of partitional information models as in [Aumann \(1976\)](#) and [Geanakoplos and Polemarchakis \(1982\)](#), which is useful when studying Bayesian learning models with different state spaces $\omega \in \Omega$. Finally, we 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 .

2.1.2. Bayesian Learning. In our analysis, we consider a model of Bayesian learning with incomplete information. Individuals have common priors over the relevant state spaces (described below) and update according to Bayes' rule in each period.

Each agent is drawn i.i.d. from an infinite population which has a π share Bayesian agents and a $1 - \pi$ share DeGroot agents. This fact is common knowledge – obviously a

⁸In the complete information model, the signal configuration s completely determines how all other agents play, and is therefore a sufficient statistic for one's belief about θ . Hence, the signal configurations $s \in S$ are the only relevant states that Bayesian agents need to learn about. In the incomplete information model (where other players in the network may not be Bayesian) the signal configuration s no longer determines how neighbors play; agents need to also learn how their neighbors process their information. Here the sufficient statistics for one's belief about θ is s as well as the configuration of types – whether each individual is Bayesian or DeGroot. See [Appendix A.2](#) for a formal discussion.

feature relevant only for the Bayesian agents – as is the structure of the entire network. Since there is incomplete information about the types of the other agents in the network, Bayesian individuals attempt to learn about the types of the other agents in the network while attempting to learn about the underlying state of the world. Formally, the relevant states are not just signal endowments, but also the types of the players in the network; i.e. $\omega = (s, \eta) \in \Omega := \{0, 1\}^n \times \{0, 1\}^n$ where $\eta_i = 1$ if and only if agent i is a Bayesian agent, and a DeGroot agent otherwise. We formalize the model in Appendix A.

The incomplete information setup is an important step beyond the complete information Bayesian environment, which restricts $\pi = 1$. For instance, if an individual believes that her neighbor does not act in a Bayesian manner, she processes the information about observed decisions accordingly; as outside observers, the econometricians might think that she is not acting as a Bayesian. This is a serious problem when testing Bayesian learning, as 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.

2.1.3. DeGroot Learning. We now briefly discuss DeGroot learning (DeGroot, 1974). DeMarzo et al. (2003), Golub and Jackson (2012), and Jackson (2008) contain an extensive reviews of DeGroot models. In our experiment, we do not consider belief-communication models; instead, we consider a DeGroot model where individuals observe each others' actions in this binary environment. We call this a DeGroot action model. It has also been called a voter model (Mossel and Tamuz, 2014). The basic idea is to maintain a parallel structure with the Bayesian environment of Gale and Kariv (2003), where the state is binary but agents can only pass on their best guesses, as in the Bayesian benchmark. In action models, individuals observe the actions of their network neighbors, whereas in communication models, individuals are able to communicate their beliefs to their neighbors. One might also call these (weighted) *majority* models; individuals choose the action that is supported by a (weighted) majority of their neighborhood.

We are interested in action models for several reasons. First, observe that the corresponding models of Bayesian learning on networks are action models, so it is the appropriate comparison. A model with communication in Bayesian learning, where agents pass their posteriors or information sets, becomes fundamentally uninteresting in an environment such as ours. Generically, if agents pass on information sets, for instance, within a number of time periods equal to the diameter of the network, each individual learns everyone else's signals (Mueller-Frank, 2014). 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 measure exactly what is (or is not) 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 only observe the actions of their neighbors, such as technology usage, microfinance adoption, statement of political preferences, etc.

Let $T = T(A)$ be a weighted matrix which parametrizes the weight that person i gives to the action of person j . We study three natural parameterizations 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

$$(2.1) \quad T_{ij}^u = \frac{A_{ij}}{d_i + 1} \text{ and } T_{ii}^u = \frac{1}{d_i + 1}$$

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

$$(2.2) \quad T_{ij}^d = \frac{d_j}{\sum_{j \in N_i} d_j + d_i} \text{ and } T_{ii}^d = \frac{d_i}{\sum_{j \in N_i} d_j + d_i},$$

where N_i is the set of neighbors of individual i .

The third model is *eigenvector weighting*.⁹ An individual places weight on her neighbor proportional to the neighbor's relative importance, given by eigenvector centrality. $T^e(A)$ is given by

$$(2.3) \quad T_{ij}^e = \frac{\xi_j}{\sum_{j \in N_i} \xi_j + \xi_i} \text{ and } T_{ii}^e = \frac{\xi_i}{\sum_{j \in N_i} \xi_j + \xi_i}$$

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.

The behavior of individuals that learn according to the DeGroot model is as follows. At time $t = 0$, individuals receive signals $s = (s_1, s_2, \dots, s_n)$, and accordingly, take actions $a_{i,0} = \mathbf{1}\{s_i = 1\}$. Let $\mathbf{a}_0 = (a_{1,0}, a_{2,0}, \dots, a_{n,0})$. At $t = 1$, beliefs are denoted by $I_1 = T\mathbf{a}_0$, and actions are chosen according to $\mathbf{a}_1 = \mathbf{1}\{I_1 > 1/2\}$. Now consider time $t = k + 1$ with lagged set of actions \mathbf{a}_k . Then, beliefs are formed as indicated by $I_{k+1} = T\mathbf{a}_k$, and actions are chosen as denoted by $a_{k+1} = \mathbf{1}\{I_{k+1} > 1/2\}$. While a very natural parametrization of learning, the DeGroot model misses the inferential features that characterize Bayesian learning. 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 the limit exploiting the linear structure, as is done with DeGroot communication models, we discuss its implications below.

⁹Thanks to Matt Jackson for suggesting this alternative.




FIGURE 2. A quilt on 7 nodes, which is network 3 in our experiment.

2.2. Fundamental Properties of Bayesian and DeGroot Action Learning. We now explore the relevant features of DeGroot action learning and Bayesian learning in the topologies we chose for our experiment, in order to understand the divergence of their predictions. This helps demonstrating why differences in the learning processes lead to very different patterns for these networks, which is what we desire in the data in order to distinguish between the learning models. To fix ideas, we concentrate on topology 3 (as shown in Figure 2) in our experiment and contrast uniform DeGroot weighting with complete information Bayesian learning.¹⁰

We focus on three properties. First, networks that have “clans,” in a sense described below, have agents who get stuck on the wrong action under DeGroot learning but not under Bayesian learning. Second, under complete information Bayesian learning, an agent whose neighborhood entirely envelops one of her neighbor’s neighborhood ignores that neighbor’s action outside of her initial action which reflected her signal. This is not true under DeGroot learning. Third, under complete information Bayesian learning, an agent whose information set is dominated by a neighbor only copies her neighbor’s previous period action. Again, this is not true under DeGroot learning.

Let us begin with a benchmark. Notice that, if agents had access to all relevant information available (i.e., everyone could observe all signals), the optimal decision would then be to guess the majority, which would make agents under the Bayesian and the uniform weighting models to coincide. We later show that, in topology 3, the network shape precludes some players to have access to information in other parts of the network under DeGroot learning, but under Bayesian learning they are able to overcome this constraint.

We define a node to be *stuck* if, from some period on, the actions she chooses are opposite to the optimal decision with full information (i.e., the actions are different from the majority of signals). Next, we show that tightly knit groups, which we call *clans* (groups that have more links to the group than to outsiders), are prone to getting stuck and cannot get new information from outside of the clan under DeGroot learning.

¹⁰We focus on uniform DeGroot weighting because this is the DeGroot model that best describes the data, as seen below.




FIGURE 3. Distribution of stuck nodes in topology 3 under Bayesian and DeGroot with uniform weighting.

Given a network $G = (V, E)$ and a subset $H \subset V$, we define $G(H) := (V(H), E(H))$ as the induced subgraph for group H . That is, we only considering edges between agents in H . We write $d_i(H)$ as the degree of the induced subgraph of group H . We say that a group $C \subset V$ is a *clan* if for all $i \in C$ we have $d_i(C) > d_i(V \setminus C)$ (including herself). There are various clans in network 3 from our experiment. Triangles, such as (134) and (236), are clans, as are their unions.

We show that under uniform DeGroot weighting, if at some point a clan chooses the wrong action, an action that is the opposite of what one who observes the majority of signals would have selected, then it continues to choose it no matter how the rest of the network behaves.

Lemma 2.1. *Assume uniform DeGroot weighting. Take $C \subset V$ such that for all $i \in C$, $d_i(C) > d_i(V \setminus C)$. Assume there exists some t such that for all $i, j \in C$: $a(i, t) = a(j, t)$. Then $a(i, t + \tau) = a$ for all $i \in C$ and $\tau > 0$.*

Proof of Lemma 2.1. The proof is by induction. Without loss of generality, suppose $a(i, t) = 1$ for all $i \in C$. Of course, for $\tau = 0$ the result is true. Suppose $a(i, t + \tau - 1) = 1$ for all $i \in C$. Let $T^u(i, t + \tau) = \frac{1}{d_i + 1} \left[a(i, t + \tau - 1) + \sum_{j \in N_i} a(j, t + \tau - 1) \right]$ be the index that defines uniform weighting (as in equation 2.1) so that if $T^u(i, t + \tau) > 1/2$ then $a(i, t + \tau) = 1$, independent of the particular tie breaking rule used. We show that

$T^u(i, t + \tau) > 1/2$:

$$\begin{aligned} T^u(i, t + \tau) &= \frac{\sum_{j \in \{N_i \cup \{i\}\} \cap C} a(j, t + \tau - 1)}{d_i + 1} + \frac{\sum_{j \in N_i \cap C^c} a(j, t + \tau - 1)}{d_i + 1} \\ &\stackrel{(i)}{=} \frac{d_i(C) + 1}{d_i + 1} + \frac{\sum_{j \in N_i \cap C^c} a(j, t + \tau - 1)}{d_i + 1} \\ &\geq \frac{d_i(C) + 1}{d_i + 1}, \end{aligned}$$

using in (i) the fact that $a(j, t + \tau - 1) = 1$ for all $j \in C$. Since $d_i = d_i(C) + d_i(V \setminus C)$ for any set $C : i \in C$, and $d_i(C) > d_i(V \setminus C)$, we then have that $\frac{d_i(C) + 1}{d_i + 1} > 1/2$, as we wanted to show. \square

The fact that clans can get stuck under DeGroot learning is a key cause in the divergence between Bayesian and uniform DeGroot models in topology 3. Figure 3 shows the probability distribution of the number of stuck nodes under Bayesian and uniform weighting DeGroot; Bayesian learning coincides with full information learning 99% of the time while this happens in the DeGroot model 63% of the time.¹¹ We can also check (by inspection of each particular signal endowment) that each instance for which two or more agents are stuck in uniform weighting DeGroot comes from the implication of Lemma 2.1.

We derive below some properties of Bayesian learning that explain the predictions of Bayesian learning for this network without the need of further calculations. These properties are enough to eliminate stuckness across most (99% of) signal endowments in this example.

Formally, let $N_i^* = N_i \cup \{i\}$ be the neighborhood of agent i (including herself), and let $a_i^t = (a_{i,1}, a_{i,2}, \dots, a_{i,t})$. In general, we have that if agent i is Bayesian, and $j \in N_i^*$, then $a_{i,t}^B = \text{Function} \left(\left(a_k^{t-1} \right)_{k \in N_i^*} \right)$. The following statement is then true.¹²

Lemma 2.2. *Suppose $N_i^* \supseteq N_j^*$ and that for all $\omega \in \Omega$, $a_{j,1}(\omega) = s_j$. Then, at $t > 2$, s_j is sufficient for (a_j^t) when explaining $a_{i,t}^B$. That is,*

$$a_{i,t}^B = \text{Function} \left(s_j, \left(a_k^{t-1} \right)_{k \in N_i \setminus \{j\}} \right).$$

Proof. The proof is straightforward and therefore omitted. \square

This property follows from the fact that, if player j 's action reveals no new information j 's choice only then matters through the signal she observed; i.e. $s_j = a_{j,1}$. This implies a

¹¹Note that the Bayesian model predicts that either no one gets stuck or everyone does, since there is convergence.

¹²This is true for larger type spaces Ω , not just when it represents the signal endowments. This includes the incomplete information model that we explore below.

conditional sufficiency test: conditional on the actions of neighbors other than j , $a_{i,t}$ should be conditionally independent of $a_{j,t}$ for all $t > 1$.¹³

This property highlights the most important difference between Bayesian and DeGroot learning, which concerns how clans behave in Bayesian relative to DeGroot learning. Under DeGroot learning, agents do not adjust for the “echo chamber effect” of overweighting the opinion of the clan. However, under Bayesian learning, agents understand this concept, and after learning the information that the clan provides, pay attention only to agents outside of the group instead. This is particularly salient in the network depicted in Figure 2, for nodes 3 and 4. Note that at $t = 2$, player 3 observes her successor nodes’ (2 and 6) initial signals. Since players 2 and 6 are not connected to any other players, player 3 has observed all relevant information that these nodes can provide her with. Therefore, from $t = 3$ on, the actions of agent 3 are only a function of the history of play of agents 1 and 4, and never a function of the decisions of 2 and 6. This is also true when allowing for incomplete information about agents payoffs, and even when we add trembles, as long as agents guess their initial signals with probability 1 at $t = 1$.

An analogous reasoning is replicated by agents 2 and 6. After $t = 1$ they know that they do not have any more relevant information: only player 3’s action may reveal information about the rest of the network. Then, in the complete information model, players 2 and 6 should just copy whatever agent 2 does, since her information set strictly dominates their own.

Lemma 2.3. *In a complete information Bayesian model, take i and j such that $N_j^* \subset N_i^*$. Then*

$$a_{j,t}^B = a_{i,t-1}^B \text{ for all } t > 2.$$

Proof. The proof is straightforward and therefore omitted. \square

The properties shown in Lemmas 2.2 and 2.3 together explain why Bayesian learning does not suffer from stuckness in this network. Lemma 2.2 implies that after learning what agents in her “clan” have observed, she should only pay attention to agents outside of the group. Moreover, agents in the family also know this, so they copy the member of the family that has access to agents outside of the group.

To illustrate Lemmas 2.2 and 2.3, consider the signal endowment depicted in Panel (A) in Figure 4. Let a “blue” signal denote a signal of 1 and a “yellow” signal denote a signal of 0. In periods $t = 1$ and $t = 2$, both Bayesian and uniform weighting DeGroot predict the same behavior, which we illustrate in Panels (A) and (B) in Figure 4, respectively. As depicted in Panel (B) in Figure 4, from $t = 3$ on, agents behaving according to the DeGroot model would always keep the same action, a consequence of Lemma 2.1, which causes 4

¹³An easier corollary of Lemma 2.2 can be derived in situations where for all players in $k \in N_i^*$ such that $k \neq j$ players have not changed their actions (including i) but agent j has. Suppose $a_{k,t-2} = a_{k,t-1}$ for all $k \in N_i^* \setminus \{j\}$ and $a_{j,t-2} \neq a_{j,t-1}$. Then, player i ’s action should not respond to the change in j ’s behavior ($a_{i,t} = a_{i,t-1}$).

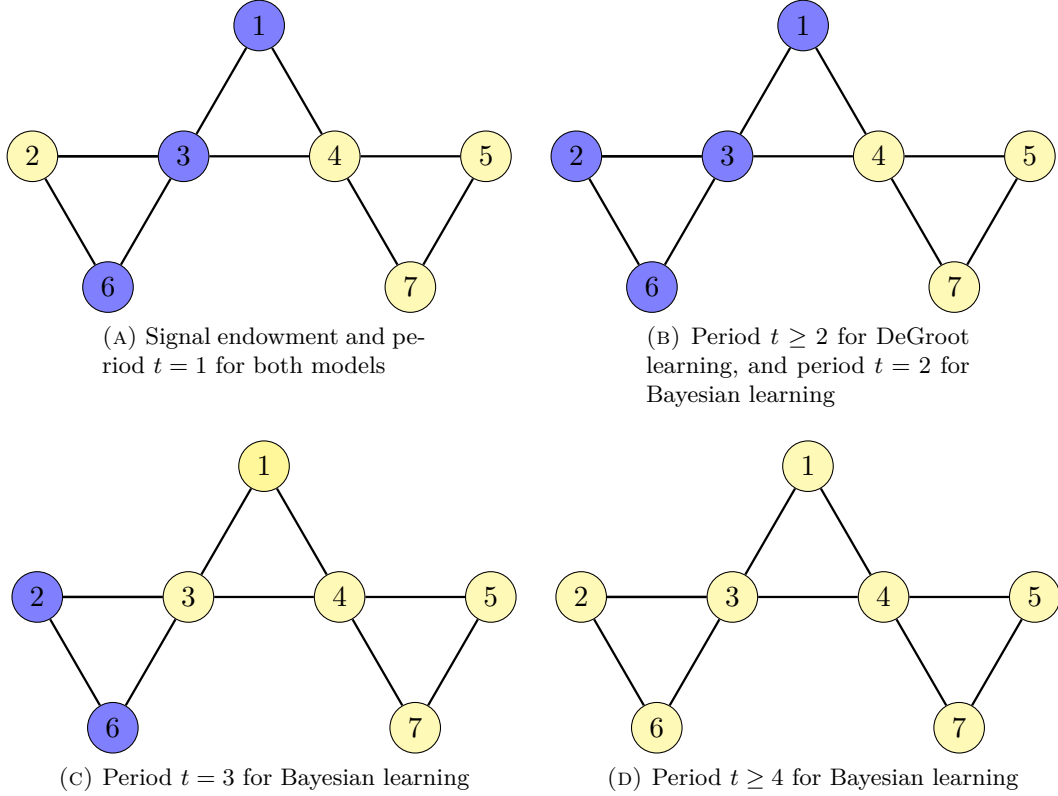


FIGURE 4

agents to be stuck, choosing the inefficient action. As illustrated in Panel (C) in Figure 4, from period $t = 3$ on, according to Lemma 2.2, agents 1 and 3, behaving as Bayesian agents, only need to learn about the initial signals obtained by agents 5 and 7, observing only player 4's decisions. After observing player 4's decision at $t = 2$, both 1 and 3 realize that the only signals consistent with her decision is that both s_5 and s_7 are equal to s_4 : if not, she would have picked $a_{4,t=2} = \text{blue}$, as 1 and 3 did. Since at $t = 2$ all agents choose according to the majority of their neighborhood, only $s_5 = s_7 = s_4$ has positive probability given their information. Knowing this, player 3's belief support is singleton, and knows that the full-information decision is $a_{3,t=3} = \text{yellow}$. Lastly, Panel (D) in Figure 4 shows that, at $t = 4$, agent 1 knows 3 knew the full signal endowment when choosing at $t = 3$, and hence $a_{1,t=4} = \text{yellow}$. With respect to players 2 and 6, we use Lemma 2.3 to predict that they copy 3's action in the previous periods.

2.3. Comparing the Asymptotic Properties of Bayesian and DeGroot Learning.

The purpose of this section is to compare how different learning models behave in achieving asymptotic learning in large networks, and thus motivate the relevance of the empirical exercise in this paper. We also compare the properties of the DeGroot communication

models, and conclude that the main difference between the DeGroot communication model and action models is the coarseness of the message space in the latter. We show that for sensible sequences of networks, both Bayesian action and DeGroot communication models lead to asymptotic efficiency (where all but a vanishing share of agents make the right decision about the state of the world) wherein DeGroot action models can yield a non-vanishing share of agents that are stuck, which highlights the importance of the empirical exercise we conduct.

Among others, DeGroot (1974), DeMarzo et al. (2003) and Golub and Jackson (2010) study DeGroot models where agents can pass on a continuous value to their neighbors as compared to a binary action. There are two interpretations of this setting: (a) the information is in the form of probabilities about a state of the world or (b) agents are estimating a state distributed with noise on a line. It is useful to interpret the model as the former case, where agents are endowed with signals that give posterior probabilities $p_i = P(\theta = 1 | s_i)$.¹⁴ In the DeGroot communication model these can be any number in $[0, 1]$, and thus, in the probability interpretation they imply non-atomic posteriors given signals. This is equivalent to having a continuum of signals, with distributions $P(s | \theta = 0)$ and $P(s | \theta = 1)$. Without loss of generality, we can assume $s_i = p_i$ directly.

In Golub and Jackson (2010) agents are able to communicate their estimate of $p_i = P(\theta | s_i)$, and aggregate their information by taking averages of the estimations of their neighbors. This is the crucial distinction between the communication model and the action model, where the former is particularly granular and the latter is considerably coarser. We study one of the particular cases they focus on, which is analogous to uniform weighting DeGroot model, where agents can communicate their estimation of the probability that $\theta = 1$ rather than a binary action, and form their next posterior belief as

$$(2.4) \quad p_{i,t} = \frac{1}{d_i + 1} \sum_{j \in N_i} p_{j,t-1}.$$

A key objective of Golub and Jackson (2010) is to find the conditions under which a sequence of growing networks achieve “wisdom of crowds” asymptotically. Formally, take a sequence of networks $G_n = (V_n, E_n)$ where $|V_n| = n$, let $p_{i,t}^{(n)}$ be defined as in 2.4 for the neighborhood $N_i(n) \subset V_n$. They define a sequence of graphs G_n to be *wise* (in the uniform weighting model) if

$$(2.5) \quad \text{plim}_{n \rightarrow \infty} \max_{i \in V_n} \left| \lim_{t \rightarrow \infty} p_{i,t}^{(n)} - P(\theta = 1) \right| \rightarrow 0.$$

Note that if a sequence of networks is wise, then when choosing $a_i \in \{0, 1\}$ as the best response to those beliefs, asymptotically they would take the correct decision, as if all

¹⁴This is consistent with the results presented in Golub and Jackson (2010), since they assume that $p_i \in [0, 1]$.

information was observed. Formally, we say that a sequence is *asymptotically efficient* if $\lim_{t \rightarrow \infty} \left| a_{i,t} - \mathbf{1} \left(p_{i,t} > \frac{1}{2} \right) \right| = 0$.

We want to compare the DeGroot communication model (with richer communication channels than the action model we consider) to the performance of the Bayesian and DeGroot action models we test in our experiments, where the beliefs $p_{i,t}$ are derived from Bayes rule or by making weighted averages of past actions, rather than past beliefs. To make this comparison, we need to consider a common environment, so we will assume that agents observe a continuum of signals, rather than the binary signals of our experiment. We see that under very general conditions, both the DeGroot model of [Golub and Jackson \(2010\)](#) and the Bayesian model achieve asymptotic efficiency, but the uniform weighting DeGroot model may not.

Proposition 2.1. *Suppose $G_n = (V_n, E_n)$ with $|V_n| = n$ is such that (i) $d_i(V_n) \leq \bar{d}$ for all i, n (i.e., they have uniformly bounded degrees), (ii) the posterior distributions $P(\theta | s_i)$ are non-atomic in s , for $\theta \in \{0, 1\}$, and (iii) signals are i.i.d across agents. Then*

- (1) *the DeGroot communication model is asymptotically efficient,*
- (2) *the Bayesian action model is asymptotically efficient, and*
- (3) *the DeGroot action model may not be asymptotically efficient.*

In particular, suppose there is a sequence of groups $A_n \subset V_n$ such that $A_n = \bigcup_{k=1}^{h_n} C_k^n$ is a strictly increasing family of disjoint clans with uniformly bounded size. Then the DeGroot action model is not asymptotically efficient.

Proof. The efficiency of the DeGroot communication model is a consequence of Corollary 1 in [Golub and Jackson \(2010\)](#). The result on Bayesian action models is the central Theorem in [Mossel et al. \(Forthcoming\)](#). The inefficiency of DeGroot action model is a simple corollary of Lemma 2.1, since as $h_n \rightarrow \infty$ and signals are non-atomic and i.i.d., then almost surely they will be at least a clan $C_{k^*}^n$ such that $a_{i,1} \neq \theta$ for all $i \in C_{k^*}^n$. \square

Is worth noting that the asymptotic efficiency result does not crucially depend on the non-atomicity of posteriors, but rather on the possibility of ties (when $p_{i,t} = 1/2$). In networks where ties do not occur (or occur with vanishing probabilities) then the asymptotic efficiency result is also true for the case of binary signals (an implicit result in [Mossel et al. \(Forthcoming\)](#), which is also shown in [Menager \(2006\)](#)). In Appendix C we also illustrate the impact of this source of asymptotic inefficiency in a particular growing class of networks, which are growing extensions of network 3 from our experiment.

[Golub and Jackson \(2010\)](#) show that the asymptotic result generalize to a greater set of weighting formulas, as long as the influence of a single agent on the average of everyone in the network is arbitrarily small as the network size grows. A similar result of inefficiency would also be true for the DeGroot action model as well, using a variant of Lemma 2.1.

Proposition 2.1 also illustrates the main discrepancy between these models in terms of social learning and why differentiating between which model is used is relevant. If agents

cannot communicate (in a credible manner) all the information they have observed (for example, can only show what they choose or agents just observe neighbors’ decisions) then the coarseness of these messages cause the averages taken to get “stuck” on a particular action, not allowing the flow of new information. This is prevented in the DeGroot communication model by allowing $p_{i,t}$ to take any number in $[0, 1]$, so small changes in beliefs are effectively passed through to other agents. Action models make more sense in situations where agents need verifiable information through costly signals (e.g., the actions they take) to believe the information to be credible or environments where the state or beliefs are complex to describe and communicate.

2.4. An Example from Indian Village Networks. The above results have shown that whether asymptotic efficiency is reached or not depends on the structure of networks in question. In this section we explore real-world network data to assess whether the problem of DeGroot action learning is relevant in realistic social network structures.

We consider data from [Banerjee et al. \(2013\)](#) consisting of detailed network data in 75 villages in Karnataka, India. We consider the networks consisting of information relationships, which in other work we have shown to correlate strongly with favor exchange relationships. We show that a DeGroot action learning process is likely to get stuck.




FIGURE 5. Both panels (A) and (B) present CDFs of the fraction of nodes that initially received the signal $1 - \theta$ that became stuck at the wrong belief for various levels of p . Panel (A) presents results where we conduct simulations using the 75 Indian village networks from [Banerjee et al. \(2013\)](#). Panel (B) presents the same results for Erdos-Renyi networks which have an average degree that matches the Indian network data.

In Panel (A) we examine simulations conducted on the 75 empirical networks and quantify the degree of stuckness. Specifically, with a signal quality of 0.55, in the median village, at least 78% of the nodes that initially received the wrong information stay stuck at the

wrong belief. Even with a signal quality of 0.66, in the median village, at least 37% of households that initially received the wrong information stay stuck. Furthermore, in over 1/3 of the villages, at least 50% of households that initially received the wrong information stay stuck. In Panel (B) we repeat the exercise but for Erdos-Renyi graphs calibrated to have an average degree that matches the empirical data. We find, similarly, that DeGroot action learning is likely to get stuck. But furthermore, by comparing the simulations under the Indian village networks and the corresponding Erdos-Renyi graphs, we can see that the problem is somewhat exacerbated for the empirical networks. For instance, 90% of empirical network simulations have at least 35% of nodes failing of to learn whereas the corresponding number is at least 18% for Erdos-Renyi graphs. This suggests evidence that, as shown in Jackson et al. (2012), the networks organized to aid informal transactions dealing with limited commitment have generated structures that are prone to misinformation traps.

3. EXPERIMENT

3.1. Setting. We conducted 95 experimental sessions with the three chosen networks across 19 villages in Karnataka, India. The experiments had 665 total subjects. The villages range from 1.5 to 3.5 hours’ drive from Bangalore. We chose the village setting because social learning through networks is of the utmost importance in rural environments; information about new technologies (Conley and Udry, 2010), microfinance (Banerjee et al., 2013), political candidates (Cruz, 2012; Cruz et al., 2015), among other things, propagates regularly through social networks.

3.2. Implementation and Game Structure. In every village, we run an average of 5 sessions, each with 7 participants, given that each of the networks of our experiment had 7 nodes. We recruited an average of 35 individuals from a random set of households from each village. We brought the individuals to a public space (e.g., marriage hall, school, dairy, barn, clusters of households) where we conducted the experiment. While individuals were recruited, the public space was divided into “stations.” In each station there was a single surveyor to monitor the single participant assigned to the station at random. This ensured that participants could not observe each other nor could they communicate. Often times, stations would be setup across several buildings.

In each village, individuals anonymously played the social learning game three times, each time with a different network structure. The three networks (see Figure 6) were played with a random order in each village to avoid order effects. At the beginning of each game, all participants were shown two identical bags, one with five yellow balls and two blue balls and the other, with five blue balls and two yellow balls. One of the two bags was chosen at random to represent the state of the world. Since there was an equal probability that either bag could be chosen, we induced priors of 1/2. As the selected bag contained

five balls reflecting the state of the world, participants anticipated receiving independent signals that were correct with probability $5/7$.

After an initial explanation of the experiment and payments, the bag for the first game was randomly chosen in front of the participants. The participants were then assigned to stations where each was shown a sheet of paper with the entire network structure of seven individuals for that game, as well as her own location in the network. The neighbors' past decisions were also communicated to subjects on sheets of paper that presented an image of the network and colored in their neighbors' guesses.

Once in their stations, after receiving their signals in round zero, all participants simultaneously and independently made their best guesses about the underlying state of the world (i.e., 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 this round, each participant was shown the round one guesses of the other participants in her neighborhood through the mentioned procedure. Agents updated their beliefs about the state of the world and then again made their best guesses about it. Once again, the game continued to the following round randomly. This process repeated until the game came to an end. Notice that, after the time zero set of signals, no more signals were drawn during the course of the game. Participants could only observe the historical decisions of their neighbors and update their own beliefs accordingly. Importantly, so that we would not bias individuals' updating against Bayesian learning individuals kept the information about the guesses of their neighbors in all previous rounds until the game concluded. After each game, participants were regrouped, the color of the randomly chosen bag was shown, and if appropriate, a new bag was randomly chosen for the next game. Participants were then sent back to their stations and the game continued as the previous one. After all three games were played, individuals were paid Rs. 100 for a randomly chosen round from a randomly chosen game, as well as a Rs. 20 participation fee. Participants then faced non-trivial incentives to submit a guess that reflected their belief about the underlying state of the world. The incentive was about three-fourths of a daily agricultural wage.

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, 2012)) make use of several 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, 2012) lack power to pit Bayesian learning against the DeGroot alternatives posited above. Panel A of Table 1 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. (2012). In two of the networks, there are no differences between the equilibrium paths of Bayesian learning and the uniform and degree weighted DeGroot

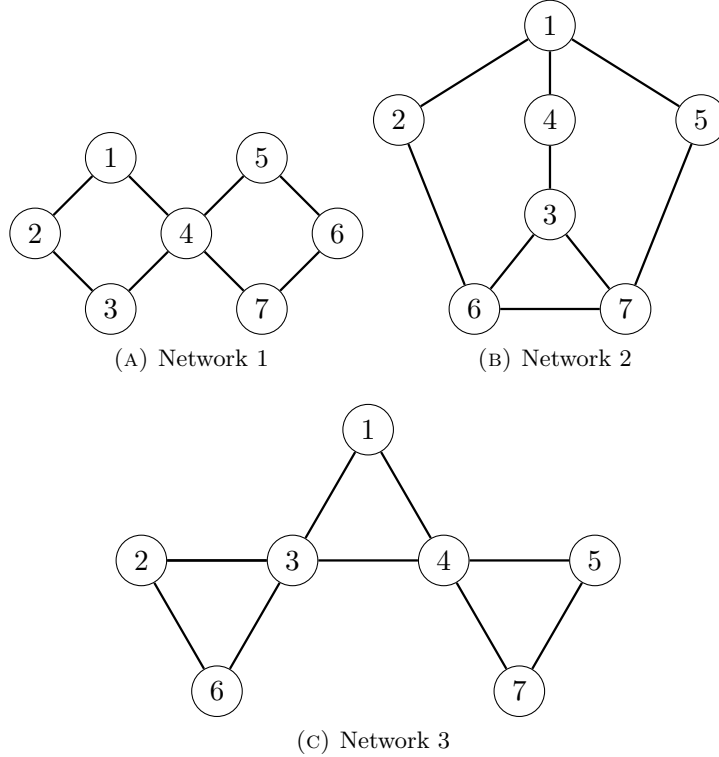


FIGURE 6. Network structures chosen for the experiment.

alternatives (with a difference with eigenvector weighting of 6.8% in one of the networks). While in the third network, the differences are on average 15.5% of the observations, these are largely an artifact of many indifferences along equilibrium path for the DeGroot models, which were broken by choosing the individuals' past action.

Given our goal of separating between Bayesian and DeGroot alternatives, we move to an environment with seven agents as opposed to three, so that we obtain more power to distinguish between these models while still maintaining computational tractability. Additionally, we consider a quality signal ($p = 5/7$) to prevent that the separation from different learning models is driven by indifferences along equilibrium path, and thus, the tie-breaking rule.¹⁵

In order to select our three networks, we initially considered all connected and undirected networks with seven nodes. Additionally, to prevent indifferences along equilibrium path, we restricted to networks where the minimum degree of a given individual is two. Further, to avoid confusing our experimental subjects, we restricted to planar networks (networks that can be drawn so that their edges intersect only at their endpoints). Next, we established a model selection criterion function. This criterion function depended on the power to detect a DeGroot alternative against a complete information Bayesian null, using

¹⁵Moving to eight agents, for instance, would be exponentially more difficult for our structural estimation.




FIGURE 7. (A) depicts the power and divergence frontier for degree weighting DeGroot. (B) shows the power and divergence frontier for uniform weighting

our pilot data to generate an estimate of the noise, as well as a divergence function. The divergence function measures the share of node-time observations for which the Bayesian model (with $\pi = 1$) and a DeGroot model pick different actions,

$$D(G) := \frac{1}{n(T-1)} \sum_{s \in S} \sum_{t=2}^T \sum_{i=1}^n \left| a_{i,t}^B(s | G) - a_{i,t}^m(s | G) \right| \cdot P(s | \theta = 1),$$

where $a_{i,t}^B(s | G)$ is the action predicted under the Bayesian model and $a_{i,t}^m(s | G)$ is the action predicted under DeGroot with m -weighting, where m is uniform, degree, or eigenvector weighting.¹⁶

Figure 7(A) depicts a scatter plot of power and divergence for network 1 in Figure 6. We see that our network 1 is the best ex-ante choice to separate between the incomplete information Bayesian and DeGroot degree weighting models. Figure 7(B) illustrates the analogous figure for network 2 in Figure 6, and highlights that it is the best ex-ante choice to separate between the Bayesian and DeGroot uniform weighting models. Lastly, the choice of network 3 in Figure 6 was motivated by our exercise in section 2, and by the fact that it performed very well in separating between the Bayesian and both the DeGroot uniform and degree weighting models.

4. TESTING THE THEORY

4.1. Estimation and Inference. In order to test how well model m fits the data in session r , we use the fraction of discrepancies between the actions taken by individuals in

¹⁶ $P(s | \theta = 1) = p^{\sum_i s_i} (1-p)^{n-\sum_i s_i}$.

the data and those predicted by the model. This is given by

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

where $D_{i,t,r}^m = |a_{i,t,r}^{obs} - a_{i,t,r}^m|$ which computes the share of actions taken by players that are not predicted by the model m .¹⁷ To examine how poorly model m predicts behavior over the entirety of the data set, we define the *divergence function* as

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

This is simply the average discrepancy taken over all sessions. Model selection is based on the minimization of this divergence measure. Note that we include dependency on π , the share of Bayesian agents believed to be in the population, since for the Bayesian model, the prediction $a_{i,t,r}^m$ depends on π . The computation of the Bayesian actions are described in Appendix A.

While 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. First, we may look at the network level, which considers the entire social learning process as the unit of observation; and second, we may study the individual level wherein the unit of observation is an individual's action at an information set.

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 that model m predicts is given 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 $t - 1$ and the ex-ante probability that a given individual is a Bayesian learner as opposed to some DeGroot-alternative learner: $a_{i,t,v}^m$ is the action predicted for agent i at time t in session r , given information set $P_{i,r,t}$ and π .

For every DeGroot alternative, we consider the probability that minimizes the divergence function:

$$\hat{\pi}^m = \operatorname{argmin}_{\pi \in [0,1]} \mathbf{D}(m; \pi)$$

where m indexes an incomplete information Bayesian model with each of the possible DeGroot alternatives: uniform, degree, and eigenvector. In order to perform inference on

¹⁷Since all models and all empirical data have a fixed first action (given by the signal endowment), the first round should not enter into a divergence metric. Thus, we restrict attention to $t \geq 2$.

the parameter, we perform a block Bayesian bootstrap that accounts for the dependence in the data of the individuals playing the same game (see a similar procedure used in Banerjee et al. (2013)).

Equipped with the minimizing value of $\hat{\pi}^m$ for $m \in \{u, d, e\}$, we are prepared to conduct our analysis. Note that procedure is done both at the network level and the individual level. In particular, in addition to simply identifying the best-fitting model, we can go one step further and ask whether incorrect beliefs about others' types can explain the data. Specifically, given $\hat{\pi}^m$, we can ask whether the divergence can be minimized further by, for instance, drawing a population of all Bayesian agents who have heterogeneous priors and lack common knowledge of Bayesian rationality, and therefore are employing $\hat{\pi}^m$ as a mistaken belief. We are able to assess whether deviations from correct beliefs can rationalize the data better than a singular DeGroot alternative. Finally, we can also look at a non-nested hypothesis test of how the model with common knowledge of Bayesian rationality compares to each of the DeGroot alternatives, and how the DeGroot models compare to each other in terms of explaining the data.

4.2. Learning at the Network Level. We begin by treating the social network and the entire path of actions as a single observation.

4.2.1. Comparing DeGroot and Complete Information Common Knowledge Bayesian Models. Before looking at the incomplete information estimation result, we begin by first looking at comparisons of the three DeGroot models and the common knowledge, complete information Bayesian model. Figure 8 presents the data in a graphical manner and Table 2 presents the results for non-nested hypothesis tests comparing each of the models in a pairwise manner.




FIGURE 8. Fraction of actions unexplained at the network level

As seen in Figure 8, across all the networks, uniform weighting fails to explain 12% of the data, degree weighting fails to explain 14% of the data, eigenvector-centrality weighting fails to explain 13.5% of the data, and complete information Bayesian learning fails to explain 18% of the data. This suggests that the DeGroot models, as well as the Bayesian learning models each explain more than 80% of the observations, but the DeGroot models do considerably better.

Turning to the pairwise comparisons of fit, we conduct a non-nested hypothesis test (Rivers and Vuong, 2002) using a nonparametric bootstrap at the session-game level, wherein we draw, with replacement, 95 session-game blocks of observations and compute the network level divergence.¹⁸ This procedure is analogous to clustering and, therefore, is conservative exploiting only variation at the block level. 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 2 presents the p -value results of the inference procedure. Note that most of the values are essentially zero. First, looking across all topologies both separately and jointly, we find evidence to reject the Bayesian model in favor of all the DeGroot alternatives. Second, we find that uniform weighting dominates every alternative across every topology both separately and jointly. Ultimately, the bootstrap provides strong evidence that the uniform-weighting DeGroot model best describes the data generating process when analyzed at the network level.

4.2.2. Incomplete Information Bayesian Learning. We now present our main results using the network level divergence. The estimation algorithm is described in Appendix A. Panel A in Table 3 displays the minimizing parameter estimate π for in incomplete information Bayesian learning model with each of the DeGroot alternatives. We robustly find that the minimizing parameter value is $\hat{\pi}^m = 0$ for every model m . This suggests that, if the common knowledge parameter π truly describes the population distribution, essentially 0% of the population is Bayesian and any Bayesian agent believes 0% of the population is Bayesian. Doing a session level bootstrap, we estimate the standard errors as 0.08, 0.12 and 0.09 respectively across uniform, degree and eigenvector alternatives. This shows that despite the uncertainty, at best only a very small share of agents could likely be Bayesian.

Our results then indicate that when we estimate a model of incomplete information learning with potentially Bayesian agents, the model that best describes the data is one that is equivalent to having no Bayesians whatsoever and instead describing each agent as DeGroot. Moreover, the results of Table 2 indicate that that the best fitting such model is one with uniform weighting.

¹⁸We have 95 village-game blocks in networks 1 and 2, and 75 for network 3. We redraw with replacement the same number that we have in our empirical data.

To give the Bayesian social learning model another shot at better describing the experimental data, we conduct a second exercise. We consider the case where all agents are Bayesian, but we relax the common prior assumption. Specifically, we allow for each agent to be Bayesian, know that she is Bayesian, but be uncertain about whether others are Bayesian or not. So each agent believes that a π share of the population is non-Bayesian, despite the fact that everyone is indeed Bayesian. We then compute the divergence minimizing π for a model where all agents are Bayesian but there is no common knowledge of Bayesian rationality but instead is a miscentered belief on the distribution of Bayesian types via heterogenous priors. Here we find the best fitting parameters across all networks to be $\hat{\pi} = 0$ for every model m (Panel C of Table 3). Unsurprisingly, however, the standard errors are larger in this case. By looking at the divergence at the optimal π , we can see that drawing individuals from the distribution given by π^m fits the data considerably better than assuming all agents are Bayesian but incorrectly believing that others could be DeGroot types (see Panels B versus D).

To summarize, as illustrated by Figure 9, whether considering common knowledge of Bayesian rationality or not, the robust best explanation of the data is the simplest model with $\pi = 0$. Here every agent is DeGroot.

4.3. 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.¹⁹

4.3.1. Complete Information Bayesian Learning. We begin by calculating the individual level divergence for the DeGroot models and the model where all agents are Bayesian and commonly know all are Bayesian.²⁰ This is depicted in Figure 10.

First, uniform weighting systematically outperforms degree weighting (0.0648 versus 0.1029), and degree weighting outperforms eigenvector weighting by a small margin (0.1029 versus 0.1097). Second, 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 94% of all individual observations. Degree and eigenvector centrality weighting models predict 90% and 89% of all individual observations, respectively. Finally, the common knowledge of Bayesian rationality model performs rather poorly, predicting 74% of all individual observations, and consequently, significantly underperforms

¹⁹It is certainly ex-ante possible that agents themselves do not each behave according to a particular model while the aggregate social group may best be described by such a model.

²⁰When an agent faces a tie, they stay with their previous action. We considered a random tie-breaking alternative as well, which does not substantively change the results. Importantly, as explained below, Figure 10 restricts to the observations in the support with positive probability under the complete information Bayesian learning model.

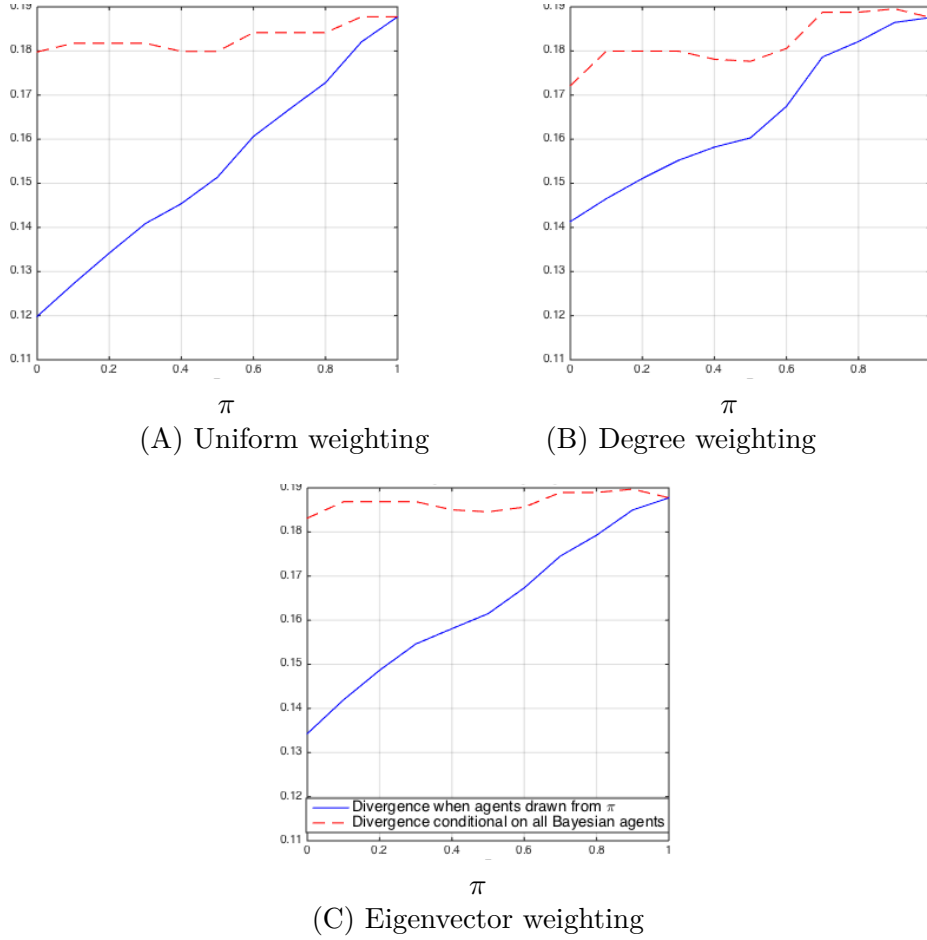


FIGURE 9. Fraction of actions unexplained by incomplete information model at various π . We show expected divergences (where we plot the share of actions unexplained under π share Bayesian agents where each agent is drawn Bayesian with probability π and DeGroot with probability $1 - \pi$). We also show divergences when all agents are Bayesian but mistakenly think that other agents could be DeGroot with probability $1 - \pi$.

at explaining the data relative to all-DeGroot models. Accordingly, Table 4 provides the hypothesis tests for the non-nested model selection procedure to show that the complete information Bayesian-learning model can be rejected in favor of DeGroot-learning alternatives.

We also provide a non-parametric test using just network 3 in our sample. Notice that in this network peripheral nodes that behave according to the Bayesian-learning model should follow the action of its parent node in the graph in any period $t > 3$. This is because the peripheral nodes' information sets are dominated by those of the parent node. Table 5 shows that only 17% of the time when the Bayesian and DeGroot models predict contradicting guesses do the agents actually take the Bayesian decision. This means that




FIGURE 10. Fraction of actions unexplained at the individual level

around 83% of the time, agents are ignoring information dominance and going with a (weighted) majority of their neighborhood.

4.3.2. Zero Probability Information Sets. When considering individual Bayesian learning, observe that there is a possibility that, empirically, Bayesian agents may arrive to an information set that has zero probability of occurrence. This is not a conceptual problem when we study the network-level divergence, because any deviation from the empirical data, irrespective of the history, is penalized. However, this is problematic for identification when we look at the individual level divergence, 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.

The degree to which this is a problem is highlighted in Table 6. We show the share of individuals, the share of sessions, and the share of observations (a given individual in a given round in a given game) that encounter zero probability information sets. In Panel A, we look at the zero probability information sets reached under the complete information common knowledge model of Bayesian learning. We find that about 45% of individuals, 70% of sessions and 40% of observations hit zero probability information sets. In Panels B, C, and D we observe that these numbers modestly decline when we allow for incomplete information about others' types. While this certainly does not eliminate the problem, it does provide immediate evidence of the potential value of allowing our Bayesian agents to doubt the Bayes-rationality of her peers.

One natural way to eliminate the zero probability information set problem entirely is to introduce disturbances. In section 4.3.4, 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 a la Choi et al. (2012). In such a

model, individuals can make mistakes with some probability, and Bayesian agents, knowing the distribution of these disturbances, integrate over this possibility when updating. We show that this approach is computationally infeasible in our context in Proposition 4.2.

Since guaranteeing full support in this model by introducing trembles creates computational problems that we describe below, let us start by discussing what we make of the zero probability events we encounter. First, we argue that 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 incomplete information Bayesian model, but we lack a theory of what individuals do once they hit zero probability events. If this is the case we may take two different approaches: we could be agnostic about the correct off equilibrium beliefs, or 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.

Instead, we take what we believe to be a reasonable approach by considering two alternatives. In the first case, we penalize each zero probability observation as an error. In the second case, we only look at observations in the support of observations with positive probability of occurrence under the Bayesian learning model (meaning we do not penalize the Bayesian model for arriving at zero probability information sets). We report results for both to paint a complete picture.

4.3.3. Incomplete Information Bayesian Learning. We now turn to studying which incomplete information Bayesian learning model best fits the data, taking these constraints into account. 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 A.

Table 7 presents the results for the case where we use all rounds of data. We find that $\hat{\pi}^u = \hat{\pi}^d = \hat{\pi}^e = 0$. Further, the standard errors are quite small: 0.014, 0.024, and 0.011, respectively for uniform, degree and eigenvector weighting. In addition, in Panel B, we see that at the optimum (as noted in Figure 10) the divergences are small.

In Table 8, we show results where we constrain ourselves to the support of observations with positive probability of occurrence under the Bayesian learning model. Here we lose a tremendous amount of information. In Panel B, we see that the fraction of observations unexplained rises steeply (almost three-fold) relative to the case where we use all

observations. This means that the DeGroot model is explaining behavior precisely when the Bayesian model, even with incomplete information, runs into a zero probability information set. Nonetheless, we find that the divergence minimizing π to $\hat{\pi}^u = \hat{\pi}^d = 0$ and $\hat{\pi}^e = 0.9$. Unsurprisingly, the standard errors are considerably larger in this case (0.32, 0.28, and 0.31, respectively) and the bootstrapped confidence intervals practically cover the entire interval. In sum, if we take the view that we must remain agnostic about the zero probability information set observations, then we do not have enough information at the individual level to distinguish the model for any π .

Finally, as before, we explore what happens if we consider the Bayesian setting but remove common knowledge of Bayesian rationality. In this case, while the objective function minimizing parameter estimates are typically far from 0, the divergences are significantly higher from those in the incomplete information case with common knowledge of distribution of types.

In summary, we have shown that, when we consider the data at the individual level, because even the incomplete information Bayesian model reaches zero probability information sets so frequently, when we consider the data inclusive of all observations, the best explanation is a simple model where all agents are DeGroot. At the same time, the share of data unexplained rises steeply when we restrict attention to only observations where we have not encountered a zero probability information set. This suggests that the DeGroot model precisely explains the data where the Bayesian models are failing. However, if we take the extreme view to not penalize the Bayesian model at all on zero probability information sets, then, when conditioning on the history, we do not have the power to distinguish between any model of individual learning with any π .

4.3.4. Bayesian Learning with Disturbances and Complexity Problems. Now we discuss expanding the support by introducing disturbances. For expositional simplicity, 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 choose. This guarantees full support: any agent can take any action given any history with positive probability.²¹

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

²¹Haile 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.

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 in the number of rounds, there is no faster algorithm.²² 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 $\Theta(T)$.²³ Moreover, it is asymptotically tight; i.e., any algorithm implementing Bayesian learning must have running time of at least $\Theta(T)$.*

Proof. See the computation in Appendix A.5. □

Specifically, the algorithm is $\Theta(n4^n T)$. Notice that if n was growing this algorithm would be exponential time, but here, n is a constant. 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 $\Theta(4^{n(T-1)})$.*

Proof. See the computation in Appendix A.5. □

In order to illustrate the complexity of these algorithms with trembles, we compare them to their deterministic counterparts. For the complete information model, the algorithm with trembles with $T = 6$ involves 1.19×10^{16} more computations than the deterministic model. With the same T , the incomplete information model involves 8.65×10^{32} more calculations than its deterministic counterpart. To see how severe the burden is, suppose that the deterministic complete information model takes 1 second to run. Then the deterministic incomplete information model (again, without trembles) takes 4 and a half hours. The trembling hand complete information model, however, takes approximately 377,346,524 years.

Note that the above exercise looked at the number of computations for a specific algorithm – one that was efficient for the case without disturbances. An objection could be made that there may be, perhaps, a more efficient algorithm. In Appendix A.5, we discuss the relationship between our current problem and work on computational complexity in Bayesian networks, which gives the reader intuition as to why we believe it is likely impossible, if not in the least very difficult, to derive efficient algorithms with trembles in this context.

²²Our environment consists of finite graphs where n does not grow in T .

²³Recall that we say $f_1(n) \in \Theta(f_2(n))$ if f_1 is asymptotically bounded above and below by f_2 , up to a multiplicative constant. Formally, if $\exists c_1, c_2 > 0, \underline{n}$ such that $\forall n > \underline{n}, c_1 \cdot |f_2(n)| < |f_1(n)| < c_2 \cdot |f_2(n)|$.

4.4. Summary. To summarize this section’s results, first we have presented evidence that the considered model of Bayesian learning result often arrives at zero probability information sets. This can be taken as evidence against these particular models. By relaxing the model to incomplete information, we can recover some, but not all, of these events. Second, we provide an argument to show 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. We lose considerable information in this case. If we penalize the Bayesian model for zero probability events, the result is particularly strong with robust estimates suggesting $\pi = 0$ irrespective of the DeGroot alternative. Finally, we show that one cannot better fit the data by adding heterogeneous priors where all agents are Bayesian but each incorrectly believes that each other may be DeGroot with probability π . The results robustly support the results we also found in the network level analysis: DeGroot models with simple weighting (such as uniform weighting) provide a good explanation of the data against a large class of alternatives and information structures, though in this case, eigenvector weighting does nearly as well.

5. WHY A LAB EXPERIMENT WITH STRUCTURAL ESTIMATION

In this section, we discuss two reduced form approaches to study the experimental data. The motivation is the following. 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. 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 on which we focus concerns double counting information and was inspired by the work of Möbius et al. (2015). Under any of the aforementioned Bayesian models, Bayesian agents should not double-count information. DeGroot agents do double-count information, however, and it is on this intuition that we build the exercise. The work of Möbius et al. (2015) explores an idea like this to look for tagged information, which is different from the question of Bayesian versus DeGroot learning. In our setting, tags are shut down.

We provide two examples of regressions that 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 explores whether individuals treat old information that cycles back to them as if it is new, additional information. The naive 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 of 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. We can also show, available upon request, that we can generate DeGroot-like reduced form parameters for a number of network topologies under true Bayesian learning, and vice versa.

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, let N_{2i} be 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.

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 the two most intuitive DeGroot alternatives.

Table 9 provides the simulation and empirical results. When looking at the empirical results, we see that, as expected, an individual's own signal being one and the share of individuals in one's neighborhood with signals of one increase the probability of the final guess being one. However, we can reject that $\beta_{3l+1} > \beta_{3l}$. While this seems to be inconsistent with the intuition that agents engage in double-counting, the simulation exercise shows that these patterns cannot necessarily be interpreted in that manner. First, columns (3) and (4) indicate that the Bayesian null does not have coefficients that are

near identical across multiplicities 1 and 2. Moreover, columns (5) and (6) show that when we look at DeGroot learning with degree weighting, it is not the case that $\beta_{3,2} > \beta_{3,1}$. More generally, 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, set of initial signals, and the particular communication channels involved.

5.2. Historical Information. Another reduced form analysis that one may conduct is addressing 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 naive 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. When looking at the empirical results, we see that as expected, an individual's own signal being one and the share of direct and new indirect neighbors with signals of one increase the probability of the final guess being one. Also, the empirical results show that the share of repeated indirect neighbors with signals of one increase the probability of the final guess being one, that is, $\alpha_4 > 0$. 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, $\alpha_4 > 0$, 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 are better fit by models where individuals construct beliefs in a Bayesian manner or by models where individuals are myopic and instead follow DeGroot rules of thumb to compute beliefs. To do so, we developed a simple experiment where we designed networks in order 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 the experimental data, we were able to study the social learning process as a whole by taking the network as the unit of observation and studying the behavior of an individual, which addresses whether an agent acts in a Bayesian manner. Furthermore, we are able to relax both common knowledge of Bayesian rationality, by allowing agents to be either Bayesian or DeGroot in their learning, modeled as agents having privately observed (epistemic) types (as in [Harsanyi \(1967\)](#)). Additionally, we relax the common prior assumptions, allowing all agents to be Bayesian (and therefore know it) yet remain unsure as to whether others are Bayesian or DeGroot.

At the network level, we find evidence that the uniform 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 experimental data, largely owing to the fact that in many circumstances, the Bayesian prescription did not differ from the DeGroot prescription.

At the individual level, we find that uniform weighting DeGroot performs the best, outperforming 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 is rather consistent with DeGroot action models where individuals myopically weight their neighbors' actions when updating their own beliefs rather from a Bayesian model. This may imply that social learning processes, in contexts where agents observe each others' actions, 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 (2011): "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78, 1201–1236.
- ANDERSON, L. AND C. HOLT (1997): "Information cascades in the laboratory," *The American Economic Review*, 87, 847–862.
- AUMANN, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4, 1236–1239.
- 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 (2013): "The Diffusion of Microfinance," *Science*, 341, 1236–1239.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades," *Journal of Political Economy*, 100, 992–1026.
- BRANDTS, J., A. EBRU GIRITLIGIL, AND R. A. WEBER (2015): "An Experimental Study of Persuasion Bias and Social Influence in Networks," *European Economic Review*, 80, 214–229.
- CELEN, B., S. KARIV, AND A. SCHOTTER (2010): "An Experimental Test of Advice and Social Learning," *Management Science*, 56, 1–15.
- CHOI, S. (2012): "A Cognitive Hierarchy Model of Learning in Networks," *Review of Economic Design*, 16, 215–250.
- 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.
- (2012): "Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data," *Review of Economic Design*, 16, 93–118.
- CHOI, S., E. GALLO, AND S. KARIV (2015): "Networks in the laboratory," *WP, March*.

- CONLEY, T. AND C. UDRY (2010): “Learning about a new technology: Pineapple in Ghana,” *The American Economic Review*, 100, 35–69.
- CORAZZINI, L., F. PAVESI, B. PETROVICH, AND L. STANCA (2012): “Influential listeners: An experiment on persuasion bias in social networks,” *European Economic Review*, 56, 1276–1288.
- CRUZ, C. (2012): “Social Networks and the Targeting of Illegal Electoral Strategies,” .
- CRUZ, C., J. LABONNE, AND P. QUERUBIN (2015): “Politician Family Networks and Electoral Outcomes: Evidence from the Philippines,” .
- 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.
- EYSTER, E. AND M. RABIN (2010): “Naive Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–243.
- (2014): “Extensive Imitation is Irrational and Harmful,” *The Quarterly Journal of Economics*, 129, 1861–1898.
- GALE, D. AND S. KARIV (2003): “Bayesian learning in social networks,” *Games and Economic Behavior*, 45, 329–346.
- GEANAKOPOLOS, J. (1994): “Common Knowledge. chapter 40 of volume 2 of the Handbook of Game Theory, edited by R. Aumann and S. Hart,” .
- GEANAKOPOLOS, J. AND H. POLEMARCHAKIS (1982): “We can’t disagree forever,” *Journal of Economic Theory*, 28, 192–200.
- GOLUB, B. AND M. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–149.
- (2012): “How homophily affects learning and diffusion in networks,” *Quarterly Journal of Economics*, 127, 1287–1338.
- HAILE, P., A. HORTAÇSU, AND G. KOSENOK (2008): “On the empirical content of quantal response equilibrium,” *The American Economic Review*, 98, 180–200.
- HARSANYI, J. C. (1967): “Games with Incomplete Information Played by “Bayesian” Players, I-III. Part I. The Basic Model,” *Management Science*, 14, 159–182.
- 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. O. (2008): *Social and Economic Networks*, Princeton University Press.

- JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (2012): “Social Capital and Social Quilts: Network Patterns of Favor Exchange,” *American Economic Review*, 102, 1857–1897.
- JADBABAIE, A., P. MOLAVI, A. SANDRONI, AND A. TAHBAZ-SALEHI (2012): “Non-Bayesian Social Learning,” *Games and Economic Behavior*, 76, 210–225.
- 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.
- LOBEL, I. AND E. D. SADLER (forthcoming): “Information diffusion in networks through social learning,” *Theoretical Economics*.
- MENAGER, L. (2006): “Consensus, communication and knowledge: An extension with Bayesian agents,” *Mathematical Social Sciences*, 51, 274–279.
- MENGEL, F. AND V. GRIMM (2014): “An Experiment on Learning in a Multiple Games Environment,” .
- MÖBIUS, M., T. PHAN, AND A. SZEIDL (2015): “Treasure Hunt: Social Learning in the Field,” .
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): “Strategic Learning and the Topology of Social Networks,” *Econometrica*, 83, 1755–1794.
- (Forthcoming): “Asymptotic learning on Bayesian social networks,” *Probability Theory and Related Fields*.
- MOSSEL, E. AND O. TAMUZ (2010): “Efficient Bayesian Learning in Social Networks with Gaussian Estimators,” *arXiv:1002.0747*.
- (2014): “Opinion exchange dynamics,” *arXiv preprint arXiv:1401.4770*.
- MUELLER-FRANK, M. (2014): “Does one Bayesian Make a Difference?” *Journal of Economic Theory*, 154, 423–452.
- MUELLER-FRANK, M. AND C. NERI (2013): “Social Learning in Networks: Theory and Experiments,” .
- OSBORNE, M. AND A. RUBINSTEIN (1994): *A Course in Game Theory*, MIT Press Books.
- RIVERS AND VUONG (2002): “Model Selection Tests for Nonlinear Dynamic Models,” *Econometrics Journal*, 5, 1–39.
- SADLER, E. D. (2014): “Bounded Rationality and the Network Economist’s Dilemma,” *Available at SSRN 2501690*.

TABLES

TABLE 1. Fraction of observations for which the complete information Bayesian model differs with the DeGroot alternative

Panel A: Networks from Choi et al. (2005, 2009)			
	[1]	[2]	[3]
Network	Uniform	Degree	Eigenvector
1	0.00%	0.00%	0.00%
2	0.00%	0.00%	6.80%
3	11.34%	22.68%	12.47%

Panel B: Networks Selected in This Paper			
	[1]	[2]	[3]
Network	Uniform	Degree	Eigenvector
1	6.95%	11.15%	6.95%
2	10.04%	10.39%	10.39%
3	9.63%	10.69%	10.69%

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). In Panel B, network 3 has identical paths for Bayesian and both uniform and degree DeGroot alternatives. We chose this network to pit Bayesian against eigenvector weighting.

TABLE 2. One-sided tests with the null of Bayesian learning against each alternative of the DeGroot model for the network-level analysis

H_0	H_1	[1]	[2]	[3]	[4]
		All networks	Network 1	Network 2	Network 3
Bayesian	Uniform	0.0000***	0.0000***	0.0000***	0.0001***
Bayesian	Eigenvector	0.0000***	0.0000***	0.0001***	0.0151**
Bayesian	Degree	0.0000***	0.0000***	0.0001***	0.0151**
Uniform	Degree	0.9999***	0.9999***	0.9482*	0.9999***
Uniform	Eigenvector	0.9999***	-	0.9458*	0.9999***
Degree	Eigenvector	0.0000***	0.0000***	-	-

Notes: The test statistic is the normalized difference in the divergence function of the null and the alternative model. We show the probability that the test statistic is less than zero, estimated via bootstrap with replacement at the session level.

TABLE 3. Incomplete information parameter estimates and corresponding divergences for the network-level analysis

Panel A: Incomplete information point estimates			
	[1]	[2]	[3]
	π^u	π^d	π^e
Point Estimate	0.00	0.00	0.00
Standard Error	(0.08)	(0.12)	(0.09)
Panel B: Incomplete information expected divergence			
	[1]	[2]	[3]
Divergence	0.1198	0.1413	0.1343
Standard Errors	(0.0096)	(0.0102)	(0.0099)
Panel C: Incomplete information point estimates, no common knowledge			
	[1]	[2]	[3]
	π^u	π^d	π^e
Point Estimate	0.00	0.00	0.00
Standard Errors	(0.2834)	(0.2620)	(0.3139)
Panel D: Incomplete information divergence, no common knowledge			
	[1]	[2]	[3]
Divergence	0.1798	0.1721	0.1831
Standard Errors	(0.0107)	(0.0106)	(0.0105)

Notes: Columns 1, 2 and 3 refer to the alternative models of uniform, degree and eigenvector centrality weighting, respectively. Standard errors are constructed by bootstrapping with replacement at the session level.

TABLE 4. One-sided tests with the null of Bayesian learning against each alternative of the DeGroot model for the individual-level analysis

H_0	H_1	[1]	[2]	[3]	[4]
		All networks	Network 1	Network 2	Network 3
Bayesian	Uniform	0.0000***	0.0000***	0.0000***	0.0000***
Bayesian	Eigenvector	0.0000***	0.0000***	0.0000***	0.0000***
Bayesian	Degree	0.0000***	0.0000***	0.0000***	0.0000***
Uniform	Degree	0.9999***	0.9999***	0.9757**	0.9999***
Uniform	Eigenvector	0.9999***	0.9999***	0.9999***	0.9999***
Degree	Eigenvector	0.8074	0.2460	0.9990***	0.6451

Notes: The test statistic is the normalized difference in the divergence function of the null and the alternative model. We show the probability that the test statistic is less than zero, estimated via bootstrap with replacement at the session level. When comparing DeGroot alternatives we use all observations, since there is no notion of zero probability information sets. When comparing complete information Bayesian to a DeGroot alternative, we only use histories on which a Bayesian agent has not reached a zero probability information set.

TABLE 5. Information dominance

H_1	[1]	[2]
	Observations	Percent Bayesian
Uniform	98	17.35%
Degree	46	17.39%
Eigenvector	98	17.35%

Notes: Observations are the number of cases where there are discrepancies between the parent node action (which, from $t > 3$ has the property that under the Bayesian model the peripheral node should follow the parent node) and the action that the alternative model prescribes for the peripheral node. The percent Bayesian presents the share of observations in which the peripheral nodes followed the Bayesian node action.

TABLE 6. Cases that encounter zero probability information sets

Panel A: Complete Information			
	[1]	[2]	[3]
Network	% Individuals	% Sessions	% Observations
1	47.97%	72.63%	43.90%
2	54.29%	73.68%	46.15%
3	33.06%	64.29%	34.66%

Panel B: Uniform Weighting Alternative			
	[1]	[2]	[3]
Network	% Individuals	% Sessions	% Observations
1	35.79%	72.63%	37.29%
2	37.14%	68.42%	39.85%
3	26.94%	64.29%	31.45%

Panel C: Degree Weighting Alternative			
	[1]	[2]	[3]
Network	% Individuals	% Sessions	% Observations
1	34.14%	67.37%	37.02%
2	38.05%	70.53%	40.70%
3	25.92%	62.86%	30.89%

Panel D: Eigenvector Weighting Alternative			
	[1]	[2]	[3]
Network	% Individuals	% Sessions	% Observations
1	36.09%	72.63%	38.00%
2	38.05%	70.53%	40.70%
3	25.92%	62.86%	30.89%

Notes: All percentages computed using the π that minimizes the share of data units (individuals, sessions or observations) that reach zero probability information sets.

TABLE 7. Incomplete information parameter estimates and corresponding divergences for the individual-level analysis on entire support

Panel A: Incomplete information point estimates			
	[1]	[2]	[3]
	π^u	π^d	π^c
Point Estimate	0.00	0.00	0.00
Standard Error	(0.0140)	(0.0243)	(0.0109)

Panel B: Incomplete information expected divergence			
	[1]	[2]	[3]
Divergence	0.0648	0.1029	0.1097
Standard Errors	(0.0064)	(0.0075)	(0.0061)

Panel C: Incomplete information point estimates, no prior			
	[1]	[2]	[3]
	π^u	π^d	π^c
Point Estimate	[0.4,0.5]	0.00	0.00
Confidence Interval (95%)	[0,0.8]	[0,0.6]	[0,0.8]

Panel D: Incomplete information divergence, no common prior			
	[1]	[2]	[3]
Divergence	0.2656	0.2554	0.2672
Standard Errors	(0.0142)	(0.0154)	(0.0142)

Notes: Columns 1, 2 and 3 refer to the alternative models of uniform, degree and eigenvector centrality weighting, respectively. Standard errors are constructed by bootstrapping with replacement at the session level over all histories where zero probability information sets are penalized. We present confidence intervals when the parameter is set identified instead of point identified.

TABLE 8. Incomplete information parameter estimates and corresponding divergences for the individual-level analysis on Bayes-feasible histories

Panel A: Incomplete information point estimates			
	[1]	[2]	[3]
	π^u	π^d	π^c
Point Estimate	0.00	0.00	0.90
Standard Error	(0.3214)	(0.2765)	(0.3102)

Panel B: Incomplete information expected divergence			
	[1]	[2]	[3]
Divergence	0.0685	0.1068	0.1177
Standard Errors	(0.0155)	(0.0145)	(0.0151)

Panel C: Incomplete information point estimates, no common prior			
	[1]	[2]	[3]
	π^u	π^d	π^c
Point Estimate	[0.9,1]	0.90	0.90
Confidence Interval (95%)	[0,1]	[0,1]	[0,1]

Panel D: Incomplete information divergence, no common prior			
	[1]	[2]	[3]
Divergence	0.2750	0.2762	0.2769
Standard Errors	(0.0148)	(0.0148)	(0.0152)

Notes: Columns 1, 2 and 3 refer to the alternative models of uniform, degree and eigenvector centrality weighting, respectively. Standard errors are constructed by bootstrapping with replacement at the session level over all histories where any Bayesian agent has not yet reached a zero probability information set. Therefore, zero probability information sets are not penalized. We present confidence intervals when the parameter is set identified instead of point identified.

TABLE 9. Overweighting information received through multiple channels

	Data		Bayesian		Degree		Uniform	
	All	Restricted	All	Restricted	All	Restricted	All	Restricted
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Signal	.4974***	.5351***	.3965***	.4429***	.5725***	.6204***	.6017***	.677***
	[.0218]	[.0279]	[.017]	[.0223]	[.0169]	[.0211]	[.0159]	[.0201]
Direct	.7735***	.7485***	.775***	.8048***	.9125***	.8913***	.8954***	.9078***
	[.0311]	[.0438]	[.0323]	[.0384]	[.0335]	[.0425]	[.0381]	[.0475]
One Way	.3157***	.2609***	.7056***	.8189***	0.0461	-0.0614	.1306***	.077*
	[.0405]	[.0542]	[.0363]	[.0486]	[.035]	[.0421]	[.0368]	[.0447]
Two Ways	.2177***	.1751***	.3582***	.4096***	.0797*	0.024	.1985***	.1841***
	[.0355]	[.0424]	[.0337]	[.0338]	[.0421]	[.0449]	[.0459]	[.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
N	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: Robust standard errors, clustered at the village by game level, in brackets. 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. Column (1) is the regression with all data. Column (2) is the regression restricting to treatments that are informative for the comparisons Bayesian - Degree and Bayesian - Uniform. * p<.1, ** p<.05, *** p<.01

TABLE 10. Re-incorporation of previously-observed historical information

	Data		Bayesian		Degree		Uniform	
	All	Restricted	All	Restricted	All	Restricted	All	Restricted
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[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]
N	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: Robust standard errors, clustered at the village by game level, in brackets. 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. Column (1) is the regression with all data. Column (2) is the regression restricting to treatments that are informative for the comparisons Bayesian - Degree and Bayesian - Uniform. * p<.1, ** p<.05, *** p<.01

APPENDIX A. BAYESIAN LEARNING ALGORITHM IN COMPLETE AND INCOMPLETE
INFORMATION MODELS

In this appendix, we describe the algorithm for computing the actions under the assumption of complete information Bayesian agents.

A.1. Setup. We follow the notation on [Osborne and Rubinstein \(1994\)](#) and [Geanakoplos \(1994\)](#), modeling agents' information in the experiment by means of *dynamically consistent models of action and knowledge* (DCMAK), a natural multi-period generalization of [Aumann \(1976\)](#). Following [Geanakoplos \(1994\)](#), a DCMAK consists of a set of *states of the world* $\omega \in \Omega$, information functions $P_{i,t} : \Omega \rightarrow 2^\Omega$, and action functions $a_{i,t} : \Omega \rightarrow \{0, 1\}$. In what follows, we will define these objects for our experimental setup, which we will use to calculate the predicted behavior of Bayesian agents $a_{i,t}(\omega)$.

A.2. States of the world. In both the complete and incomplete information models, we model agents information as partitions over $\omega \in \Omega$, where $\omega = (\omega_1, \omega_2, \dots, \omega_n)$ is the vector of agents' initial private information. In the *complete information case*, $\omega_i \in \{0, 1\}$ denotes whether they observe a yellow ball ($\omega_i = 1$) or blue ($\omega_i = 0$); i.e., $\Omega = S = \{0, 1\}^n$. In the *incomplete information case* (where players don't know if their neighbors are Bayesian or DeGroot types) we model the state of the world as $\omega_i = (s_i, \eta_i)$ where $s_i \in \{0, 1\}$ is the color of the observed ball, and $\eta_i \in \{0, 1\}$ denotes agent *i's type*: she is either a Bayesian type ($\eta_i = 1$) who guesses the most likely state following Bayes' rule, or a DeGroot agent ($\eta_i = 0$) who decides her guess based on a weighted average of her neighbors' guesses (we consider the 3 variants mentioned: uniform, degree and eigenvector centrality weighting).

On both complete and incomplete information models, Bayesian agents have a common prior belief over states $\omega \in \Omega$, conditional on the realization of $\theta \in \{0, 1\}$ (i.e., which bag has been chosen), which we denote by $\rho(\omega | \theta)$. In the complete information model $\omega = s$ and all agents are of the Bayesian type $\eta_i = 1$ at all states, and hence

$$(A.1) \quad \rho(\omega | \theta) = \rho(s | \theta) := p_\theta^{\sum_j s_j} (1 - p_\theta)^{n - \sum_j s_j} \text{ for } \theta \in \{0, 1\}, \omega \in \Omega,$$

where $p_\theta = P(s_i = 1 | \theta)$. In our experiment, $p_\theta = 5/7$ if $\theta = 1$ and $p_\theta = 2/7$ if $\theta = 0$. In the incomplete information model, $\omega = (s, \eta)$ where $\eta = (\eta_1, \eta_2, \dots, \eta_m)$ is a type profile (the complete description of every agent's type), and we assume that agents have a common prior over agents' types as well, which are assumed to be i.i.d. across agents and independent of observed signals;

$$(A.2) \quad \rho(s, \eta | \theta) := p_\theta^{\sum_j s_j} (1 - p_\theta)^{n - \sum_j s_j} \left[\pi^{\sum_j \eta_j} (1 - \pi)^{n - \sum_j \eta_j} \right],$$

where $\pi := P(\eta_i = 1)$. The set of all type configurations is denoted by $H = \{0, 1\}^n$, and in this model, $\Omega := S \times H = \{0, 1\}^n \times \{0, 1\}^n$.

A.3. Recursive definition of information and action functions. The function $P_{i,t}(\omega) \subseteq \Omega$ denotes the information set of agent i at round t , under state ω . At round $t = 1$, agent i only observes ω_i out of state ω , and hence, her information set is:

$$(A.3) \quad P_{i,1}(\omega) := \{\omega' \in \Omega : \omega'_i = \omega_i\},$$

i.e., the possible states of the world are those compatible with the private information she has received (which includes her signal $s_i \in \{0, 1\}$ and her type). Based on this information, all agents choose to match their signal; i.e.,

$$(A.4) \quad a_{i,1}(\omega) := s_i.$$

For $t > 1$ we compute $P_{i,t}(\omega)$ and $\mathbf{a}_{i,t}(\omega)$ inductively, for each $\omega \in \Omega$. In our experimental setup, at round t agent i observes all the actions taken by her neighbors $j \in N(i)$ (including herself) up to $s = t - 1$. Therefore, the states of the world that are consistent with agent i 's observations (her information set) are

$$(A.5) \quad P_{i,t}(\omega) := \{\omega' \in \Omega : \omega'_i = \omega_i \text{ and } a_{j,s}(\omega') = a_{j,s}(\omega) \text{ for all } j \in N(i), s \leq t - 1\}.$$

Clearly, we have $P_{i,t}(\omega) \subseteq P_{i,t-1}(\omega)$ for all $i, \omega \in \Omega$ (i.e., $P_{i,t}(\cdot)$ corresponds to a filtration).²⁴ The round t action function $a_{i,t}(\omega)$ is then given by:

$$(A.6) \quad a_{i,t}(\omega) := \begin{cases} \mathbf{1} \{I_{i,t}(\omega) > \frac{1}{2}\} & \text{if } I_{i,t}(\omega) \neq \frac{1}{2} \\ a_{i,t-1}(\omega) & \text{if } I_{i,t}(\omega) = \frac{1}{2} \end{cases},$$

where $I_{i,t}(\omega)$ is the ‘‘belief index’’ at state ω , which depends on the agents’ type. If agent i is Bayesian (i.e., under the complete information model, or if $\tau_i = 1$ in the incomplete one), then $I_{i,t}(\omega) := \mathbb{P}(\theta = 1 \mid P_{i,t}(\omega))$, which is calculated using Bayes rule conditioning on the event $B = P_{i,t}(\omega)$:

$$(A.7) \quad \mathbb{P}(\theta = 1 \mid P_{i,t}(\omega)) := \frac{\sum_{\omega' \in P_{i,t}(\omega)} \rho(\omega' \mid \theta = 1)}{\sum_{\omega' \in P_{i,t}(\omega)} \rho(\omega' \mid \theta = 1) + \rho(\omega' \mid \theta = 0)}$$

When i is not Bayesian at ω , then $I_{i,t}(\omega) := \sum_{j=1}^n T_{ij} a_{j,t-1}(\omega)$ where $[T_{ij}]_{ij}$ are the DeG-root weights (which are T^u, T^d or T^e as in equations (2.1), (2.2), and (2.3), depending on the assumed *alternative* type).

A.4. Numerical Implementation. The algorithm used is based on the inductive step defined above, calculating iteratively the objects $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$ for all i, t and ω , for both the complete and incomplete information models.

Algorithm 1. *Bayesian Learning Algorithm (Complete Information case)*

Input: An n -person network $G = (V, E)$ with adjacency matrix $A_{n \times n}$

²⁴We can also define $P_{i,t}$ recursively, starting at $P_{i,1}$ as in A.3, and for $t \geq 1$ let $P_{i,t}(\omega) := P_{i,t-1}(\omega) \cap \{\omega' \in \Omega : a_{j,t-1}(\omega') = a_{j,t-1}(\omega) \text{ for all } j \in N(i)\}$

Outputs: Information and action functions $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$

Step 1: Initialize algorithm by defining:

- (1) State space $\Omega = \{0, 1\}^n$
- (2) Measures $\rho(\omega | \theta)$ according to A.2, for $\theta \in \{0, 1\}$
- (3) Information functions $P_{i,1}(\omega)$ and actions $a_{i,1}(\omega)$ according to A.3 and A.4 for all $i = 1, \dots, n$ and $\omega \in \Omega$.

Step $t > 1$: Given $(P_{i,s}(\omega), a_{i,s}(\omega))_{i=1, \dots, n, s=1, \dots, t-1, \omega \in \Omega}$ calculate $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$ for all i and $\omega \in \Omega$ according to A.5 and A.6, where $I_{i,t}(\omega) = P(\theta = 1 | P_{i,t}(\omega))$ as in equation (A.7).

Algorithm 2. Bayesian Learning Algorithm (Incomplete Information case)

Inputs:

- (1) An n -person network $G = (V, E)$ with adjacency matrix $A_{n \times n}$
- (2) A row stochastic matrix of DeGroot weights $T_{n \times n}$
- (3) Probability $\pi \in [0, 1]$

Output: Information and action functions $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$

Step 1: Initialize algorithm by defining:

- (1) State space $\Omega = S \times H = \{\omega = (s, \eta) \text{ where } s \in S := \{0, 1\}^n, \eta \in H := \{0, 1\}^n\}$
- (2) Measures $\rho(\omega | \theta) = \rho(s, \eta | \theta)$ according to A.2, for $\theta \in \{0, 1\}$
- (3) Information functions $P_{i,1}(\omega)$ and actions $a_{i,1}(\omega)$ according to A.3 and A.4 for all $i = 1, \dots, n$ and $\omega \in \Omega$.

Step $t > 1$: Given $(P_{i,s}(\omega), a_{i,s}(\omega))_{i=1, \dots, n, s=1, \dots, t-1, \omega \in \Omega}$ calculate $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$ for all i and $\omega \in \Omega$ according to A.5 and A.6, where $I_{i,t}(\omega) = P(\theta = 1 | P_{i,t}(\omega))$ if $\eta_i = 1$ and $I_{i,t}(\omega) = \sum_j T_{ij} a_{j,t}(\omega)$ if $\eta_i = 0$.

It is worth noting that an alternative way of modeling the knowledge structure is by including the true state θ in the description of the state of the world; i.e., define $\omega = (\theta, s)$ in the complete information case, and $\omega = (\theta, s, \eta)$ in the incomplete information case, which would need the definition of just one common prior $\rho(\omega)$, instead of having to define it conditional on θ . While this would perhaps be a better fit for most epistemic models, the description of the algorithm is slightly easier in our model, given the fact that $\omega = s$ in the complete information model, and $\omega = (s, \eta)$ in the incomplete information models are, respectively, sufficient statistics for the actions sequence of players, since θ is never in any information set of any of the players, significantly reducing the relevant state space (in fact, these are the minimal state spaces we can consider, exactly because of sufficiency).

A.5. Complexity. 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. This means that the computational problem is at least as hard as NP-complete problems.²⁵ 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 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 $N \cdot T$. 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 associated learning network is growing exponentially 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.

Let Ω_t be the the set of states that agent i has to integrate over at time t . The basic algorithm (in this general version) involves two states: the indicator function of the set $P_i(\omega)$ for each $\omega \in \Omega_t$ and the action function $a_{i,t}(\omega)$. We define

$$\sigma_t(i, \omega, \omega') := \begin{cases} 1 & \text{if } \omega' \in P_{i,t}(\omega) \\ 0 & \text{otherwise} \end{cases}$$

and

$$\alpha_t(i, \omega, t) := a_{i,t}(\omega)$$

to calculate the objects $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$ numerically, as in appendix A.4. To calculate them, we then have to loop across $\#(\Omega_t) \times \#(\Omega_t)$ states for each (i, t) to update σ_t to σ_{t+1} and $\#(\Omega_t)$ to update α_t . The number of operations is then $\sum_t \sum_i (\sum_{w \in W_t} (k + \sum_{\hat{w} \in W_t} k))$ where k is the number of computations done in each step. In the deterministic complete information model (without trembles), $\Omega_t = S = \{0, 1\}^n$ and then

$$\text{Computations} = nT (2^n) (1 + 2^n) k = \Theta(T4^n).$$

²⁵A problem is said to be NP-complete if (a) it is in 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.

In the incomplete information model, it is $nT(4^n)(1+4^n)k$. The ratio between the complete and incomplete information models is then

$$\frac{nT(4^n)(1+4^n)k}{nT(2^n)(1+2^n)k} = 2^n \frac{1+4^n}{1+2^n} \approx 4^n$$

So, for a network of $n = 7$, the relative complexity of the incomplete information model approximately 16,258.

The trembling hand complete information model needs agents to integrate over $2^{n(t-1)}$ states, at least, in each round; since there is no longer a deterministic mapping between information sets and signal profiles, agent i needs to integrate over the actions of other agents. Although agent i actually does not observe the information of $n - d_i$ agents, for rounds $t \geq 3$ we have to have to calculate her beliefs about those agents' information sets. The partitional model presented in appendix A.4 does not suffer this problem, by calculating beliefs on all states, which we do here as well. Therefore $\#(W_t) = 2^{n(t-1)}$ and the number of calculations is

$$k \sum_i \sum_{t=1}^{t-1} 2^{n(t-1)} (k + 2^{n(t-1)}) = k \sum_i \left(\sum_{t=1}^{T-1} 2^{n(t-1)} + \sum_{t=1}^{T-1} 2^{2n(t-1)} \right)$$

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

Therefore, the ratio between the two is approximately

$$\frac{n2^{n(T-1)} [k + 2^{n(T-1)}]}{nT(2^n)(k + 2^n)} \approx \frac{1}{T} \times 4^{n(T-2)},$$

and for the incomplete information model, the equivalent ratio is

$$\frac{n4^{n(T-1)} (1 + 4^{n(T-1)})}{nT(4^n)(1 + 4^n)}.$$

ONLINE APPENDICES

APPENDIX B. ONLINE APPENDIX: FILTERING

Here we describe a filter to estimate, in the incomplete information model, the probability that an agent is Bayesian given the data that the researcher has in hand, under the network level approach. The objective of this section is to estimate

$$F_i := \mathbb{P}(i \text{ is Bayesian} \mid \pi, \text{ experimental data in room } r).$$

To simplify the exposition, we drop the subscripts for the particular room r , the alternative model used, m and the prior probability of players being Bayesian, $\pi \in [0, 1]$. Let $[k] = \{1, 2, \dots, k\}$ for any $k \in \mathbb{N}$, $T = T_r$ be the number of rounds played, and $\{a_{i,t}^{obs}\}_{i \in [n], t \in [T]}$ denote the (panel) data set of player choices, which include the signals drawn at $t = 1$ (since $s_i^{obs} \equiv a_{i,1}^{obs}$). Given a type profile $\eta \in H$, define the observation level divergence $D_{i,t}(\eta) := |a_{i,t}(s^{obs}, \eta) - a_{i,t}^{obs}|$ for all $i \in [n]$ and $t > 1$. The room level divergence of type profile η is defined as $D(\eta) := \frac{1}{n(T-1)} \sum_{i,t > 1} D_{i,t}(\eta)$.

In the network level approach, we treat the data generating process as follows: a state $\omega \in \Omega = S \times H$ is drawn at $t = 1$, and agents play according to the equilibrium path sequence $a_{i,t}(\omega)$. The researcher knows the actual signal endowment $s = s^{obs}$ (i.e., the color of the balls drawn in the experiment), and the panel of actions chosen by all agents, which we denote by $\{a_{i,t}^{obs}\}_{i \in [n], t \in [T]}$. Formally, the information set of the researcher, $P^R \subseteq \Omega$, restricts the set of states to those consistent with the observed data; i.e. $P^R := \{\omega = (s, \eta) \in S \times H : a_{i,t}(\omega) = a_{i,t}^{obs} \text{ for all } i \in [n], t \leq T\}$. Therefore, the only remaining uncertainty is the type profile $\eta \in H$, and the researcher information can be better described as restrictions over possible type profiles, denoted by

$$H^R := \text{proj}_H P^R = \left\{ \eta \in H : a_{i,t}(s^{obs}, \eta) = a_{i,t}^{obs} \text{ for all } i \in [n], t \in [T] \right\}.$$

In this appendix, we are particularly interested in calculating the posterior probability that agent i is Bayesian, given a prior probability π ,

$$F_i(\pi) := \mathbb{P}(\eta_i = 1 \mid \{a_{i,t}^{obs}\}, \pi).$$

If there exists some type profile η^* that is consistent with the observed data (i.e. $H^R \neq \emptyset$), then this can be easily calculated using Bayes' rule:

$$\begin{aligned} F_i(\pi) &= \mathbb{P}(\eta_i = 1 \mid \eta \in H^R) = \frac{\sum_{\eta \in H^R: \eta_i=1} \mathbb{P}(\eta \mid s^{obs})}{\sum_{\eta \in H^R} \mathbb{P}(\eta \mid s^{obs})} \\ &= \frac{\sum_{\eta \in H^R: \eta_i=1} (\pi/1 - \pi)^{n\bar{\eta}}}{\sum_{\eta \in H^R} (\pi/1 - \pi)^{n\bar{\eta}}}. \end{aligned}$$

Since $P(\eta) = \pi^{\sum_j \eta_j} (1 - \pi)^{n - \sum_j \eta_j} = (1 - \pi)^n \left(\frac{\pi}{1 - \pi}\right)^{n\bar{\eta}}$, where $\bar{\eta} := \frac{1}{n} \sum_j \eta_j$. While tractable, this exercise is impossible in most of our experiments: as Table 6 shows, in at least 63% of experiments, no type profile $\eta^* \in H$ exist that could rationalize the observed data (i.e $H^R = \emptyset$) regardless of network topology chosen and the alternative DeGroot action type for the incomplete information model.

If we maintain the network level approach interpretation, an alternative is to add “trembles” to ensure full support over possible action sequences $\{a_{i,t}\}_{i,t}$. Formally, let $\mathbf{a}_{i,t} \in \{0, 1\}$ be the (random) choice of agent i at round t . The key object we need to model is the probability of observing the dataset $\{a_{i,t}^{obs}\}_{i,t}$ if the state is $\omega = (s^{obs}, \eta)$

$$\mathcal{A}(\eta) := P\left(\mathbf{a}_{i,t} = a_{i,t}^{obs} \text{ for all } i \in [n], t \in [T] \mid \omega = (s^{obs}, \eta)\right)$$

so $\mathcal{A}(\eta)$ is the likelihood function of type profile η (treated as a parameter) given the observed data. Given $\mathcal{A}(\eta)$, we can calculate F_i by first getting the posterior distribution over type profiles, using Bayes’ rule

$$P\left(\eta \mid \{a_{i,t}^{obs}\}_{i,t}\right) = \frac{P(\eta) \mathcal{A}(\eta)}{\sum_{\eta \in H} P(\eta) \mathcal{A}(\eta)}$$

and then calculating the probability of $\eta_i = 1$ under it,

$$\begin{aligned} F_i &:= P\left(\eta_i = 1 \mid \{a_{i,t}^{obs}\}_{i,t}\right) = \sum_{\eta: \eta_i=1} P\left(\eta \mid \{a_{i,t}^{obs}\}_{i,t}\right) \\ &= \frac{\sum_{\eta: \eta_i=1} (\pi/1 - \pi)^{n\bar{\eta}} \mathcal{A}(\eta)}{\sum_{\eta} (\pi/1 - \pi)^{n\bar{\eta}} \mathcal{A}(\eta)}. \end{aligned}$$

Our leading example of a model of trembles will be the *independent trembles model*, parametrized by a constant $\epsilon \in (0, 1)$. Formally, given a type profile $\eta \in H$, we assume **(1)** $\mathbf{a}_{i,1} = s_i^{obs}$ with probability 1 for all i **(2)** $\mathbf{a}_{i,t} = a_{i,t}(s^{obs}, \eta)$ with probability $(1 - \epsilon)$, which is i.i.d across players and rounds. Formally,

$$\mathcal{A}(\eta) := \epsilon^{\sum_{i,t>1} D_{i,t}(\eta)} (1 - \epsilon)^{n(T-1) - \sum_{i,t>1} D_{i,t}(\eta)},$$

where $D_{i,t}(\eta) := |a_{i,t}^{obs} - a_{i,t}(s^{obs}, \eta)|$ is the theoretical observation divergence under state $\omega = (s^{obs}, \eta)$. It can be expressed in a more familiar way as

$$\mathcal{A}(\eta) = \left[\epsilon^{D(\eta)} (1 - \epsilon)^{1 - D(\eta)}\right]^{n(T-1)} = (1 - \epsilon)^{n(T-1)} \left(\frac{\epsilon}{1 - \epsilon}\right)^{n(T-1)D(\eta)}.$$

A nice feature is that the posterior distribution over type profiles is directly related to its measured divergence:

$$P\left(\eta \mid \{a_{i,t}^{obs}\}_{i,t}\right) \propto P(\eta) \times \left(\frac{\epsilon}{1 - \epsilon}\right)^{n(T-1)D(\eta)},$$

and hence,

$$(B.1) \quad F_i = \frac{1}{\Lambda(\pi, \epsilon)} \sum_{\eta: \eta_i=1} \left(\frac{\pi}{1-\pi} \right)^{n\bar{\eta}} \left(\frac{\epsilon}{1-\epsilon} \right)^{n(T-1)D(\eta)},$$

where $\Lambda(\pi, \epsilon) := \sum_{\eta} (\pi/1-\pi)^{n\bar{\eta}} (\epsilon/1-\epsilon)^{n(T-1)D(\eta)}$ is proportional to the probability of the researcher's information set over type profiles H^R ²⁶. One may observe that type profiles that do not explain the data well (if $D(\eta)$ is high) are given a lesser weight in the posterior distribution if $\epsilon < \frac{1}{2}$, so that the filtered probability of i being Bayesian depends crucially on how much is this evidenced by the data.

APPENDIX C. ONLINE APPENDIX: STUCK NODES AND COMPLETE BINARY QUILTS

C.1. Setup. We present a simple growing sequence of networks that illustrates the results of Proposition 2.1. In this example, there is asymptotic efficiency under communication DeGroot models and action Bayesian models, but asymptotic learning fails and consensus does not occur with action DeGroot models

We say a network $T = (V, E)$ is a *binary tree* if it is a rooted tree where every node has either two children, or none (i.e., a *terminal node*).²⁷ The *level* of a node $j \in V$ is the shortest path distance to the root node i_0 . The set of nodes of level d is denoted by L_d .²⁸

Define the depth of a network d to be the maximum distance possible from the root node. A *complete binary tree of depth d* is a binary tree $T_d = (\hat{V}_d, \hat{E}_d)$ where every terminal node has depth d . Note that in this network, every node (other than the root) has 1 sibling (a node that has the same parent) to which it is not connected.

Given T_d , we define a *complete binary quilt of depth d* to be a network $Q_d = (V_d, E_d)$, super-graph of T_d , where $V_d = \hat{V}_d$ and $E_d = \hat{E}_d \cup \{ij : i \text{ and } j \text{ are siblings}\}$, as illustrated in Figure C.1.




FIGURE C.1. Complete binary tree T_d and complete binary quilt Q_d ($d = 3$)

²⁶Namely, $P(H^R) = \Lambda(\pi, \epsilon) \times [(1-\pi)\epsilon^{T-1}]^n$.

²⁷A terminal node, or leaf, is a node in a tree with only one neighbor (i.e., it has degree 1).

²⁸Formally, $L_d := \{j \in V : d(i_0, j) = d\}$ where $d(i, j)$ is the length of the shortest path between i and j (which exists, given that T is a connected graph).

The motivation for the graph structure comes from Jackson et al. (2012), who study network architecture that arise as equilibria in favor-exchange games. They show that these networks are social quilts. A social quilt is a patchwork of substructures (e.g., triangles) pasted together in specific ways: the substructures do not share edges, only vertices. From the applied perspective, this speaks to how the incentives to overcome contractual incompleteness, or lack of formal insurance, influence how well a society may engage in social learning. If graphs are constructed as equilibria of risk-sharing or favor-exchange games, then they may have such quilt-like substructures. However, at the same time, because of fixed costs in building relationships, the resulting pattern of relationships will also be the network on which information travels.




FIGURE C.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 correct signal. In Panel B, in the first period, 4 nodes receive the correct signal, and after one node switches, 3 are stuck.

This family of networks is prone to having stuck nodes, just as we saw when studying topology 3, which is a particular sub-case. Panel A of Figure C.2 illustrates the problem. Assume that for some subtree of level $d - 1$ of Q_d (i.e., where the siblings are terminal nodes), which connects to the rest of the network through its parent node (as in Figure C.2), we have the initial signal endowment shown in Panel A of Figure C.2. 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 periods, since they follow the DeGroot uniform weighting rule as in equation 2.1. Note that in large enough networks, taking an action that goes against the majority will be almost surely be the same as taking the wrong action (since almost surely, most agents will have observed the right signal as $p > 1/2$). We will use this equivalent definition of stuckness for large networks in this application (i.e., a node i is stuck if $\exists T_i : a_{i,t} = 1 - \theta$ for all $t \geq T_i$)

Take a sequence of networks $\{Q_d\}_{d \in \mathbb{N}}$ of growing complete binary quilts. As $d \rightarrow \infty$, there will be a non-vanishing fraction of subtrees with this initial configuration of signals. Following the argument of Panel A of Figure C.2, take the subgraph induced by levels $d - 2, d - 1$ and d of Q_d ; i.e., subgraph Q' induced by $V' = L_{d-2} \cup L_{d-1} \cup L_d$. It consists of the union of 2^{d-2} connected components of 7 players, like the ones in panel A and B of Figure C.2. If d is big enough for the law of large numbers to kick in, there will be approximately a fraction of $p^5 (1 - p)^2$ of those components with signal configurations as in Panel A, and each of these will have at least 2 players getting stuck choosing the wrong action. This already gives a non-zero lower bound on the fraction of agents that will never match their action to the true state of nature, even if the network is arbitrarily big. In this case, the fraction of agents that get stuck in components like Panel A is at least $f = \frac{1}{4}p^5 (1 - p)^2$ as $d \rightarrow \infty$.²⁹

This example has demonstrated the following result. We say that a sequence of networks exhibits asymptotic learning if all but a vanishing share of nodes correctly guess the state of the world.

That asymptotic learning occurs with the Bayesian action model follows from Mossel et al. (Forthcoming), and that it occurs with DeGroot communication for this model follows from Corollary 1 of Golub and Jackson (2010). The result for the DeGroot action model is apparent from the previous example. To illustrate the severity of Proposition 2.1, in Figure C.3 we show lower bounds on the expected fraction of nodes that are stuck. These bounds are calculated by the argument following Figure C.2, replicating the bounds obtained for the signal configurations of Panel A and B for *all* possible signal configurations of each component. Even with high quality signals ($p = 0.7$), at least 16% of nodes become stuck and do not asymptotically learn.

C.2. Estimating Upper and Lower Bounds on the Fraction Stuck.

²⁹To see this, one can show that (1) $\# \{V\} = 2^{d+1} - 1$; (2) there are 2^{d-2} components in V' , each of 7 nodes; and (3) at least 2 players out of Panel A-type components get stuck. Then, if d is large enough, we have that $f \approx 2 \times p^5 (1 - p^2) 2^{d-2} / (2^{d+1} - 1) = \frac{1}{4}p^5 (1 - p)^2 \times (2^{d+1} / 2^{d+1} - 1)$. Since the second term converges to 1 as $d \rightarrow \infty$, we get the desired expression.

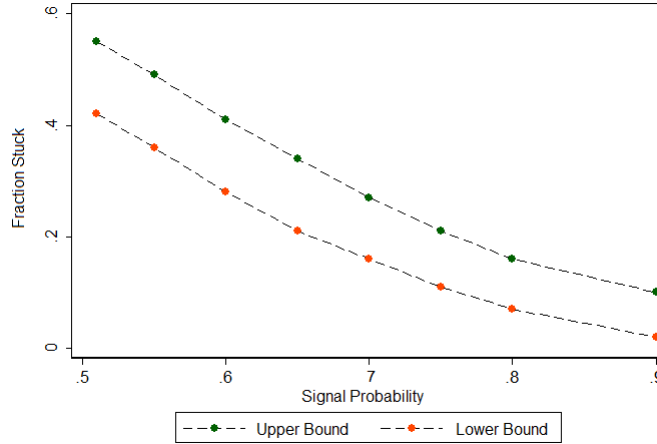



FIGURE C.3. Bounds for fraction stuck


 FIGURE C.4. Subgraphs $O_d(k)$ of Q_d ($d = 3$).

C.2.1. *Preliminary tools.* We define $S_d = \{i \in V_d : i \text{ gets stuck}\}$ and let $N_d = \#(V_d)$, the number of nodes in Q_d . Our object of interest is the random variable

$$\mathcal{F}_d = \text{Fraction of nodes in } Q_d \text{ that gets stuck} \equiv \frac{\#S_d}{N_d}.$$

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

$$(C.1) \quad \underline{\mathcal{F}} = \liminf_{d \rightarrow \infty} \mathcal{F}_d \text{ and } \overline{\mathcal{F}} = \limsup_{d \rightarrow \infty} \mathcal{F}_d$$

which is well defined for all realizations of the sequence \mathcal{F}_d , 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 $\underline{\mathcal{F}} \geq F$ and $\overline{\mathcal{F}} \leq F$ almost surely; i.e., $\mathbb{P}\{\underline{\mathcal{F}} \leq F \leq \overline{\mathcal{F}}\} = 1$.

Define

$$O_d(k) := \bigcup_{s=0}^{s=k} L_{d-s}$$

as the subgraph formed by the last $k + 1$ levels of Q_d . This subgraph is disconnected, with a lot of connected components, as pictured in Figure C.4. Let C be a generic connected component of $O_d(k)$; i.e., the induced subgraph of one of the connected components of $O_d(k)$. For example, Panels A and B of Figure C.2 correspond to connected components of $O_d(2)$. Notice that in this example, each component of $O_d(2)$ is itself a complete binary quilt of level 2. It is easy to see that in fact, the connected components of $O_d(k)$ are *also complete binary quilts*; i.e., $C = \hat{Q}_k$ (binary quilts of depth $d' = k$). The key property of these subgraphs (and the reason for studying this family of networks) 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 . Next, we see how this property allows us to study the local behavior of the component, abstracting away from the rest of the network, in order to get bounds on the fraction of players stuck.

Define

$$\Psi_d(k) = \frac{\#\{O_d(k) \cap S_d\}}{\#\{O_d(k)\}}$$

as the fraction of stuck nodes in $O_d(k)$,

$$\underline{\Psi}(k) := \liminf_{d \rightarrow \infty} \frac{\#\{O_d(k) \cap S_d\}}{\#\{O_d(k)\}},$$

and

$$\bar{\Psi}(k) := \limsup_{r \rightarrow \infty} \frac{\#\{O_r(k) \cap S_r\}}{\#\{O_r(k)\}},$$

which are also all well-defined random variables. These are the tightest asymptotic lower and upper bounds on the fraction of nodes in $O_d(k)$ that get stuck.

LEMMA C.1. *For all $k \in \mathbb{N}$, $\lim_{d \rightarrow \infty} \frac{\#\{O_d(k)\}}{N_d} = \frac{2^{k+1}-1}{2^{k+1}}$.*

Proof. Let $L_d =$ number of nodes in level r . Because of Q_d being a complete binary network, we have the following recursion for L_d :

$$L_{d+1} = 2L_d \text{ and } L_0 = 1$$

from which it follows that

$$(C.2) \quad L_d = 2^d.$$

Therefore, the number of nodes in $\#\{O_d(k)\}$ is then

$$\#\{O_d(k)\} = \sum_{s=0}^k \#\{L_{d-s}\} = 2^{d-k} \sum_{s=0}^k 2^s = 2^{d-k} \left(\frac{1-2^{k+1}}{1-2} \right) = 2^d \left(\frac{2^{k+1}-1}{2^k} \right).$$

We also need to calculate $N_d := \#(V_d)$. Again, because of how Q_d grows, it can be also easily shown that

$$(C.3) \quad N_d = 2^{d+1} - 1.$$

Now, we can state the result. Observe that

$$\frac{\#\{O_d(k)\}}{N_d} = \frac{2^d}{2^{d+1}-1} \left(\frac{2^{k+1}-1}{2^k} \right) \xrightarrow{d \rightarrow \infty} \frac{2^{k+1}-1}{2^k}.$$

as we wanted to show. \square

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

PROPOSITION C.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,

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

and

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

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

$$\begin{aligned} \mathcal{F}_d &= \frac{\#\{O_d(k)\}}{N_d} \left(\frac{\#\{O_d(k) \cap S_d\}}{\#\{O_d(k)\}} \right) + \frac{\#\{S_d - O_d(k)\}}{N_d} \\ &\geq \frac{\#\{O_d(k)\}}{N_d} \left(\frac{\#\{O_d(k) \cap S_d\}}{\#\{O_d(k)\}} \right), \end{aligned}$$

so, for all realizations,

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

This, together with the fact that $\underline{\Psi}(k) \geq \underline{\psi}(k)$ almost surely, completes 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 $\overline{\psi}$. See that condition C.4 and C.5 are bounds on $\underline{\mathcal{F}}$ and $\overline{\mathcal{F}}$, which do not depend on k . Therefore, these are bounds for all k . Moreover, the higher k , the tighter the bound.

C.2.2. Bounding the fraction stuck under uniform DeGroot. Without loss of generality, we assume that the true state of nature is $\theta = 1$, which implies that as $d \rightarrow \infty$ the fraction of nodes with true signals is $p > \frac{1}{2}$. The idea is simple: take a component $C = (V_C, E_C) \subset O_d(k)$. 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 C.4). Let $S_C = \{0, 1\}^{N_C}$ be the set of signal

endowments for nodes in C . We look for a lower bound $\underline{\psi}_k(s)$ for each signal endowment realization in C such that, when signal endowment is s , the fraction of stuck nodes in C is larger than $\underline{\psi}_k(s)$ fraction of stuck nodes in C if the endowment is $s \geq \underline{\psi}_k(s)$. If we can find such $\underline{\psi}_k(s)$, then we can use a law of large numbers to argue that

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

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

$$\bar{\Psi}(k) \geq \bar{\psi}(k) \equiv \mathbb{E}_{s \in S_C} \left\{ \bar{\psi}_k(s) \right\} \text{ almost surely.}$$

Imagine first that the signal endowment of the upper triangle in C is $(0, 0, 0)$. Then, using Lemma 2.1 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 $s \in S_C$, which is independent of the realization of the signal endowment of other components on $O_d(k)$

When the signal endowment of the upper triangle in C is different from $(0, 0, 0)$, we make use of the other property we know from C : 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 stuck, 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 $\theta = 1$, 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 $\underline{\psi}_k(s)$ and $\bar{\psi}_k(s)$.

There are two ways of calculating $\mathbb{E}_{s \in S_C} \left\{ \underline{\psi}_k(s) \right\}$:


- (1) **Doing it explicitly:** This can be done for $k = 2$ and $k = 3$, because $\# \{S_C(k = 4)\} = 128$.
- (2) **Monte-Carlo:** As k goes bigger, it is computationally unfeasible to calculate the expected value of $\underline{\psi}_k(s)$ explicitly, since

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


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

The bounds presented in Figure C.3 were calculated using the explicit approach in (1), replicating the idea behind Panel A of Figure C.2 for all possible signal configurations for components in $O_d(3)$ (i.e., which components are complete binary quilts of depth $k = 3$). Following Lemma C.1, as $d \rightarrow \infty$, nodes in $O_d(3)$ account for approximately $2^{3+1} - 1/2^{3+1} = 93.75\%$ of all nodes, so we are looking at most nodes by just looking at the last 3 levels.


APPENDIX D. ONLINE APPENDIX: FIGURES



Panel A: Distinguishing between Bayesian and DeGroot uniform weighting




Panel B: Distinguishing between Bayesian and DeGroot degree weighting




Panel C: Distinguishing between Bayesian and DeGroot eigenvector weighting


FIGURE D.1. Divergence versus power frontier for network 1 selection



Panel A: Distinguishing between Bayesian and DeGroot uniform weighting



Panel B: Distinguishing between Bayesian and DeGroot degree weighting



Panel C: Distinguishing between Bayesian and DeGroot eigenvector weighting

FIGURE D.2. Divergence versus power frontier for network 2 selection