

**Stanford** | Center on Global Poverty  
and Development

# Testing Models of Social Learning on Networks: Evidence from Two Experiments

Arun G. Chandrasekhar  
Horacio Larreguy  
Juan Pablo Xandri

July, 2018

Working Paper No. 571

# TESTING MODELS OF SOCIAL LEARNING ON NETWORKS: EVIDENCE FROM TWO EXPERIMENTS

ARUN G. CHANDRASEKHAR<sup>‡</sup>, HORACIO LARREGUY<sup>§</sup>, AND JUAN PABLO XANDRI<sup>\*</sup>

**ABSTRACT.** We theoretically and empirically study an incomplete information model of social learning. Agents initially guess the binary state of the world after observing a private signal. In subsequent rounds, agents observe their network neighbors' previous guesses before guessing again. Types are drawn from a mixture of learning models—Bayesian, where agents face incomplete information about others' types, and DeGroot, where agents follow the majority of their neighbors' previous period guesses. We study (1) learning features of both types in our incomplete information model; (2) what network structures lead to failures of asymptotic learning; (3) whether realistic networks exhibit such structures.

We conducted lab experiments with 665 subjects in Indian villages, and 350 students from ITAM in Mexico. We conduct a reduced form analysis and then structurally estimate the mixing parameter, finding the share of Bayesian agents to be 10% and 50% in the village and student samples, respectively.

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

**JEL CLASSIFICATION CODES:** D83, C92, C91, C93

## 1. INTRODUCTION

Information and opinions about new technologies, job opportunities, products, and political candidates, among other things, are largely transmitted through social networks. However, the information individuals receive from others often contains noise that individuals need to filter. A priori, individuals are likely to differ in their sophistication (Bayes-rationality) or naivete about how they engage in social learning. That is, they might vary in the extent that they are able to assess how much independent information is contained among the friends with whom they communicate and also vary in whether they assess the naivete of their friends.

---

*Date:* First Version: August 2011, This Version: July 16, 2018.

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 Devika Lakhote, Gowri Nagaraj, Adriana Paz, Mounu Prem, Alejandra Rogel, Diego Dominguez, José Ramón Enríquez, Piotr Evdokimov, Andrei Gomberg and Juan Pablo Micozzi for their assistance. We thank the Russell Sage Behavioral Economics Grant, the NSF GRFP (Chandrasekhar), and the Bank of Spain and Caja Madrid (Larreguy) for financial support.

<sup>‡</sup>Stanford University, Department of Economics; NBER; JPAL.

<sup>§</sup>Harvard University, Department of Government.

<sup>\*</sup>Princeton University, Department of Economics.

In many settings, individuals cannot or do not transmit their beliefs to an arbitrarily fine degree nor their information sets.<sup>1</sup> Here individuals process coarser information from their network neighbors when engaging in learning. This may be for many reasons, including but not limited to settings where learning operates through observations of others' actions or the costs of communicating very complex information are too high and therefore only summaries are transmitted.<sup>2</sup>

We study such a coarse communication environment, where individuals receive signals about an unknown binary state of the world in the first period and, in subsequent rounds, they communicate to their neighbors their best (binary) guess about the state of the world.<sup>3</sup> In this ubiquitous setting, if all agents are Bayesian, and there is common knowledge of this, under mild assumptions, learning will be asymptotically efficient in large networks (Gale and Kariv (2003), Mossel and Tamuz (2010), and Mossel, Sly, and Tamuz (2015) in a strategic learning environment). However, if all agents update their guess as the majority of their neighbors' prior guesses—as modeled by the DeGroot action model of learning (DeGroot, 1974), also known as the majority voting model (Liggett, 1985)—then it is possible that a non-trivial set of agents will end up stuck making the wrong decision. In practice, it might be that there is a mix of sophisticated (Bayesian) and naive (DeGroot) action learners, and that Bayesians are aware of this and incorporate it in their calculations. Such an incomplete information model, its relevance, and implications for asymptotic learning have not been studied in our environment.

This paper develops an incomplete information model of social learning on a network in which agents can potentially be Bayesian or DeGroot, and agents have common knowledge of the distribution of Bayesian or DeGroot types in the population. Bayesian agents then learn in an environment of incomplete information. The model nests the two extreme cases—complete information common knowledge of Bayes-rationality and all DeGroot—and is a hybrid that serves as an empirically relevant

---

<sup>1</sup>In other cases, communicating very fine information may be more reasonable such as in discussing political opinions. Surely this is context specific.

<sup>2</sup>Examples of coarse information include observing whether one's neighbor intends to join microfinance, the kind of crop or fertilizer used, or the candidate for whom the neighbor may vote.

<sup>3</sup>DeGroot (1974); DeMarzo, Vayanos, and Zwiebel (2003); Gale and Kariv (2003); Golub and Jackson (2010); Mossel and Tamuz (2010); Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012); Feldman, Immorlica, Lucier, and Weinberg (2014), among others, study models where signals are endowed in the first period and then agents repeatedly communicate to learn the state of the world. Under DeGroot communication learning asymptotic convergence occurs via a law of large numbers argument and linear algebra results give bounds on the speed of convergence. This model has been justified as the naive application of a one-period Bayes-rational procedure of averaging.

benchmark as, a priori, agents are likely to be heterogeneous in the degree of sophistication of learning styles. In what follows we use “incomplete information Bayesian” and “hybrid” model interchangeably.

We then study the data from two lab experiments we conducted, one in 2011 in the field with 665 Indian villagers and another in 2017 with 350 university students of ITAM in Mexico City, to examine whether subject learning behavior is consistent with DeGroot, Bayesian, or a mixed population. In each experiment, we randomly placed seven subjects into a connected network designed to maximize the ability to distinguish between the learning models, and let agents anonymously interact with each other in a social learning game. We study reduced patterns of learning, informed by the theory, and structurally estimate the mixing parameter of the model via maximum likelihood estimation, which we demonstrate delivers a consistent estimate. Conducting the experiment in two distinct locations, with a particularly different educational background, enables us to apply our methods broadly and also consider whether learning behavior differs by context.<sup>4</sup>

Our theoretical and empirical results are as follows. Beginning with the theory, we first identify four learning patterns that distinguish Bayesian and DeGroot agents in the incomplete information model. In particular, we identify a key network feature that sets apart the learning types, which we denote a *clan*. This is a set of individuals who each have more links among themselves than to the outside world. The first pattern we demonstrate is that, if a clan is comprised entirely of DeGroot agents, and they ever agree with each other about the state of the world in a period, they will never change their opinions in all future periods, even if they are wrong. We denote these agents as stuck. The second pattern is that, in the complete information Bayesian model, if agent  $i$ 's neighborhood is contained in agent  $j$ 's,  $i$  always copies  $j$ 's prior period guess. The third pattern is that, even under incomplete information, any Bayesian agent who at any point learns the true majority of initial signals never changes her guess. The fourth pattern is that no Bayesian  $j$ , even under incomplete information, ever responds to an agent  $i$  whose neighborhood is contained in  $j$ 's.

Our second theoretical result is to contrast the incomplete information Bayesian model with the conventional complete all-Bayesian model and the DeGroot communication model where agents can pass on their exact beliefs. For any sequence of growing

---

<sup>4</sup>Importantly, our lab experimental analysis abstracts from other possibly important, often unobserved determinants of learning behavior. As with any other lab experiment, and more generally work emphasizing causal identification, we prioritize studying fundamentals of learning behavior to analyzing the effect of confounders of such fundamentals.

networks with uniformly bounded degree, the all-Bayesian model leads to asymptotic learning, whether agents behave myopically (choosing the short-run best response) or strategically (playing a Nash equilibrium on a repeated game where agents may have incentives to experiment). Further, if DeGroot agents can pass on arbitrarily fine information, then they too all learn the state of the world. In contrast, if the sequence has a non-vanishing share of finite clans in the limit, then the incomplete information Bayesian model exhibits failure of asymptotic learning and a non-vanishing share of agents become stuck guessing the wrong state of the world forever. This is true even if Bayesian agents play strategically rather than myopically.

Third, we address whether realistic networks have clans, which lead to failure of asymptotic learning in the incomplete information Bayesian model. We study a mixture of two canonical models – random geometric graphs and Erdos-Renyi graphs – to model sparse – (each agent has few links) and clustered (agents’ friends tend to be themselves friends) networks, which are hallmarks of realistic network data (Penrose, 2003; Erdős and Rényi, 1959). We show that across the entire parameter space, the share of clans is uniformly bounded from below and therefore learning under the incomplete information Bayesian model must lead to a failure of asymptotic learning.

Turning to our empirical results, we begin with a reduced form analysis of the differing patterns of Bayesian and DeGroot learning that we derived. Subjects behave largely consistent with DeGroot learning in the Indian village sample but exhibit mixed learning behavior in the Mexican college sample. Specifically, first, 94.6% of the subjects in the Indian experiment who are in a clan that comes to consensus remain stuck on the wrong guess when the all-Bayesian model would suggest they change their guess. However, in the Mexican experiment this number is 30.3%. Second, over 82.9% of subjects in the Indian sample who have an information set that is dominated by a network neighbor fail to simply copy their neighbor (which is what a Bayesian would do). In contrast, this failure occurs 54.5% of the time in the Mexican data. Third, nearly 94.5% of subjects in the Indian experiment respond to changes in their neighbors’ guesses despite learning the majority of initial signals in the entire network, which Bayesian agents would never do. This happens 60.0% of the time in the Mexican sample. Fourth, 93.1% in the Indian sample respond to changes in the behavior of agents with a dominated information set, which again Bayesian agents would never do. In contrast, this happens 61.4% of the time in the Mexican sample.

We then turn to the structural estimation of the mixing parameter in the incomplete information model. There are two perspectives that we can take with the data. The

first is that the unit being analyzed is the entirety of the dynamic of social learning at the network level. The focus is not describing an individual’s behavior per se but rather the social learning process as a whole. The second is at the individual level, where we look at each agent in each period, given a history. We estimate the parameters of the model under both perspectives, but we prefer the network-level perspective since we are interested in explaining the overall social learning process.

We find similar results from both perspectives. The network-level estimator indicates that the mixture parameter is 0.1 (with a standard error of 0.130) in the Indian network data, whereas the mixture parameter is 0.5 (with a standard error of 0.184) in the Mexican data. At the individual level, the mixture parameter is 0.1 (standard error 0.113) in the Indian network data, while the parameter is 0.4 (standard error 0.268) in the Mexican data. The individual-level analysis naturally runs into the problem of zero-probability events happening since agents do not internalize the fact that others may make mistakes. We deal with this by specifying that the data generating process, and thus the model we estimate, terminates once an agent hits a zero probability event. Another way one might have dealt with this would be by adding trembles and studying quantal response equilibria [Choi et al. \(2012\)](#). We show in our setting that this is computationally infeasible.<sup>5</sup> So, instead, we model agents that do not internalize mistakes and only consider the experimental data before agents reach zero-probability events.

We contribute to an active both theoretical and experimental discussion about social learning on networks. The theoretical work closest to ours are [Feldman et al. \(2014\)](#), [Mossel and Tamuz \(2010\)](#), [Mossel et al. \(2014\)](#), and [Mossel et al. \(2015\)](#). [Mossel and Tamuz \(2010\)](#) consider a sequence of growing graphs and show that, in our setting, if all agents are Bayesian, then in the limit there will be consensus on the true state. [Mossel et al. \(2015\)](#) show that the same result holds true if agents behave strategically, where agents follow Nash equilibrium strategies of a repeated game. Turning to the DeGroot action model, [Feldman et al. \(2014\)](#) study a variation with asynchronous learning behavior, where in each period a single node is chosen uniformly at random to update, and show that, as long as the network is sufficiently expansive, there will also be asymptotic learning. [Mossel et al. \(2014\)](#) focus on a synchronous setting and study when networks are such that the majority reach the right opinion, and study unanimity in the special case of  $d$ -regular graphs, and like

---

<sup>5</sup>A back of the envelope calculation shows that using QRE in our setting with 7 agents would take 377,346,524 years in a case where the same computation with 3 agents this would take 4.5 hours.

Feldman et al. (2014) show that consensus on the right opinion will be reached for sufficiently expansive graphs.

Our theoretical results extend these results. Studying the synchronous case with incomplete information on irregular graphs, the concept we identify—clans—allows us to relate our result on asymptotic learning failures to graph conductance and a different notion of expansiveness (Chung, 1997). We show that a large family of flexible network formation models that reflect real-world data will have non-trivial clan shares to characterize the pervasiveness of asymptotic learning failures. Finally, our model examines the robust implications of Bayesian agents’ behavior in the presence of incomplete information about others’ types, which is entirely new to the literature.

Failures of learning on networks have also been well-studied outside of the repeated communication setting. A segment of the literature sets up the problem as sequential learning. They study asymptotic learning in an environment consisting of a directed network where agents move once and see a subset of their predecessors’ choices (Banerjee, 1992; Bikhchandani et al., 1992; Smith and Sorensen, 2000; Acemoglu et al., 2011; Lobel and Sadler, 2015; Eyster and Rabin, 2014).<sup>6</sup> The environment is quite different from our repeated communication setting since, as noted by Eyster and Rabin (2014), “(a) networks are undirected and (b) players take actions infinitely often, learning from one another’s past actions”. Thus, the asymptotic learning failures identified by our main theoretical results are conceptually different.

We also contribute to the experimental literature that began by documenting stylized facts about learning. Choi et al. (2005) demonstrated that, in networks of three nodes, the data is consistent with Bayesian behavior. Similarly, Corazzini et al. (2012) show that, in networks of four nodes, where agents receive signals and the goal is for them to estimate the average of the initial signals and agents could pass on real-numbers (arbitrarily fine information), the eventual averages reflected double counting, consistent with DeGroot-like behavior. Meanwhile, Möbius et al. (2015) conduct a field experiment to pit DeGroot learning against a Bayes-like – though decidedly non-Bayesian – alternative where agents may “tag” information (pass on information about the originator) to dampen double-counting. Finally, subsequent to our work, Mueller-Frank and Neri (2013) and Mengel and Grimm (2014) also conducted lab experiments to look at Bayesian versus DeGroot learning. Mueller-Frank

---

<sup>6</sup>Eyster and Rabin (2014) show that “social learning-rules that are strictly and boundedly increasing in private signals as well as neglect redundancy, in which no player anti-imitates any other have, with positive probability, that society converges to the action that corresponds to certain beliefs in the wrong state.”

and Neri (2013) develop a general class of models called Quasi-Bayesian updating and show that long-run experimental outcomes are consistent with this model. Crucial differences from our work include the fact that neither paper allows for agents to be drawn from a mixed population and the latter (and initially the former) did not reveal the network to their subjects, making the inference problem more complicated.

Our contribution to the literature is to first directly address whether social learning patterns reflect for the most-part DeGroot-like or Bayesian-like behavior, or even the behavior of a heterogenous population in an incomplete information Bayesian environment. This framework nests all prior work with pure DeGroot or complete information Bayesian models at the extremes. We take seriously the idea that, a priori, heterogeneity in the sophistication of learning types is likely. We identify a number of distinguishing features between learning types that are even robust to incomplete information. We demonstrate the relationship between social learning failures and structures of the underlying network such as the presence of clans or information dominating neighborhoods. Finally, we by conducting our experiment in two contrasting setting, we show that the mixing parameter might be context dependent, which has important implications for research and policy making.

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework and patterns of behavior by DeGroot and Bayesian agents in our incomplete information model. Section 3 contains the experimental setup. In Section 4 we explore the raw data and show reduced form results from the perspective of the contrasting learning patterns developed in the theory section. Section 5 describes the structural estimation procedure and the main results of such estimation. Section 6 concludes. All proofs are in Appendix A.

## 2. THEORY

We develop an incomplete information model of social learning on networks where every agent in the network is drawn to be either a Bayesian type or a DeGroot type with independent and identically distributed probability ( $\pi$ ). This nests the pure DeGroot model ( $\pi = 0$ ), the complete information Bayesian model ( $\pi = 1$ ), and the incomplete information model with ( $\pi \in (0, 1)$ ).

### 2.1. Setup.

2.1.1. *Environment.* We consider an undirected, unweighted graph  $G = (V, E)$  with a vertex set  $V$  and an edge list  $E$  of  $n = |V|$  agents. In an abuse of notation, we use



$G$  as the adjacency matrix as well, with  $G_{ij}$  being an indicator of whether  $ij \in E$ . We let  $N_i = \{j : G_{ij} = 1\}$  be the neighborhood of  $i$  and let  $N_i^* := N_i \cup \{i\}$ .

Every agent has a type  $\eta_i \in \{D, B\}$  — DeGroot (D) or Bayesian (B). We assume this type is drawn independently and identically distributed by a Bernoulli with probability  $\pi$  of  $\eta_i = B$ . This describes how the agent processes information, either according to the DeGroot model or by using Bayesian updating. The process by which  $\eta_j$  are drawn is commonly known as is the structure of the entire network.<sup>7</sup> Thus, it is in an incomplete information Bayesian model.

Individuals in the network attempt to learn about the underlying state of the world,  $\theta \in \{0, 1\}$ . Time is discrete with an infinite horizon, so  $t \in \mathbb{N}$ . Then in every period  $t = 1, 2, \dots$ , every agent takes an action  $a_{i,t} \in \{0, 1\}$ , which is the guess about the underlying state of the world.

At  $t = 0$ , and only at  $t = 0$ , every agent receives an independent and identically distributed signal

$$s_i = \begin{cases} \theta & \text{with probability } p \\ 1 - \theta & \text{with probability } 1 - p, \end{cases}$$

for some  $p \in (\frac{1}{2}, 1)$ . Let  $s = (s_1, \dots, s_n)$  denote the initial signal configuration. Then, at the start of every period  $t > 1$ , every agent observes the history of play by each neighbor  $j$ . Let  $A^{t-1}$  denote the set of actions  $\{a_{i\tau}\}_{i=1, \tau=1}^{n, t-1}$ , so every  $i$  in period  $t$  observes  $A_i^{t-1}$ .

2.1.2. *Learning.* Consider a DeGroot agent ( $\eta_i = D$ ). This agent, in an action model, follows the majority of her guess and her neighborhood's guesses in the prior period. We assume that for ties the agent simply follows her prior period's guess.<sup>8</sup>

Formally, for  $t > 1$

$$a_{i,t} = \begin{cases} 1 & \text{if } \frac{\sum_{j=1}^n a_{j,t-1} \cdot G_{ij} + a_{i,t-1}}{|N_i^*|} > \frac{1}{2} \\ 0 & \text{if } \frac{\sum_{j=1}^n a_{j,t-1} \cdot G_{ij} + a_{i,t-1}}{|N_i^*|} < \frac{1}{2} \\ a_{i,t-1} & \text{if } \frac{\sum_{j=1}^n a_{j,t-1} \cdot G_{ij} + a_{i,t-1}}{|N_i^*|} = \frac{1}{2}. \end{cases}$$

Next, consider a Bayesian agent ( $\eta_i = B$ ). Since there is incomplete information about the types of the other agents in the network, Bayesian individuals attempt to learn about the types  $\eta_j$  of all other agents in the network while learning about the

<sup>7</sup>As we make clear later, this feature is only relevant for Bayesian agents as DeGroot agents only rely on the prior period actions of their neighbors.

<sup>8</sup>We designed the experiments so that minimize the possibility of such ties.

underlying state of the world in order to make their most informed guess about it in every period.<sup>9</sup> Formally, the relevant states for the Bayesian agent are not just the signal endowments but also the types of players in the network. Thus, we take the state of the world to be  $\omega = (s, \eta) \in \Omega := \{0, 1\}^n \times \{0, 1\}^n$ . We formalize the model in Appendix B.

**2.2. Patterns of behavior by DeGroot and Bayesian agents.** We examine several distinguishing patterns of learning behavior by DeGroot and Bayesian agents in our setting, which we make use of when analyzing the experimental data. To that end, it is helpful to start by defining the concept of *stuckness* and *clan*. Denote a node to be *stuck* if, from some period on, the actions she chooses are the opposite to the optimal decision with full information (the majority of signals in  $s$ ).

Define a *clan* as a set of nodes who are more connected among themselves than to those outside the group. Given network  $G$ , and a subset  $H \subset V$ , the induced subgraph is given by  $G(H) = (V(H), E(H))$  consisting of only the links among the subgroup  $H$ . Let  $d_i(H)$  be the degree of node  $i \in H$  counted only among partners within  $H$ . A group  $C \subset V$  is a *clan*, if for all  $i \in C$ ,  $d_i(C) \geq d_i(V \setminus C)$  and  $|C| \geq 2$ . An example of a clan with three nodes in Panel C of Figure 2 is that of nodes (2,3,6). The entire set of nodes comprises a clan. Nodes (1,3,6) do not constitute a clan.

With these concepts defined, our first result describes how a set of all-DeGroot agents in a clan who reaches a consensus opinion at any point cannot ever change their mind for any  $\pi < 1$ .

**PROPOSITION 1.** Assume all agents in a clan  $C$  are DeGroot and there exists  $t \geq 1$  such that for all  $i, j \in C$ ,  $a_{it} = a_{jt} =: a$ . Then  $a_{i,t+\tau} = a$  for all  $i \in C$  and  $\tau > 0$ .

*Proof.* See Appendix A. □

This immediately implies that among a set of DeGroot learners, if any clan agrees on the wrong state of the world at any period, then all nodes in the clan are forever stuck on that state irrespective of the other agents' types. The result does not depend on  $\pi$  and applies for  $\pi < 1$  since DeGroot agents simply average their neighbors' past period behavior.<sup>10</sup> See Figure 1 for an illustration.

<sup>9</sup>We do not consider the possibility that players engage in experimentation in early rounds. While such is a theoretical possibility, anecdotal evidence from participants suggests that it does not fit our experimental data. In addition, the theoretical and experimental literature assumes away experimentation (see, e.g., Choi et al. (2012), Mossel et al. (2015)).

<sup>10</sup>With probability  $(1 - \pi)^{|C|}$  a clan  $C$  consists of all DeGroot agents.

Next, we turn to the all-Bayesian case ( $\pi = 1$ ). Here, since all agents are Bayesian and  $\pi = 1$  is commonly known, this is the complete information Bayesian case. Consider nodes 3 and 2 in Panel C of Figure 2. Note that  $N_2^* \subset N_3^*$ . In this case, the information set of node 3 dominates that of node 2. Our second result illustrates that, when  $\pi = 1$ , any agent who is informationally dominated by another always copies the dominating agent's prior period guess. In this example, node 2 always should copy node 3's prior period guess.

**PROPOSITION 2.** Assume that all agents are Bayesian ( $\pi = 1$ ). Consider any  $i$  and  $j$  such that  $N_i^* \subset N_j^*$ . Then  $a_{i,t} = a_{j,t-1}$  for all  $t > 2$ .

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

We then turn to the intermediate case where  $\pi \in (0, 1)$ , and thus there is incomplete information for Bayesian agents. We consider two learning patterns of these agents. First, if a Bayesian agent ever learns the majority of the signals that were initially distributed, her guess should never change, irrespective of  $\pi$ . Note that, in this case, every signal need not be learned – only that a majority is either 1 or 0. Second, any agent  $i$  who is Bayesian will never respond to the actions of any other agent  $j$  whose neighborhood is informationally dominated by the neighborhood of  $i$ , irrespective of  $\eta_j$  and  $\pi$ , after period 2. In our above example, this means that, if node 3 is Bayesian, she should never respond to any behavior by node 2 after period 2.

**PROPOSITION 3.** Consider any  $\pi \in (0, 1)$  and suppose  $\eta_i = B$ .

- (1) If at any period  $\tau$  agent  $i$  learns the majority  $\hat{\theta} := \mathbf{1} \left\{ \frac{1}{n} \sum_j s_j > \frac{1}{2} \right\}$ , then  $a_{i,t} = \hat{\theta}$  for all  $t \geq \tau$ , irrespective of any future sequence of actions by  $i$ 's neighbors.
- (2) If  $N_j^* \subseteq N_i^*$ , and that for all  $\omega \in \Omega$ ,  $a_{j,1}(\omega) = s_j$ , then for  $t > 2$ ,  $s_j$  is sufficient for  $A_j^t$  when explaining  $a_{i,t}$ . That is,  $a_{i,t} = \text{Function} \left( s_j, \left( A_k^{t-1} \right)_{k \in N_i \setminus \{j\}} \right)$ .

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

In both cases, the intuition is that, if  $j$ 's action reveals no new information to  $i$ , then  $j$ 's choice only matters through the initial signal:  $s_j = a_{j,1}$ . There is, therefore, a conditional sufficiency requirement: conditional on the actions of neighbors other than  $j$ ,  $a_{i,t}$  should be conditionally independent of  $a_{j,t}$  for all  $t > 2$ .

Below, when we present our raw data, we revisit these properties in our reduced form results. We check the extent to which clans get stuck, informationally dominated

agents tend to copy the dominating agents, and agents who necessarily learn the majority or necessarily dominate other agents ignore the other agents going forward.

**2.3. Asymptotic efficiency: Long-run behavior in large networks.** Having studied differences in action learning behavior in the DeGroot model and Bayesian model, both with complete ( $\pi = 1$ ) and incomplete information ( $\pi \in (0, 1)$ ), we assess how efficient learning is in the long run in a large network under these models.

We compare the complete information Bayesian action model ( $\pi = 1$ ) with a purely incomplete information Bayesian action model ( $\pi < 1$ ). We also include another benchmark—the DeGroot model where the agents can communicate their “beliefs” rather than a coarser “action” (or guess)—as this is what has been classically studied in the literature. Note that this model is considerably richer in terms of what it can communicate, far richer than the DeGroot action (or majority) and Bayesian action models, where agents can only transmit 0 or 1 in our model.

The main point that we illustrate is that, for a sensible sequence of networks, both the complete information Bayesian action and DeGroot communication models lead to asymptotic efficiency where all but a vanishing share of agents converge to a right guess about the state of the world. However, in the incomplete information model, there is a non-vanishing share of agents that are stuck with the wrong guess.

To make these comparisons precise, we need to nest the different models. Assume without loss of generality that every agent  $i$  receives a signal  $p_{i,0}$  that is continuously distributed in  $[0, 1]$ . In the DeGroot communication model, these  $p_{i,0}$ ’s and their subsequent averages can be communicated directly. That is,

$$p_{i,t} = \frac{\sum_{j \in N_j} p_{j,t-1} + p_{i,t-1}}{|N_i^*|}.$$

The action then is  $a_{i,t} = \mathbf{1}\{p_{i,t} > \frac{1}{2}\}$  and we say the limit action  $a_{i,\infty} = \lim_{t \rightarrow \infty} \mathbf{1}\{p_{i,t} > \frac{1}{2}\}$ .

Meanwhile, in the Bayesian and DeGroot action models, we can think of  $p_{i,0} = \mathbb{P}(\theta = 1 | s_i)$  — so the signal delivered at  $t = 0$  equivalently generates a posterior. Information transmission and updating are as before, with  $a_{i,t}$  depending on  $A_i^{t-1}$  using either Bayesian updating or DeGroot updating via majority.

We consider a sequence of networks  $G_n = (V_n, E_n)$  where  $|V_n| = n$ , letting  $p_{i,t}^{(n)}$  be defined as above for the DeGroot communication model and  $a_{i,t}^{(n)}$  analogously be defined for each model. We say that a sequence is *asymptotically efficient* if for all  $\epsilon > 0$  we have

$$\lim_{n \rightarrow \infty} \max_{i \leq n} \mathbb{P} \left\{ \lim_{t \rightarrow \infty} |a_{i,t}^{(n)} - \theta| \leq \epsilon \right\} = 1$$

Under very general conditions, we show that both the DeGroot communication model and the complete information Bayesian action model achieve asymptotic efficiency, but the incomplete information version may not.

**THEOREM 1.** Suppose  $G_n = (V_n, E_n)$  with  $|V_n| = n$  is such that (i) there is a uniform bound on degree:  $d_i(V_n) \leq \bar{d}$  for all  $i, n$ ; (ii) the posterior distributions  $\mathbb{P}(\theta|s_i)$  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 complete information Bayesian action model is asymptotically efficient, and
- (3) the incomplete information Bayesian action model with  $\pi < 1$  may not be asymptotically efficient. In particular, suppose there exist  $k < \infty$  such that  $X_n := \#\{i : i \text{ is in a clan of size } k\} / n$  is positive in the limit. Then the DeGroot action model is not asymptotically efficient.

*Proof.* See Appendix A. □

Note 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), also shown in Menager (2006)). Both results 2 and 3 remain true if Bayesian agents play a Nash equilibrium in the normal form of the game instead of the assumption of myopic behavior by Bayesian agents Mossel, Sly, and Tamuz (2015) show that the same asymptotic efficiency result remains true in a strategic setting. Moreover, in the incomplete information setting, the result showing asymptotic inefficiency relies on bounding the random number of clans formed only by agents with DeGroot types, so while Bayesian types may be playing according to an equilibrium of the normal form game, DeGroot types may not.

The theorem illustrates an important discrepancy between these models in terms of social learning and why differentiating between which model best describes agents is relevant. If agents cannot communicate all the information they have observed, the coarseness of these messages causes beliefs to get “stuck” on a particular action, not allowing the flow of new information. This is particularly true of DeGroot agents in the incomplete information model, as seen in Proposition 1. Because  $\pi$  is fixed in  $n$ , the share of finite-sized clans that are all-DeGroot does not vanish in  $n$ , and therefore the share of nodes that get stuck is non-vanishing in  $n$ . 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.

**2.4. What kinds of networks exhibit stuckness?** Having examined the properties of DeGroot, complete information Bayesian, and incomplete information Bayesian models, we address whether DeGroot-like behavior may cause lack of asymptotic learning in practice. In particular, we study whether realistic network structures would exhibit stuckness in a setting in which the share of DeGroot agents is non-zero ( $\pi < 1$ ). Specifically, we explore whether, as the number of nodes  $n \rightarrow \infty$  asymptotically, there is a non-vanishing share of clans. If so, the share of all-DeGroot clans is bounded from below, and consequently, there is a non-vanishing lower bound on the share of agents getting stuck.

We begin by examining a stochastic network formation model that mimics the structure of real-world networks — the resulting graphs tend to be both sparse and clustered.<sup>11</sup> This model is very similar to a number of network formation models in the statistics and econometrics literature (see, e.g., [Fafchamps and Gubert \(2007\)](#); [Graham \(2017\)](#); [McCormick and Zheng \(2015\)](#)) and has a random utility interpretation. We then relate stuckness to a concept from graph theory called conductance and related properties of expansion. This allows showing that readers to show can check an eigenvalue of a standard transformation of the adjacency matrix  $G$  to evaluate whether stuckness, and consequently lack of asymptotic learning, is possible in networks of interest.

**2.4.1. A Hybrid Model of Random Geometric Graphs and Erdos-Renyi Graphs.** We build a simple, but general, model to capture sparse and clustered network structures, which resemble those in the real world. Our model toggles between two starkly different canonical models of network formation: Random Geometric Graphs (henceforth RGG) and Erdos-Renyi (henceforth ER) ([Penrose, 2003](#); [Erdős and Rényi, 1959](#)). The basic idea is that, in an RGG, nodes have positions in some latent space and tend to be linked when they are close enough, capturing latent homophily. Meanwhile, in an ER, nodes are independently linked.

We demonstrate that, for such a mixture, the share of clans is non-vanishing. Practically, this means that, if  $\pi < 1$ , we should expect that, in networks with real-world like structures, there will remain pockets of individuals who become irreversibly

---

<sup>11</sup>See Online Appendix [E](#) for simulations that demonstrate this on empirical network data.

convinced of the wrong state of the world and are unable to shift away from this, if they behave in a DeGroot action (or majority) manner.

Let  $\Omega \subset \mathbb{R}^2$  be a latent space.<sup>12</sup> We say an *RGG-ER mixture* is  $(\alpha, \beta)$ -mixed if the network, where  $(\alpha, \beta) \in [0, 1]^2$ , is formed as follows. There exists a Poisson point process on  $\Omega$ , which determines which points of the latent space will receive a node.  $n$  nodes are drawn according to this point process, with uniform intensity  $\lambda > 0$ . Note this means that for any subset  $A \subset \Omega$

$$n_A \sim \text{Poisson}(\nu_A) \text{ where } \nu_A := \lambda \int_A dy.$$

Given the draw of  $n$  nodes, the network forms as follows.

(1) RGG component:

- If  $d(i, j) \leq r$  then  $i$  and  $j$  are linked with probability  $\alpha$ .
- If  $d(i, j) > r$  then there is no link between  $i$  and  $j$ .

(2) ER component:

- Every pair  $ij$  with  $d(i, j) > r$  is linked i.i.d. with probability  $\beta$ .

There is a simple random utility interpretation of this model similar. Let the latent utility to  $i$  of a link to  $j$  be, for  $\theta > \gamma$ ,

$$u_i(j) = \theta \cdot \mathbf{1}\{d(i, j) \leq r\} + \gamma \cdot \mathbf{1}\{d(i, j) > r\} - \epsilon_{ij}$$

where  $\epsilon_{ij} \sim F(\cdot)$ , and  $F(\cdot)$  is Type I extreme value. Assume  $\epsilon_{ij} = \epsilon_{ji}$  so that, if one wants the link, so does the other so that mutual consent is satisfied. Then

$$\mathbb{P}(G_{ij} = 1) = F(\theta) \cdot \mathbf{1}\{d \leq r\} + F(\gamma) \cdot \mathbf{1}\{d > r\}.$$

Setting  $\alpha = F(\theta)$  and  $\beta = F(\gamma)$  gives us the  $(\alpha, \beta)$ -mixture of RGG and ER.

We assume  $\alpha \geq \beta$  in our applications. Note that, if  $\alpha = 1$  and  $\beta = 0$ , then this is just a standard RGG. If  $\alpha < 1$  and  $\beta = 0$ , this is what is called a soft RGG. And if  $\alpha = \beta > 0$ , then this is a standard ER.

Formally, we are interested in a sequence of networks  $(\Omega_k, \alpha_k, \beta_k)$  such that  $\text{vol}(\Omega_k) \rightarrow \infty$ .<sup>13</sup> Since we study sparse graphs, we need to ensure that both components contribute a sparse number of links. The RGG component is sparse by definition for fixed  $r$ . For the ER component, this requires  $\beta_k = O\left(\frac{1}{\text{vol}(\Omega_k)}\right)$ .

We derive lower and upper bounds on a particular class of clans, which we call *local clans*. A local clan is a group located in a ball of radius  $r/2$ , so that the probability of

<sup>12</sup>The result is easily generalized for  $\mathbb{R}^h$ ,  $h \geq 2$

<sup>13</sup>For example, if  $\Omega = [0, k]^h$ —here we take  $k = 2$ —then  $\text{vol}(\Omega) = k^h$  and we take  $k \rightarrow \infty$ .

any two nodes in this group having a link, is  $\alpha_k$ . A local clan has an expected number of nodes within the ball that is constant along the growing sequence of nodes.

**THEOREM 2.** Consider any  $(\alpha_k, \beta_k)$ -mixture of Random Geometric and Erdos-Renyi graphs, and a sequence  $\{\Omega_k\}_{k \in \mathbb{N}}$  such that  $\text{vol}(\Omega_k) \rightarrow \infty$ . If  $\beta_k = O\left(\frac{1}{\text{vol}(\Omega_k)}\right)$  and  $\alpha_k \rightarrow \alpha > 0$ , then the share of nodes belonging to local clans remains positive as  $k \rightarrow \infty$ . If  $\alpha_k \rightarrow 0$ , the shares of such nodes vanishes as  $k \rightarrow 0$ .

The proof is a direct corollary to Proposition A.2 in Appendix A.<sup>14</sup> The proposition illustrates that, in stochastic network formation models that mimic the structure of real-world networks where graphs tend to be both sparse and clustered, if  $\pi < 1$ , a non-trivial share of nodes is likely to get stuck on the wrong state of the world, and thus asymptotic learning is unfeasible.

2.4.2. *Assessing the Presence of Clans.* A natural question to ask is whether the existence of clans is related to some well-studied structural property of the network, such as expansiveness. We show that the existence of clans relates to a measure in graph theory called conductance and a related spectral property of expansiveness. We show that any sequence of graphs that is sufficiently expansive in a specific sense cannot have a non-vanishing share of clans.

The conductance (Cheeger constant in Chung’s terminology (Chung, 1997)) of a graph is

$$\phi(g) := \min_{S: 0 < \text{vol}(S) \leq \frac{1}{2} \text{vol}(V)} \frac{\partial S}{\text{vol}(S)},$$

where  $\partial S := \sum_{i \in S} d_i(V \setminus S)$  is the number of links from within set  $S$  to  $V \setminus S$  and  $\text{vol}(S) = \sum_{i \in S} d_i(S) + \sum_{i \in S} d_i(V \setminus S)$ .

**PROPOSITION 4.** Any graph sequence  $(g_n)$  that has a non-vanishing share of clans of uniformly bounded size has  $\lim_{n \rightarrow \infty} \phi(g_n) < \frac{1}{2}$ .

*Proof.* See Appendix A. □

The conductance of the graph is well-known to be difficult to compute (NP-Complete). We appeal to the bounds in the literature called the Cheeger inequality. This uses the spectrum of the Laplacian of  $g$ ,  $L := I - D^{-1/2} g D^{-1/2}$ , where  $D = \text{diag}\{d_1(g), \dots, d_n(g)\}$ , to bound  $\phi(g)$  by the second-smallest eigenvalue of the Laplacian:

$$\frac{\lambda_2(L)}{\sqrt{2}} \leq \phi(g) \leq \sqrt{2\lambda_2(L)}.$$

<sup>14</sup>We also derive closed form bounds on the probability that a given node is part of a local clan in Proposition A.1 in Appendix A.



We say a graph is an  $h$ -Laplacian expander if  $\lambda_2(L) \geq h$ .

**COROLLARY 1.** Consider a sequence of graphs  $g_n$  which are  $\frac{\sqrt{2}}{2}$ -Laplacian expanders. Then no graph in the sequence contains a clan.

This means that one can simply assess whether agent stuckness or lack of asymptotic learning is a likely problem of a graph by looking at the values of the Laplacian’s spectrum. Any graph that is a  $\frac{\sqrt{2}}{2}$ -expander cannot have clans. As such, a sufficient condition to rule out the possibility of stuckness is that the graph has sufficiently high expansion properties.

### 3. TWO EXPERIMENTS

**3.1. Setting.** First, in 2011, we conducted 95 experimental sessions with a total of 665 subjects across 19 villages in Karnataka, India. The villages range from 1.5 to 3.5 hours’ drive from Bangalore. We initially 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 (Alt et al., 2017; Cruz et al., 2017), among other things, propagates regularly through social networks.

Second, in 2017, we conducted 50 experimental sessions with a total of 350 subjects from the pool of undergraduate students at ITAM in Mexico, who were largely junior, and from economics and political science. We chose this second setting to study how learning patterns among urban, well-educated individuals may differ from those of the rural poor. ITAM is one Mexico’s most highly-ranked institutions of higher education and systematically places every year its undergraduates in economics and political science in top PhD programs in the United States and Europe.

**3.2. Game Structure.** In both settings, the game structure followed the setup of the model. Every set of 7 subjects anonymously played the learning game on three different network structures, displayed in Figure 2, which were designed to distinguish between Bayesian and DeGroot learning behavior. Positions were randomized and order of the networks for each set of 7 subjects was also randomized.

As reflected by Figure 3, every subject received a signal about the binary state of the world correct with probability  $5/7$ . Then agents in each period submitted their guess about the state of the world. In every period, every agent could see the entire history of guesses of their network neighbors, including their own, before making their subsequent guesses.

Subjects were paid for a randomly chosen round from a randomly chosen game. In India, subjects were paid Rs. 100 if they guessed the state correctly, as well as a Rs. 20 participation fee — just under a day’s wage. In Mexico, students were paid \$100 pesos if they guessed the state correctly and an additional \$50 participation fee, which amounts to slightly above the hourly rate for working as research assistants.

**3.3. Implementation.** In every Indian village, 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.” Each station had a single staff member to monitor the single participant assigned to the station at random to ensure that participants could not observe each other or communicate. Often stations would be across several buildings.

At ITAM, we recruited 7 undergraduates for each experimental session through emails to the economics and political science mailing lists. Students were congregated in a spacious classroom and placed throughout the room also in stations so that they were unable to see and communicate with other participants. In contrast to India, each experimental session was run by two staff members since we could not afford staff to monitor each of the participants individually. However, we observed no instances of students trying to talk to another or look at other participants’ signals or guesses.

The experimental protocol was identical both in India and Mexico. 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.

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. 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 sheets of paper that presented an image of the network and colored in their neighbors' guesses. 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 decision of their neighbors and update their own beliefs accordingly. Importantly, individuals kept the information about the guesses of their neighbors in all previous rounds until the game concluded. The reason was that we intend to test social learning, but not the ability of participants to memorize past guesses.

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 the corresponding amount for a randomly chosen round from a randomly chosen game, as well as their participation fee. Participants then faced non-trivial incentives to submit a guess that reflected their belief about the underlying state of the world.

#### 4. REDUCED FORM RESULTS

We assess whether the learning patterns described in Section 2.2 hold in our experimental data. Table 1 presents the results for our experiment among Indian villagers and Table 2 presents the results for our experiment in Mexico with undergraduate students from ITAM.

Recall that we had identified four key patterns. First, irrespective of  $\pi < 1$ , if there are a set of DeGroot agents in a clan, once a clan comes to consensus, all agents remain stuck. To assess the prevalence of this feature in the experimental data, Panel A presents the share of clans that remain stuck given that the Bayesian model would have predicted a change along the path. Here an all DeGroot model would predict that the share is 1.

Second, when all agents are Bayesian (and have common knowledge of this), then any agent whose information set is dominated by that of another must follow the other agent's prior behavior for all  $t > 2$ . We look at the share of times that those informationally dominated agents fail to copy their informationally dominating neighbors, which the complete information Bayesian case predicts to be 0.

Third, even in an incomplete information setup, if any Bayesian agent learns through any history the majority of the initial signals, then the agent must play this for all future periods. Panel C, column 2, turns to the share of times an agent learns the majority of signals and yet changes guess along the path in a manner consistent with DeGroot learning.

Finally, even in an incomplete information setup, any Bayesian agent must never condition her decision on the prior period action of an informationally dominated agent and instead should restrict to only using her initial signal. In column 1 of Panel C, we look at the share of times an agent fails at this, which should be 0% of the time for Bayesian agents.

In the Indian village data, we find evidence that is consistent with DeGroot behavior and inconsistent with Bayesian behavior. In Panel A of Table 1, we show that the share of clans that remain stuck, conditioning on the cases where Bayesian agents would have changed along the path, is 0.946. Then in Panel B, we show that 82.9% an agent fails to copy an informationally dominating agent. In Panel C, column 2, we find that in 94.5% of the instances where an agent should have learned the majority, the agent changed her opinion in the direction suggested DeGroot updating. Finally, in Panel C, column 1, we find that 93.1% of the time, agents inefficiently respond to informationally dominated neighbors' actions.

The data from the experiment with ITAM students in Mexico exhibit considerably different patterns. In fact, there is evidence consistent with both DeGroot and Bayesian behavior, indicating that there is likely a more heterogeneous mix of agents from the perspective of our hybrid model. Panel A shows that stuckness occurs only 30.3% of the time. In Panel B we find that informationally dominated agents fail to copy dominating agents 54.5% of time. Column 1 of Panel C shows that, when agents learn the majority, they change their guesses in the manner that DeGroot learners would 60% of the time, and column 2 indicates that agents inefficiently respond to informationally dominated neighbors' actions 61.4% of the time.

Taken together, we see that the Indian village population behaves consistently with the (all) DeGroot model and inconsistently with any Bayesian model (or at least one of  $\pi \approx 0$ ). Meanwhile, the undergraduate student population at ITAM, who are considerably more educated, behave in a manner reflecting a possibly mixed population. The results demonstrate that context affects whether we should consider pure DeGroot model, a pure Bayesian model, or a hybrid model. Furthermore, it suggests that the sorts of misinformation traps that we have described—those that

arise when networks have clans and some individuals exhibit DeGroot-style learning behavior— might be much more of a problem for the village population. The Indian villagers, therefore, are more vulnerable to misinformation traps, whereas some set of agents in the more educated population may be able to overcome them.

## 5. STRUCTURAL ESTIMATION

We now turn to our structural estimation. Our primary approach is to consider the entirety of the social learning outcome as the subject of our study. So we take, therefore, as the object to be predicted is the entire matrix of actions  $A^T = (a_{i\tau})_{i=1, \tau=1}^{n, T}$  and study which  $\pi$  best explains the data. 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 method then maintains that the predicted action under a given model is not path-dependent and is fully determined by the network structure and the set of initial signals. We denote this approach as the *network-level estimation*.

An alternative approach is to look at *individual-level estimation*. In this case, 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 learner.

We estimate the parameters of the model under both approaches, but we prefer the network-level approach. Our aspiration is not simply to study how a given individual updates but rather, when we focus on a community as a whole, how should we think about the macro-dynamics of and the model that governs the social learning process.

For both estimation approaches, we estimate the model parameters that maximize the likelihood of observing the experimental data. In every period  $\tau$  there is an action taken by  $i$ ,  $a_{i\tau}^*$ . The type of the agent and the history determines the action. Given a history  $A^{t-1} = (a_{i\tau})_{i=1, \tau=1}^{n, t-1}$ , there is a prescribed action under the model of behavior which can depend on the agent's type  $\eta_i$ , the history of observed play, and the prior probability that an agent is Bayesian:  $a_{it}^*(A^{t-1}; \eta, \pi)$ . Agents also make mistakes with probability  $\epsilon$ , and thus, given the prescribed option, the observed data for the econometrician is

$$a_{it} = \begin{cases} a_{it}^* & \text{with probability } 1 - \epsilon \\ 1 - a_{it}^* & \text{with probability } \epsilon \end{cases}$$

for any  $t \geq 2$ . The history of play is then that of observed actions, which can differ from the prescribed action. As noted, for computational feasibility, we assume that mistakes, and thus such differences, are not internalized by agents.

The matrix  $A_v^T = [a_{it,v}]$  is the dataset for a given village  $v$  (or session in the case of the Mexican sample). Suppressing  $v$  until it is needed, the likelihood is

$$\mathcal{L}(\pi, \epsilon; A^T) = \mathbb{P}(A^T | \pi, \epsilon) = \mathbb{P}(a_T | A^{T-1}, \pi, \epsilon) \cdot \mathbb{P}(a_{T-1} | A^{T-2}, \pi, \epsilon) \cdots \mathbb{P}(a_1 | \pi, \epsilon).$$

Notice that  $\mathbb{P}(a_1 | \pi)$  and  $\mathbb{P}(a_2 | \pi)$  are both independent of  $\pi$ , because they are independent of  $\eta$ : in period 1 every agent plays their signal and in period 2 every agent plays the majority.

Let  $x_{it} = \mathbf{1}\{a_{it} = a_{it}^*(A^{t-1}; \eta, \pi)\}$ , which determines whether the observed action matches that which was prescribed by the model given the history, type vector, and parameter value. So,<sup>15</sup>

$$\mathbb{P}(a_t | A^{t-1}, \pi, \epsilon) = \prod_{i=1}^n \sum_{\eta} (1 - \epsilon)^{x_{it}} \epsilon^{1-x_{it}} \mathbb{P}(\eta | \pi).$$

Taking logs, we can write the estimator as

$$\hat{\pi} = \operatorname{argmax}_{\pi \in [0,1]} \ell(\pi; A^T, \epsilon) = \sum_{v=1}^V \sum_{t=3}^T \sum_{i=1}^n \log \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta | \pi) \right\}.$$

In Appendix C we prove consistency of the estimator and show simulations that indicate that our estimators are consistent. Specifically, first, under each  $\pi = \{0, 0.1, \dots, 0.9, 1\}$ , we generate learning data from the model. Then, using both network-level and individual-level estimation, we show that our estimators consistently recover the true parameter used for data generation.

The intuition for identification is as follows. The maximum likelihood estimator sets the score function (the derivative of the log-likelihood) to zero. So it assesses all configurations of learning types of the 7 nodes such that, at the parameter value, the likelihood of this configuration changes steeply if  $\pi$  changes minimally. The reduced form results of learning patterns that we describe above provide considerable contrast between Bayesian and DeGroot learning thus contributing to the estimation. For example, take the setting of a clan where, if all the agents were Bayesian learners, they would realize they do not need to respond to the changes in actions of their informationally dominated neighbors. Here the likelihood of stuckness declines steeply

<sup>15</sup>It is worth noting that if we could pass the logarithm then this is

$$\sum_{t=3}^T \sum_{i=1}^n \sum_{\eta} (x_{it}[A^{t-1}; \eta, \pi] \log(1 - \epsilon) + (1 - x_{it}[A^{t-1}; \eta, \pi]) \log \epsilon) \cdot \mathbb{P}(\eta | \pi)$$

and for small  $\epsilon$  this is just proportional to the expected share of actions that deviate from the model under a given  $\pi$  and a given type-configuration.

as  $\pi$  increases, as in that case it becomes increasingly likely that *some* agents are Bayesian and therefore become unresponsive to the behavior of those agents.

In order to perform inference on the parameters, we compute standard errors over the parameter estimates via a block bootstrap procedure that accounts for the dependence in the data of the individuals playing the same game and session. Specifically, we draw with replacement the same number of session-game blocks of observations that we have in each of our experimental samples and compute the parameters that maximize the corresponding likelihood.<sup>16</sup>

Before turning to our structural estimates of  $\pi$ , we estimate  $\epsilon$ , which is common irrespective of network-level or individual-level estimation. Note that for any node  $i$  in any graph  $v$  where the majority of a node's own signal and her neighbors' signals is uniquely determined, both the Bayes and DeGroot models, irrespective of  $\pi$ , prescribe the majority as an action in the second period. Therefore, recalling that  $N_i^* = \{j : g_{ij} = 1\} \cup \{i\}$ ,

$$\hat{\epsilon} := \frac{\sum_v \sum_j \mathbf{1} \{a_{j2} \neq \text{majority}(a_{j1} : j \in N_i^*)\} \cdot \mathbf{1} \{\text{unique majority}(a_{j1} : j \in N_i^*)\}}{\sum_v \sum_j \mathbf{1} \{\text{unique majority}(a_{j1} : j \in N_i^*)\}}.$$

By standard arguments, this is a consistent and asymptotically normally distributed estimator since this is comprised of a set of Bernoulli trials. Panel B of Table 3 shows that in both cases  $\hat{\epsilon}$  is similar: 0.1288 (standard error 0.007) in the Indian sample and 0.134 (standard error 0.013) in the Mexican sample. This means that about 87% of the time, agents act in accordance with the prescribed action under the incomplete information Bayesian model irrespective of the  $\pi$ .

**5.1. Main Structural Results: Network-Level Estimation.** Next we turn to the estimation of  $\pi$ . Under the network level approach, we treat the entire path of actions by all agents in the network as a single observation. From the network level perspective we take  $a_{it}^*(A^{t-1}; \eta, \pi) = a_{it}^*((a_{i0})_{i=1}^n; \eta, \pi)$ . This means that given a signal endowment, a type endowment, and a parameter  $\pi$ , the path of all agents' prescribed action in each period is deterministic. Then the observed actions are just independently identically distributed perturbations of these prescribed actions.

Panel B of Table 3 presents the results of our structural estimation from the network level perspective. In column 1, we present the results for the data from the sample of Indian villages and, in column 2, the data from the Mexican student sample. In the Indian sample, we find that the share of Bayesian agents is low— $\hat{\pi}$  is 0.1 (standard

<sup>16</sup>This procedure is analogous to clustering and, therefore, is conservative by exploiting only variation at the block level.

error 0.130) —and we cannot reject that is statistically different from zero. This is consistent with the reduced form results in Table 1, which, for example, indicate that over 90% of behavior, when Bayesian and DeGroot models disagree on the guesses that ought to be made, is consistent with DeGroot behavior.

Similarly, the structural estimates of  $\pi$  are consistent with our prior reduced form results from Table 2 for the Mexican data. There we find mixed behavior, where for example informationally dominated agents fail to copy dominating agents 54.5% of the time, and we estimate that  $\hat{\pi} = 0.5$  (standard error 0.184).

**5.2. Individual Level Estimation.** Having looked at the estimates of  $\pi$  from a network-level approach, we turn our attention the estimates from an individual-level approach.<sup>17</sup> From this perspective  $a_{it}^*(A^{t-1}; \eta, \pi) = a_{it}^*(A_i^{t-1}; \eta_i, \pi)$  depends on the entire observed history  $A_i^{t-1}$ , the agent’s type  $\eta_i$ , and the commonly known  $\pi$ .

Observe that because there is an  $\epsilon$  that is not internalized by the agents, it is possible for them to reach a zero-probability event. We, therefore, define the model as proceeding until any agent hits a zero probability event, and the terminal round  $T^*$ , which is a function of the signal and type endowment as well as the sequence of shocks, is endogenously determined as the round prior to any agent hitting a zero-probability event. The model is silent after this and therefore we treat the data in the same way. This constitutes a well-defined data-generating process and has a well-defined likelihood. We elaborate on this in Appendix C and demonstrate consistency. In practice, we consider the data until  $T = 3$  as 58% of the sessions had at least one agent hit a zero probability information set at  $T = 4$ .

We present our results in Panel C of Table 3. In column 1, we present the results for the data from the sample of Indian villagers and, in column 2, the data from the Mexican student sample. In the Indian sample, we find that the share of Bayesian agents is low:  $\hat{\pi} = 0.1$  (standard error 0.113). This is consistent with the reduced form results in Table 1. Similarly, the structural estimates of  $\pi$  are consistent with our prior reduced form results from Table 2 for the Mexican student data. In this sample, we estimate  $\hat{\pi} = 0.4$  (standard error 0.268).

**5.2.1. Trembles and Quantal Response.** In our model, agents arrive at zero probability events.<sup>18</sup> Since agents do not internalize that others can make errors, they may arrive

<sup>17</sup>It 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.

<sup>18</sup>At  $T = 4$  58% of sessions arrive at a zero probability event and this is considerably worse thereafter.



at histories that are not rationalizable. Our individual-level estimation circumvents this by defining the model to terminate at the first zero-probability information set.

A natural way to eliminate the zero-probability information set problem is to introduce disturbances (e.g., trembles or as quantal-response equilibrium as used by (Choi et al., 2012)). Individuals can make mistakes with some probability, and Bayesian agents, knowing the distribution of these disturbances, incorporate this. Unfortunately, this approach is computationally unfeasible for networks beyond a trivial size. To see this, consider the simpler case where  $\pi = 1$ , and thus there is common knowledge of this. Let us consider the cases with and without trembles.

**PROPOSITION 5.** The algorithm for computing Bayesian learning with no disturbances is  $\Theta(T)$ .<sup>19</sup> 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. □

Specifically, the algorithm is  $\Theta(n4^n T)$ . If  $n$  was growing this algorithm would be exponential time, but in our case,  $n$  is constant. Second, we show that the extension of this algorithm to an environment with disturbances is computationally intractable.

**PROPOSITION 6.** Implementing the Bayesian learning algorithm with disturbances has a computational time complexity of  $\Theta(4^{n(T-1)})$ .

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

To see the computational burden of introducing trembles, we compare them to their deterministic counterparts. For the  $\pi = 1$  model, the algorithm with trembles with  $T = 6$  involves  $1.19 \times 10^{16}$  more computations than the deterministic model. With the same  $T$ , the incomplete information model ( $\pi \in (0, 1)$ ) involves  $8.65 \times 10^{32}$  more calculations than its deterministic counterpart. Suppose that the deterministic  $\pi = 1$  model takes 1 second to run. Then the deterministic incomplete information model (again, without trembles) takes 4 and a half hours. The trembling hand  $\pi = 1$  model, however, takes approximately 377,346,524 years.

## 6. CONCLUSIONS

We study a model of incomplete information Bayesian learning on social networks in which individuals' aim to guess a binary state of the world after a series of rounds

<sup>19</sup>Recall 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  $\forall n > \underline{n}, c_1 \cdot |f_2(n)| < |f_1(n)| < c_2 \cdot |f_2(n)|$ .

of repeated communication in which they transmit their guesses to their network neighbors each period. Agents are heterogeneous in their sophistication of learning: they can either be Bayesian or DeGroot type, and Bayesian agents are cognizant of the distribution of Bayesian types in the population. This model nests the prior models in the literature that study similar settings.

We identify key network patterns that separate DeGroot and Bayesian learning behavior in the presence of incomplete information. One such concept is that of clans—subgraphs of nodes who each have more links within the group as compared to beyond the group. We show that realistic networks tend to have clans, which lead to failure of asymptotic learning: that non-trivial shares of agents will never learn the truth. This result is even robust to strategic behavior on the part of Bayesian agents.

Our empirical results demonstrate that the incomplete information model fits the data well, but the mixing parameter varies by context. In the Indian village setting, our estimates indicate that there are approximately 10% Bayesian agents whereas the data is best explained in the Mexican student setting by a share of 50% Bayesian agents. These contrasting results point at the importance of contextual factors to understand how people engage in social learning. It is possible, for instance, that more vulnerable populations may be more subject to DeGroot-type learning behavior. Future work could systematically assess the relevance of various contextual factors.

An interesting direction to take this line of inquiry is to think about the relationship between the network-formation and social learning processes. As noted by [Jackson et al. \(2012\)](#), among others, relational contracting motives generate the need for triadic closure: friends of friends tend to be friends themselves. In fact, the social quilts of [Jackson et al. \(2012\)](#) consist of only node-adjacent clans, of which our Network 3 is an example. If agents only communicate with their favor-exchange networks,<sup>20</sup> and the incomplete information learning setup describes agents' behavior, asymptotic learning would fail rampantly in communities that had to overcome contracting problems. Practically speaking, this suggests that vulnerable communities such as villages which need to organize themselves to share risk would precisely be those where we would expect misinformation traps.

---

<sup>20</sup>A natural reason for this comes from multiplexing motives. Fixed costs may be required to maintain links and therefore it makes sense naturally to use links for multiple reasons such as through layering financial, informational, and social links.

## REFERENCES

- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): “Bayesian Learning in Social Networks,” *Review of Economic Studies*, 78, 1201–1236. 1
- ALT, J. E., A. JENSEN, H. LARREGUY, D. D. LASSEN, AND J. MARSHALL (2017): “Contagious Political Concerns: Identifying Unemployment Information Shock Transmission Using the Danish Population Network,” Working Paper. 3.1
- AUMANN, R. J. (1976): “Agreeing to disagree,” *The annals of statistics*, 1236–1239. B.1
- BANERJEE, A. (1992): “A simple model of herd behavior,” *The Quarterly Journal of Economics*, 797–817. 1
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. JACKSON (2013): “The Diffusion of Microfinance,” *Science*, 341, 1236498. 3.1, E, 7
- 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. 1
- 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. 1
- (2012): “Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data,” *Review of Economic Design*, 16, 93–118. 1, 9, 5.2.1
- CHUNG, F. R. (1997): *Spectral graph theory*, 92, American Mathematical Soc. 1, 2.4.2
- CONLEY, T. AND C. UDRY (2010): “Learning about a new technology: Pineapple in Ghana,” *The American Economic Review*, 100, 35–69. 3.1
- 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. 1
- CRUZ, C., J. LABONNE, AND P. QUERUBIN (2017): “Politician Family Networks and Electoral Outcomes: Evidence from the Philippines,” *The American Economic Review*, 107, 3006–3037. 3.1
- DEGROOT, M. (1974): “Reaching a consensus,” *Journal of the American Statistical Association*, 69, 118–121. 1, 3
- DEMARZO, P., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions\*,” *Quarterly journal of economics*, 118, 909–968. 3

- ERDÖS, P. AND A. RÉNYI (1959): “On random graphs, I,” *Publicationes Mathematicae (Debrecen)*, 6, 290–297. 1, 2.4.1
- EYSTER, E. AND M. RABIN (2014): “Extensive imitation is irrational and harmful,” *The Quarterly Journal of Economics*, 129, 1861–1898. 1, 6
- FAFCHAMPS, M. AND F. GUBERT (2007): “The formation of risk sharing networks,” *Journal of Development Economics*, 83, 326–350. 2.4
- FELDMAN, M., N. IMMORLICA, B. LUCIER, AND S. M. WEINBERG (2014): “Reaching Consensus via Non-Bayesian Asynchronous Learning in Social Networks,” *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 192. 3, 1
- GALE, D. AND S. KARIV (2003): “Bayesian learning in social networks,” *Games and Economic Behavior*, 45, 329–346. 1, 3
- GEANAKOPOLOS, J. (1994): “Common knowledge,” *Handbook of game theory with economic applications*, 2, 1437–1496. B.1
- GOLUB, B. AND M. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–149. 3, A
- GRAHAM, B. S. (2017): “An econometric model of network formation with degree heterogeneity,” *Econometrica*, 85, 1033–1063. 2.4
- 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. 6, E
- JADBABAIE, A., P. MOLAVI, A. SANDRONI, AND A. TAHBAZ-SALEHI (2012): “Non-Bayesian Social Learning,” *Games and Economic Behavior*, 76, 210–225. 3
- LIGGETT, T. M. (1985): “Interacting particle systems, volume 276 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences],” . 1
- LOBEL, I. AND E. D. SADLER (2015): “Information diffusion in networks through social learning,” *Theoretical Economics*, 10, 807–851. 1
- MCCORMICK, T. H. AND T. ZHENG (2015): “Latent surface models for networks using Aggregated Relational Data,” *Journal of the American Statistical Association*, 110, 1684–1695. 2.4
- MENAGER, L. (2006): “Consensus, communication and knowledge: An extension with Bayesian agents,” *Mathematical Social Sciences*, 51, 274–279. 2.3

- MENDEL, F. AND V. GRIMM (2014): “An Experiment on Learning in a Multiple Games Environment,” . 1
- MÖBIUS, M., T. PHAN, AND A. SZEIDL (2015): “Treasure Hunt: Social Learning in the Field,” . 1
- MOSSEL, E., J. NEEMAN, AND O. TAMUZ (2014): “Majority dynamics and aggregation of information in social networks,” *Autonomous Agents and Multi-Agent Systems*, 28, 408–429. 1
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): “Strategic Learning and the Topology of Social Networks,” *Econometrica*, 83, 1755–1794. 1, 9, 2.3
- (Forthcoming): “Asymptotic learning on Bayesian social networks,” *Probability Theory and Related Fields*. 2.3, A
- MOSSEL, E. AND O. TAMUZ (2010): “Efficient Bayesian Learning in Social Networks with Gaussian Estimators,” *arXiv:1002.0747*. 1, 3
- MUELLER-FRANK, M. AND C. NERI (2013): “Social Learning in Networks: Theory and Experiments,” . 1
- OSBORNE, M. J. AND A. RUBINSTEIN (1994): *A course in game theory*, MIT press. B.1
- PENROSE, M. (2003): *Random geometric graphs*, 5, Oxford university press. 1, 2.4.1
- SMITH, L. AND P. SORENSEN (2000): “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398. 1

FIGURES




FIGURE 1. Contrast between DeGroot learning, where agents 2,3,6, and 1 remain stuck forever, and complete information Bayesian learning where, because all agents are Bayesian and this is commonly known, all agents converge to the truth (yellow in this example).




FIGURE 2. Network structures chosen for the experiment.

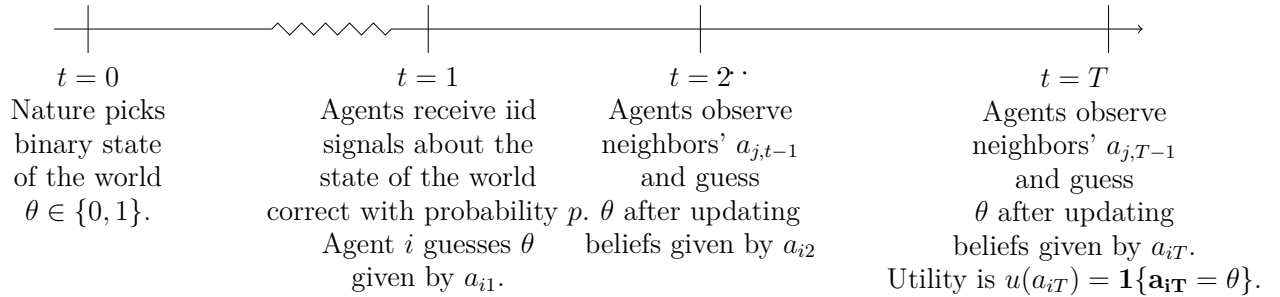


FIGURE 3. Timeline

## TABLES

TABLE 1. Reduced form patterns: India

<i>Panel A: Stuckness</i>		
(1)		
Share of clans that remain stuck on the wrong guess given that the Bayesian model would have predicted a change along the path (DeGroot predicts 1)		
VARIABLES		
Constant	0.946 (0.0303)	
Observations	74	
<i>Panel B: Information dominance</i>		
(1)		
Share of times information dominated agent fails to copy dominating agent (Complete Information Bayesian predicts 0)		
VARIABLES		
Constant	0.829 (0.0380)	
Observations	140	
<i>Panel C: Information revelation</i>		
(1)		(2)
Share of times an agent responds inefficiently to neighbor's actions (Bayesian predicts 0)		Share of times an agent necessarily learns the majority of signals and yet changes guess along path given that the DeGroot assessment would have changed guess (Bayesian predicts 0)
VARIABLES		
Constant	0.931 (0.0203)	0.945 (0.0249)
Observations	159	73

Notes: Standard errors are reported in parentheses. Panel A corresponds to the feature that, in DeGroot models, a clan that is stuck remains so until the end. Panel B is motivated by the fact that an agent should never respond to the behavior of someone whose information set is a subset under a Bayesian model, which is robust to incomplete information (column 2). Similarly, a Bayesian agent in a complete Bayesian world should only copy their information dominating neighbor and do nothing else (column 1). Panel C looks at the feature that, irrespective of whether agents are Bayesian or DeGroot, in round 2 they will play the majority, and therefore it is possible for Bayesian agents, even in an incomplete information world, to learn the majority of signals in certain cases and thus they should then stick to this guess.



TABLE 2. Reduced form patterns: Mexico

<i>Panel A: Stuckness</i>		
VARIABLES	(1) Share of clans that remain stuck on the wrong guess given that the Bayesian model would have predicted a change along the path (DeGroot predicts 1)	
Constant	0.303 (0.144)	
Observations	33	
<i>Panel B: Information dominance</i>		
VARIABLES	(1) Share of times information dominated agent fails to copy dominating agent (Complete Information Bayesian predicts 0)	
Constant	0.545 (0.0660)	
Observations	112	
<i>Panel C: Information revelation</i>		
VARIABLES	(1) Share of times an agent responds inefficiently to neighbor's actions (Bayesian predicts 0)	(2) Share of times an agent necessarily learns the majority of signals and yet changes guess along path given that the DeGroot assessment would have changed guess (Bayesian predicts 0)
Constant	0.614 (0.0862)	0.600 (0.117)
Observations	57	35

Notes: Standard errors are reported in parentheses. Panel A corresponds to the feature that, in DeGroot models, a clan that is stuck remains so until the end. Panel B is motivated by the fact that an agent should never respond to the behavior of someone whose information set is a subset under a Bayesian model, which is robust to incomplete information (column 2). Similarly, a Bayesian agent in a complete Bayesian world should only copy their information dominating neighbor and do nothing else (column 1). Panel C looks at the feature that, irrespective of whether agents are Bayesian or DeGroot, in round 2 they will play the majority, and therefore it is possible for Bayesian agents, even in an incomplete information world, to learn the majority of signals in certain cases and thus they should then stick to this guess.

TABLE 3. Structural estimates

<i>Panel A: <math>\epsilon</math></i>		
	(1)	(2)
	$\epsilon_{\text{India}}$	$\epsilon_{\text{Mexico}}$
	0.1288	0.134
	(0.007)	(0.013)
<i>Panel B: Network Level</i>		
	(1)	(2)
	$\pi_{\text{India}}$	$\pi_{\text{Mexico}}$
	0.1	0.5
	(0.130)	(0.184)
<i>Panel B: Individual Level</i>		
	(1)	(2)
	$\pi_{\text{India}}$	$\pi_{\text{Mexico}}$
	0.1	0.4
	(0.113)	(0.268)

Notes: Block-bootstrapped standard errors at the session level for  $\pi$  and at the agent level for  $\epsilon$  are reported in parentheses.

## APPENDIX A. PROOFS

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

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

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.  $\square$

*Proof of Theorem 1.* For (1), Golub and Jackson (2010) study a model where initial beliefs (or signals)  $p_{i,t}^{(n)} \in [0, 1]$  are independently distributed, with some common mean  $\mu = \mathbb{E} \left[ p_{i,0}^{(n)} \right]$  and common finite variance, and agents update their beliefs according to a DeGroot model with weighting matrix  $\mathbf{T}(n)$ . They show that, if  $\mathbf{T}(n)$  corresponds to an uniform DeGroot weighting model (as the one we use), then  $\text{plim}_{n \rightarrow \infty} \max_{i \leq n} |p_{i,\infty}^{(n)} - \mu| = 0$ , where  $p_{i,\infty}^{(n)} = \lim_{t \rightarrow \infty} p_{i,t}^{(n)}$  (see Corollary 2 in their paper). This corresponds to their definition of *wise sequences* of DeGroot weighting matrices (Definition 3). For our application, let us assume (without loss of generality) that  $\theta = 1$ , and  $p_{i,0} = \mathbb{P}(\theta = 1 \mid s_i)$ . Given the state, the distribution of  $p_{i,0}$  is binary (with values  $p$  and  $(1-p)$ , with probabilities  $p$  and  $(1-p)$ , respectively), independent and identically distributed across agents, with mean  $\mu = \mathbb{E} [p_{i,0} \mid \theta = 1] = p^2 + (1-p)^2$ . Therefore, if agents communicate their initial posterior beliefs (after observing their original signals) then  $\text{plim}_{n \rightarrow \infty} \max_{i \leq n} |p_{i,\infty}^{(n)} - \mu| = 0$ . Since  $p > 1/2$  we also have  $\mu > \frac{1}{2}$ , and hence whenever  $p_{i,\infty} \approx \mu$  then  $a_{i,\infty} = \theta$ , implying then that  $\lim_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{i \leq n} |a_{i,\infty}^{(n)} - \theta| \leq \epsilon \right\} = 0$  for all  $\epsilon > 0$ , a stronger result than the one we use. An analogous result works if  $\theta = 0$  (with  $\mu = 2p(1-p)$ ).

Result (2) on Bayesian action models is the central Theorem in Mossel et al. (Forthcoming).

For Result (3), the assumption that  $X_n \rightarrow x > 0$  implies that, for large enough  $n$ , there exist a number  $h_n \in \mathbb{N}$  of clans of  $k$  members who are disjoint; i.e., for every  $n$ , there exist sets  $\{C_{j,n}\}_{j=1}^{j=h_n} \subset V_n$  such that (a)  $\bigcap_{j=1}^{j=h_n} C_{j,n} = \emptyset$  and (b)  $C_{j,n}$  is a clan with

$|C_{j,n}| = k$  for all  $j$ . Moreover,  $h_n \rightarrow \infty$  (since each  $C_{j,n}$  has only  $k < \infty$  members). Define  $h = \liminf_{n \rightarrow \infty} h_n/n$ . In the incomplete information Bayesian model, a clan  $C \subseteq V$  of  $k$  members where every agent is a DeGroot type and every agent gets the wrong signal (i.e.  $\eta_i = D$  and  $s_i = 1 - \theta$  for all  $i \in C$ ) plays the wrong action forever, a corollary of Proposition 1. This happens with probability  $\alpha := (1 - p)^k (1 - \pi)^k > 0$ . Therefore, in the limit a fraction  $\alpha$  of all the disjoint clans satisfy this property, implying that at least a share  $\alpha h k$  of agents chooses  $a_{i,t} = 1 - \theta$  at every  $t \in \mathbb{N}$ , thus showing the desired result.  $\square$

*Proof of Proposition 4.* Consider any set  $S$  such that  $0 < \text{vol}(S) \leq \frac{1}{2} \text{vol}(V)$ . Observe

$$\frac{\partial S}{\text{vol}(S)} = \frac{\frac{\sum_{i \in S} d_i(V \setminus S)}{\sum_{i \in S} d_i(S)}}{1 + \frac{\sum_{i \in S} d_i(V \setminus S)}{\sum_{i \in S} d_i(S)}}.$$

Consider a clan  $C$  satisfying the above requirement. Because it is a clan,  $d_i(C) > d_i(V \setminus C)$  for every  $i$ . This immediately implies that

$$\frac{\frac{\sum_{i \in C} d_i(V \setminus C)}{\sum_{i \in C} d_i(C)}}{1 + \frac{\sum_{i \in C} d_i(V \setminus C)}{\sum_{i \in C} d_i(C)}} < \frac{1}{2}$$

and, therefore, since we are taking the minimum over all sets and  $C$  is just one such set,  $\phi(g) < \frac{1}{2}$ . By assumption the share of clans of bounded size is positive in the limit, and thus the clans satisfy the requirement above along the sequence of  $g_n$ . Thus, the result follows.  $\square$

*Proof of Corollary 1.* Since

$$\frac{\lambda_2(L)}{\sqrt{2}} \leq \phi(g) \leq \sqrt{2\lambda_2(L)}$$

if  $\frac{\lambda_2(L)}{\sqrt{2}} > \frac{1}{2}$  or  $\lambda_2(L) > \frac{\sqrt{2}}{2}$  then  $\phi(g) > \frac{1}{2}$  which proves the result.  $\square$

### A.1. Proofs of Propositions 5 and 6.

*Proof of Proposition 5.* Let  $\Omega_t$  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 B.4. To calculate them, we then have to loop across  $\#(\Omega_t) \times \#(\Omega_t)$  states for each  $(i, t)$  to update  $\sigma_t$  to  $\sigma_{t+1}$  and  $\#(\Omega_t)$  to update  $\alpha_t$ . 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

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

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 is approximately 16,258.

*Proof of Proposition 6.* 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 B.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

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

Therefore, the ratio between the two is approximately

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

and for the incomplete information model, the equivalent ratio is

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

□

## A.2. Proof of Theorem 2.

A.2.1. *Bounds for Finite Graphs.* For the given metric space  $(\Omega, d)$ , we denote  $B(i, r)$  to be the open ball centered at  $i \in \Omega$  with radius  $r > 0$ . The model is *Euclidean* if  $\Omega \subseteq \mathbb{R}^h$  is an open set and  $d(i, j) := \sqrt{\sum_{i=1}^k (x_i - y_i)^2}$ . The results in this section uses an Euclidean model with  $h = 2$  and uniform Poisson intensity;  $f(i) = 1$  for all  $i \in \Omega$ . However, all results are easily generalizable for any intensity function  $f$ , and non-Euclidean models (we clarify this below) with higher dimensions. For any measurable  $A \subseteq \Omega$  define the random variable  $n_A = \{\text{number of nodes } i \in A\}$ . The Poisson point process assumption implies that  $n_A \sim \text{Poisson}(\lambda\mu(A))$ , where  $\mu(\cdot)$  is the Borel measure over  $\mathbb{R}^h$ . For any potential node  $j \in \Omega$ , define  $d_j(A) := \{\text{number of links } j \text{ has with nodes } i \in A\}$ .  $d_j = d_j(\Omega)$  denotes the total number of links  $j$  has (i.e., its degree).

Define  $\nu := \mathbb{E}\{n_A\}$  with  $A = B(i, r)$  as the expected number of nodes in a “local neighborhood,” which is  $\nu := \lambda\pi r^2$  in the Euclidean model with  $h = 2$ .<sup>21</sup> Define also the volume of  $\Omega$  simply as its measure; i.e.,  $\text{vol}(\Omega) := \mu(\Omega)$ . It is also useful to define  $\omega := \text{vol}(\Omega) / \nu$ , so that the expected number of nodes on the graph can be expressed as  $\mathbb{E}[n_\Omega] = \lambda \text{vol}(\Omega) = \nu \times \omega$ .

A *local clan* is a non-trivial clan  $C \subset V$  (i.e., with  $\#C \geq 2$ ) where the probability of a link forming between any pair  $\{i, j\} \in C$  is  $\alpha$ . A necessary condition for  $C$  to be a local clan is that  $C \subset L := B(i, \frac{r}{2})$  for some  $i \in \Omega$ . With the above definitions,  $C \subseteq L$  is a local clan if  $\#C \geq 2$  and, for all  $j \in C$ ,  $d_j(L) \geq d_j(\Omega \setminus L)$ . The goal of this section is to provide lower and upper bounds for the event

$$\mathcal{B}_L := \{g = (V, E) : C = V \cap L \text{ is a local clan}\} = \left\{ g = (V, E) : \#C \geq 2 \text{ and } \bigwedge_{j \in C} \{d_j(L) \geq d_j(\Omega \setminus L)\} \right\}.$$

<sup>21</sup>If  $h > 2$ ,  $\nu := \lambda \times (R\sqrt{\pi})^h / \Gamma(1 + h/2)$

**PROPOSITION A.1.** Suppose  $\omega > \frac{9}{4}$  and take  $i \in \Omega$  such that  $B\left(i, \frac{3}{2}r\right) \subseteq \Omega$ , and let  $L = B\left(i, \frac{r}{2}\right)$ . Then,

(A.1)

$$\mathbb{P}\{g = (V, E) : C = V \cap L \text{ is a local clan}\} \geq \sum_{n=2}^{\infty} \left(\frac{\nu}{4}\right)^n \frac{e^{-\nu/4}}{n!} F^*(n-1)^n \times \alpha^{n(n-1)/2} > 0$$

where  $F^*(\cdot)$  is the cdf of a Poisson random variable  $d^*$  with expected value  $\mathbb{E}(d^*) = \left(2\alpha + \left(\omega - \frac{9}{4}\right)\beta\right) \times \nu$ . Moreover,

$$(A.2) \quad \mathbb{P}\{g = (V, E) : C = V \cap L \text{ is a local clan}\} \leq \sum_{d=1}^{\infty} \left(\alpha \frac{\nu}{4}\right)^d \frac{e^{-\alpha\nu/4}}{d!} \hat{F}(d)$$

where  $\hat{F}(\cdot)$  is the marginal cdf of  $d_j(\Omega \setminus L)$  for any  $j \in C$ , a Poisson distribution with  $\mathbb{E}[d_j(\Omega \setminus L)] = \left(\frac{3}{4}\alpha + (\omega - 1)\beta\right) \times \nu$

*Proof.* See Online Appendix D □

From Proposition A.1 we get simpler upper and lower bounds, which are useful when proving Theorem 2. Specifically, if  $\alpha\nu < 4$ , we can bound the probability of this event by

$$(A.3) \quad \left(\frac{\nu}{4}\right)^2 \frac{e^{-\nu/4}}{2} F^*(1)^2 \alpha \leq \mathbb{P}(\mathcal{B}_L) \leq e^{-\alpha\nu} \frac{\alpha\nu}{4 - \alpha\nu}.$$

This implies that, if  $\alpha \approx 0$ ,  $\mathbb{P}(\mathcal{B}_L) \approx 0$ , which we use in the next subsection.

**A.2.2. Sparsity and Asymptotics.** For a given mixed model, the degree of any given node is given by the random variable  $d_j = d_j(\Omega) = d(\Omega)$ . Since  $d_j = d_j(B(j, r)) + d_j(\Omega \setminus B(j, r))$  is a sum of independent Poisson random variables, so  $d_j$  is also Poisson, with expectation

$$\mathbb{E}(d_j) = \alpha\lambda\mu[B(j, r)] + \beta\lambda\{\mu(\Omega) - \mu[B(j, r)]\} = [\alpha + (\omega - 1)\beta] \times \nu.$$

We now consider a sequence of models  $(\Omega_k, \alpha_k, \beta_k)$  with  $\omega_k \rightarrow \infty$ . A sequence is sparse if  $\mathbb{E}(d_j) \rightarrow d_\infty < \infty$  as  $\omega \rightarrow \infty$ . For that to be the case, we then need that

$$\lim_{k \rightarrow \infty} [\alpha_k + (\omega_k - 1)\beta_k] \times \nu = d_\infty,$$

which can only happen if  $\beta_k = O\left(\omega_k^{-1}\right)$ ; i.e.,  $\beta_k/\omega_k \rightarrow \rho_\infty$ . We also look only for sequences with  $\alpha_k \rightarrow \alpha_\infty$ , so that  $d_\infty = \alpha_\infty + \rho_\infty$ .

In the next Proposition, we show the main result of this section, which has Theorem 2 as a direct corollary.

**PROPOSITION A.2.** Consider a sequence  $(\alpha_k, \beta_k)$ , where  $\omega_k \rightarrow \infty$ ,  $\beta_k/\omega_k \rightarrow \rho_\infty > 0$  and  $\alpha_k \rightarrow \alpha_\infty$ , and let  $L := B\left(i, \frac{3}{2}r\right) \subseteq \Omega_k$  for large enough  $k$ . Then

$$\lim_{k \rightarrow \infty} \mathbb{P}(g = (V, E) : \text{under } (\alpha_k, \beta_k) \text{ } C = V \cap L \text{ is a local clan}) \begin{cases} > 0 & \text{if } \alpha_\infty > 0 \\ = 0 & \text{if } \alpha_\infty = 0 \end{cases}.$$

*Proof.* Denote  $d_k^* = \text{Poisson} \left[ \left( 2\alpha + \left( \omega_k - \frac{9}{4} \right) \beta_k \right) \times \nu \right]$  with cdf  $F_k^*(\cdot)$ . Then  $F_k^*(d) \rightarrow_{k \rightarrow \infty} F_\infty^*(d)$  is the cdf of  $d_\infty^* = \text{Poisson} [(2\alpha_\infty + \rho_\infty) \times \nu]$ . Moreover, for a large enough  $k$  so that  $L \subseteq \Omega_k$ ,

$$\mathbb{P}(g = (V, E) : \text{under } (\alpha_k, \beta_k) \ C = V \cap L \text{ is a local clan}) \geq \left( \frac{\nu}{4} \right)^2 \frac{e^{-\nu/4}}{2} F_k^*(1)^2 \times \alpha_k,$$

so that

$$\lim_{k \rightarrow \infty} \mathbb{P}(g = (V, E) : \text{under } (\alpha_k, \beta_k) \ C = V \cap L \text{ is a local clan}) \geq \left( \frac{\nu}{4} \right)^2 \frac{e^{-\nu/4}}{2} F_\infty^*(1)^2 \times \alpha_\infty.$$

Since  $F_\infty^*(1) = [1 + (2\alpha_\infty + \rho_\infty) \nu] e^{-(2\alpha_\infty + \rho_\infty)\nu} > 0$ , this limit is strictly bigger than zero when  $\alpha_\infty > 0$ .

When  $\alpha_\infty = 0$ , we need to show that the upper bound A.3 is 0, showing that no local clans can appear in the limit. Expression A.3 implies that for any  $k : \alpha_k < \nu/4$  (which exists, since  $\alpha_k \rightarrow 0$ ), then

$$\lim_{k \rightarrow \infty} \mathbb{P}(g = (V, E) : \text{under } (\alpha_k, \beta_k) \ C = V \cap L \text{ is a local clan}) \leq \lim_{k \rightarrow \infty} e^{-\alpha_k \nu/4} \times \frac{\alpha_k \nu}{4 - \alpha_k \nu} = 0.$$

□

## APPENDIX B. BAYESIAN LEARNING ALGORITHM IN INCOMPLETE INFORMATION MODELS

In this appendix, we describe the algorithm for computing the actions in the complete and incomplete information Bayesian model.

**B.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 define these objects for our experimental setup, which we use to calculate the predicted behavior of Bayesian agents  $a_{i,t}(\omega)$ .

**B.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 incomplete information model, 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 an average of her neighbors' (and her previous) guesses. Both  $s_i$  and  $\eta_i$  are drawn i.i.d. across agents and types and signals are independent of each other as well. 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  $\rho(\omega \mid \theta)$ . Then

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

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

Let  $p_\theta = \mathbb{P}(s_i = 1 \mid \theta)$ . In our experiment,  $p_\theta = 5/7$  if  $\theta = 1$  and  $p_\theta = 2/7$  if  $\theta = 0$ .

**B.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  $\omega$ . At round  $t = 1$ , agent  $i$  only observes  $\omega_i$  out of state  $\omega$ , and hence, her information set is:

$$(B.2) \quad P_{i,1}(\omega) := \{\omega' \in \Omega : \omega'_i = \omega_i\},$$

In words, 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 initially choose to match their signal; i.e.,

$$(B.3) \quad a_{i,1}(\omega) := s_i.$$

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 (i.e., her information set) are

$$(B.4) \quad P_{i,t}(\omega) := \{\omega' \in \Omega : \omega'_i = \omega_i \text{ and } a_{j,s}(\omega') = a_{j,s}(\omega) \text{ for all } j \in N(i), s \leq t - 1\}.$$

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).<sup>22</sup> The round  $t$  action function  $a_{i,t}(\omega)$  is then given by:

$$(B.5) \quad a_{i,t}(\omega) := \begin{cases} \mathbf{1} \left\{ I_{i,t}(\omega) > \frac{1}{2} \right\} & \text{if } I_{i,t}(\omega) \neq \frac{1}{2} \\ a_{i,t-1}(\omega) & \text{if } I_{i,t}(\omega) = \frac{1}{2} \end{cases},$$

where  $I_{i,t}(\omega)$  is the ‘‘belief index’’ at state  $\omega$ , which depends on the agents’ type. If agent  $i$  is Bayesian (i.e., under the complete information model, or if  $\eta_i = 1$  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  $P_{i,t}(\omega)$ :

$$(B.6) \quad \mathbb{P}(\theta = 1 \mid P_{i,t}(\omega)) := \frac{\sum_{\omega' \in P_{i,t}(\omega)} \rho(\omega' \mid \theta = 1)}{\sum_{\omega' \in P_{i,t}(\omega)} \rho(\omega' \mid \theta = 1) + \rho(\omega' \mid \theta = 0)}$$

<sup>22</sup>We can also define  $P_{i,t}$  recursively, starting at  $P_{i,1}$  as in B.2, and for  $t \geq 1$  let  $P_{i,t}(\omega) := P_{i,t-1}(\omega) \cap \{\omega' \in \Omega : a_{j,t-1}(\omega') = a_{j,t-1}(\omega) \text{ for all } j \in N(i)\}$

When  $i$  is not Bayesian at  $\omega$ , then  $I_{i,t}(\omega) := \sum_{j=1}^n T_{ij} a_{j,t-1}(\omega)$ , where  $[T_{ij}]_{ij}$  are the DeGroot weights.

**B.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  $\omega$ .

**ALGORITHM 1.** Bayesian Learning Algorithm

**Inputs:**

- (1) An  $n$ -person network  $G = (V, E)$  with adjacency matrix  $A_{n \times n}$ ;
- (2) A row stochastic matrix of DeGroot weights  $T_{n \times n}$ ; and
- (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 B.1, for  $\theta \in \{0, 1\}$ ; and
- (3) Information functions  $P_{i,t}(\omega)$  and actions  $a_{i,t}(\omega)$  according to B.2 and B.3 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 B.4 and B.5, where  $I_{i,t}(\omega) = \mathbb{P}(\theta = 1 | P_{i,t}(\omega))$  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  $\theta$  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  $\theta$ . 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  $\theta$  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.

## APPENDIX C. CONSISTENCY OF STRUCTURAL ESTIMATION

**C.1. Setup.** There are  $V$  villages, each with  $n$  individuals who are arranged in a network. Our asymptotic sequence will take  $V \rightarrow \infty$ .<sup>23</sup>

<sup>23</sup>In what follows we use the terminology of our experiment in India but we could just as well have  $v$  index session with a total of  $V$  sessions.

Every network of  $n$  individuals will play a learning game as follows. Each of  $n$  individuals have a type (Bayesian or DeGroot), so  $\eta_i \in \{B, D\}$ . This type is drawn iid with probability  $\pi = \mathbb{P}(\eta_i = B)$  before the start of the game. This is commonly known by all agents.<sup>24</sup>

Our goal is to estimate  $\pi$  from the data generated in our experiment.

At time 0, there is a vector of binary signals  $s = (s_1, \dots, s_n)$  drawn iid conditional on the state ( $\theta \in \{0, 1\}$ ). Agents are trying to learn  $\theta$ . The signals are distributed

$$s_i = \begin{cases} \theta & \text{with probability } p = 5/7 \\ 1 - \theta & \text{with probability } 1 - p. \end{cases}$$

The agents are engaging in a learning task wherein in every period, given the history, they take their best guess about the state of the world (1 or 0). Agents observe all their own previous actions as well as those of their network neighbors from prior periods. The type space here is therefore the cross between agent type (Bayes or DeGroot) and signal endowment. Let  $\omega = (\eta, s)$ . Note that the most information an agent could theoretically use to assess the value of  $\theta$  is  $(s_1, \dots, s_n)$ .

In every period  $\tau$  there is an action taken by  $i$ ,  $a_{i\tau}^*$ . The type of the agent and the history determines the action. Given a history  $A^{t-1} = (a_{i\tau})_{i=1, \tau=1}^{n, t-1}$ , there is a prescribed action under the model of behavior which can depend on the agent's type  $\eta_i$ , the history of observed play, and the prior probability that an agent is Bayesian:<sup>25</sup>

$$a_{it}^* \left( A^{t-1}; \eta, \pi \right).$$

Then given the prescribed option, the observed data for the econometrician (and agents) is

$$a_{it} = \begin{cases} a_{it}^* & \text{with probability } 1 - \epsilon \\ 1 - a_{it}^* & \text{with probability } \epsilon \end{cases}$$

for any  $t = 2, \dots, T$ . Note that the history is the history of observed actions, which can differ from the prescribed action. We assume that this mistake is not internalized by agents. For the network level approach, we can take any  $T \geq 3$  whereas for the individual level approach, assume  $T = 3$ .<sup>26</sup>

<sup>24</sup>Note that DeGroot agents are mechanical and don't use this information so it really matters for Bayesian agents

<sup>25</sup>In the network level approach this is  $a_{it}^* \left( A^{t-1}; \eta, \pi \right) = a_{it}^* \left( (a_{i0})_{i=1}^n; \eta, \pi \right)$  and in the individual level approach this is  $a_{it}^* \left( A^{t-1}; \eta, \pi \right) = a_{it}^* \left( A_i^{t-1}; \eta_i, \pi \right)$ .

<sup>26</sup>As discussed in the text, we say that the model is defined until the first  $t$  at which some  $i$  encounters a zero probability information set, which we denoted as  $T^*$ . This cannot happen at  $T = 3$  so for simplicity consider this to be the case which defines a valid sample from which to construct a consistent estimator.

The matrix  $A_v^T = [a_{it,v}]$  is the dataset for a given village  $v$ . Suppressing  $v$  until it is needed, the likelihood is

$$\mathcal{L}(\pi, \epsilon; A^T) = \mathbb{P}(A^T | \pi, \epsilon) = \mathbb{P}(a_T | A^{T-1}, \pi, \epsilon) \cdot \mathbb{P}(a_{T-1} | A^{T-2}, \pi, \epsilon) \cdots \mathbb{P}(a_1 | \pi, \epsilon).$$

Notice that  $\mathbb{P}(a_1 | \pi)$  and  $\mathbb{P}(a_2 | \pi)$  are both independent of  $\pi$ , because they are independent of  $\eta$ : in period 1 every agent plays their signal and in period 2 every agent plays the majority (subject to a fixed tie breaking rule).

**C.2. Estimation of  $\epsilon$ .** Observe that for any graph  $v$  for any node  $i$  such that the majority of their neighbors and their own signal is unique, both the Bayes and DeGroot models, irrespective of  $\pi$ , prescribe the majority. Therefore, recalling that  $N_i^* = \{j : g_{ij} = 1\} \cup \{i\}$ ,

$$\hat{\epsilon} := \frac{\sum_v \sum_j \mathbf{1}\{a_{j2} \neq \text{majority}(a_{j1} : j \in N_i^*)\} \cdot \mathbf{1}\{\text{unique majority}(a_{j1} : j \in N_i^*)\}}{\sum_v \sum_j \mathbf{1}\{\text{unique majority}(a_{j1} : j \in N_i^*)\}}.$$

By standard arguments  $\hat{\epsilon} \xrightarrow{p} \epsilon$  and  $\sqrt{V} \frac{\hat{\epsilon} - \epsilon}{\widehat{\epsilon}(1 - \widehat{\epsilon})} \rightsquigarrow \mathcal{N}(0, 1)$ , since this is just a set of Bernoulli trials.

**C.3. Estimation of  $\pi$ .** For simplicity of exposition we take  $\epsilon$  as known, though in practice this will be a two-step estimator.

We can now consider

$$\mathcal{L}(\pi; A^T, \epsilon) = \prod_{t=3}^T \mathbb{P}(a_t | A^{t-1}, \pi, \epsilon).$$

It is useful to expand the term noting that  $A^1 = s$ ,

$$\begin{aligned} \mathbb{P}(a_t | A^{t-1}, \pi, \epsilon) &= \prod_{i=1}^n \mathbb{P}(a_{it} | A^{t-1}, \pi, \epsilon) \\ &= \prod_{i=1}^n \sum_{\eta} \mathbb{P}(a_{it} | A^{t-1}, \eta, \pi, \epsilon) \mathbb{P}(\eta | \pi) \end{aligned}$$

by independence and then

$$\begin{aligned} \mathbb{P}(a_{it} | A^{t-1}, \eta, \pi, \epsilon) &= \mathbf{1}\{a_{it} = a_{it}^*\} \mathbb{P}(a_{it} = a_{it}^* | a_{it}^* (A^{t-1}), A^{t-1}, \eta, \pi, \epsilon) \mathbb{P}(a_{it}^* | A^{t-1}, \eta, \pi, \epsilon) \\ &\quad + \mathbf{1}\{a_{it} \neq a_{it}^*\} \mathbb{P}(a_{it} \neq a_{it}^* | a_{it}^* (A^{t-1}), A^{t-1}, \eta, \pi, \epsilon) \mathbb{P}(a_{it}^* | A^{t-1}, \eta, \pi, \epsilon) \\ &= \mathbf{1}\{a_{it} = a_{it}^* (A^{t-1}; \eta, \pi)\} \cdot (1 - \epsilon) + \mathbf{1}\{a_{it} \neq a_{it}^* (A^{t-1}; \eta, \pi)\} \cdot \epsilon. \end{aligned}$$

Let  $x_{it} = \mathbf{1}\{a_{it} = a_{it}^* (A^{t-1}; \eta, \pi)\}$ , which computes whether the observed action matches that which was prescribed by the model given the history, type vector, and parameter value.

So,<sup>27</sup>

$$\mathbb{P}\left(a_t|A^{t-1}, \pi, \epsilon\right) = \prod_{i=1}^n \sum_{\eta} (1 - \epsilon)^{x_{it}} \epsilon^{1-x_{it}} \mathbb{P}(\eta|\pi).$$

Recalling  $x_{it} = x_{it}[A^{t-1}; \eta, \pi]$ , we can consider the log likelihood for a given  $v$ ,

$$\ell_v(\pi; A^T, \epsilon) = \sum_{t=3}^T \sum_{i=1}^n \log \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) \right\}.$$

And, since villages are independent, the full log likelihood is

$$\ell(\pi; A^T, \epsilon) = \sum_{v=1}^V \sum_{t=3}^T \sum_{i=1}^n \log \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) \right\}.$$

Then, let us define

$$\log f(A_v|\pi) := \sum_{t=3}^T \sum_{i=1}^n \log \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) \right\}.$$

**C.4. Consistency of  $\hat{\pi}$ .** In what follows, for simplicity assume that  $\epsilon$  is known. Standard arguments will allow us to extend the below to joint consistency. Now we demonstrate that

$$\hat{\pi} := \operatorname{argmax}_{\pi} \hat{Q}_V(\pi) := \operatorname{argmax}_{\pi} \frac{1}{V} \sum_v \log f(A_v|\pi).$$

generates a consistent estimator of  $\pi$ . The limit object is  $Q_0(\pi) := \mathbb{E}[\log f(A_v|\pi)]$ .

**PROPOSITION C.1.** Under the above assumptions,  $\hat{\pi} \rightarrow_p \pi_0$  as  $V \rightarrow \infty$ .

*Proof.* This serves only as a sketch, but follows the arguments of Theorem 2.1 in Newey and McFadden (1994). First, by the arguments of Lemma 2.2 in Newey and McFadden (1994), there is a unique maximum of  $Q_0(\pi)$  at the true value  $\pi_0$ , since

$$Q_0(\pi_0) - Q_0(\pi) = \mathbb{E}_{\pi_0} \left[ -\log \frac{f(A_v|\pi)}{f(A_v|\pi_0)} \right] > -\log \mathbb{E}_{\pi_0} \left[ \frac{f(A_v|\pi)}{f(A_v|\pi_0)} \right] = 0$$

by the information inequality.

Second, we can take compactness as given since  $\pi \in [0, 1]$ .

Third, the objective is continuous in  $\pi$  with probability one. To see this, notice that  $\mathbb{P}(\eta|\pi)$  is continuous in the parameter since it consists of binomial draws with probability  $\pi$ . Further,  $x_{it}[A^{t-1}; \eta, \pi]$  is continuous a.e. in  $\pi$  because it is a step function.

Lastly, we need to establish that the finite sample objective function converges uniformly in probability to its limit. To show that, we argue that  $\hat{Q}_V(\pi) := \frac{1}{V} \sum_v \log f(A_v|\pi)$  is

<sup>27</sup>It is worth noting that if we could pass the logarithm then this is

$$\sum_{t=3}^T \sum_{i=1}^n \sum_{\eta} (x_{it}[A^{t-1}; \eta, \pi] \log(1 - \epsilon) + (1 - x_{it}[A^{t-1}; \eta, \pi]) \log \epsilon) \cdot \mathbb{P}(\eta|\pi)$$

and for small  $\epsilon$  this is a reweighted divergence.

stochastically equicontinuous and converges pointwise. Pointwise convergence is self-evident. To show stochastic equicontinuity, we check the Holder inequality which is a sufficient condition. Consider any two  $\pi$  so we that have

$$\log \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) \right\} - \log \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi']} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi']} \cdot \mathbb{P}(\eta|\pi') \right\},$$

which is

$$\begin{aligned} & \left| \log \left\{ \frac{\sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi)}{\sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi']} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi']} \cdot \mathbb{P}(\eta|\pi')} \right\} \right| \\ & \leq 0 + |\pi - \pi'| \left| \frac{\frac{\partial}{\partial \pi} \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) \right\}}{\sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi)} \right|. \end{aligned}$$

Then,

$$\begin{aligned} & \frac{\partial}{\partial \pi} \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) \right\} \\ & = \sum_{\eta} \left[ \frac{\partial}{\partial \pi} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \right] \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \mathbb{P}(\eta|\pi) + \left\{ \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \left[ \frac{\partial}{\partial \pi} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \right] \cdot \mathbb{P}(\eta|\pi) \right\} \\ & + \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \frac{\partial \mathbb{P}(\eta|\pi)}{\mathbb{P}(\eta|\pi)} \cdot \mathbb{P}(\eta|\pi). \end{aligned}$$

Then, the first two terms are 0 a.e. and therefore certainly bounded by 1, and the final term is just

$$\mathbb{E} \left[ (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \text{Score}(\eta|\pi) \right] \leq n2^n,$$

which is a constant since  $n$  is fixed. This follows from

$$\begin{aligned} & \mathbb{E} \left[ (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \text{Score}(\eta|\pi) \right] \\ & = \sum_{\eta} (1 - \epsilon)^{x_{it}[A^{t-1}; \eta, \pi]} \epsilon^{1-x_{it}[A^{t-1}; \eta, \pi]} \cdot \left( z\pi^{z-1} (1 - \pi)^{n-z} + (n - z)\pi^z (1 - \pi)^{n-z-1} \right) \leq n2^n. \end{aligned}$$

So, we have a parameter-independent bound that satisfies the Holder condition.  $\square$

**C.5. Simulations.** We now show that, if we generate data with parameters  $(\pi, \epsilon)$ , we can use our estimator to recover both parameters. Figure 4 shows the results. We have generated data with  $\epsilon = 0.13$  (the level estimated in both datasets) and  $\pi \in \{0, 0.1, \dots, 0.9, 1\}$ . We show that across the board the objective function is maximized exactly at the right parameter value both in the network- and individual-level estimations.

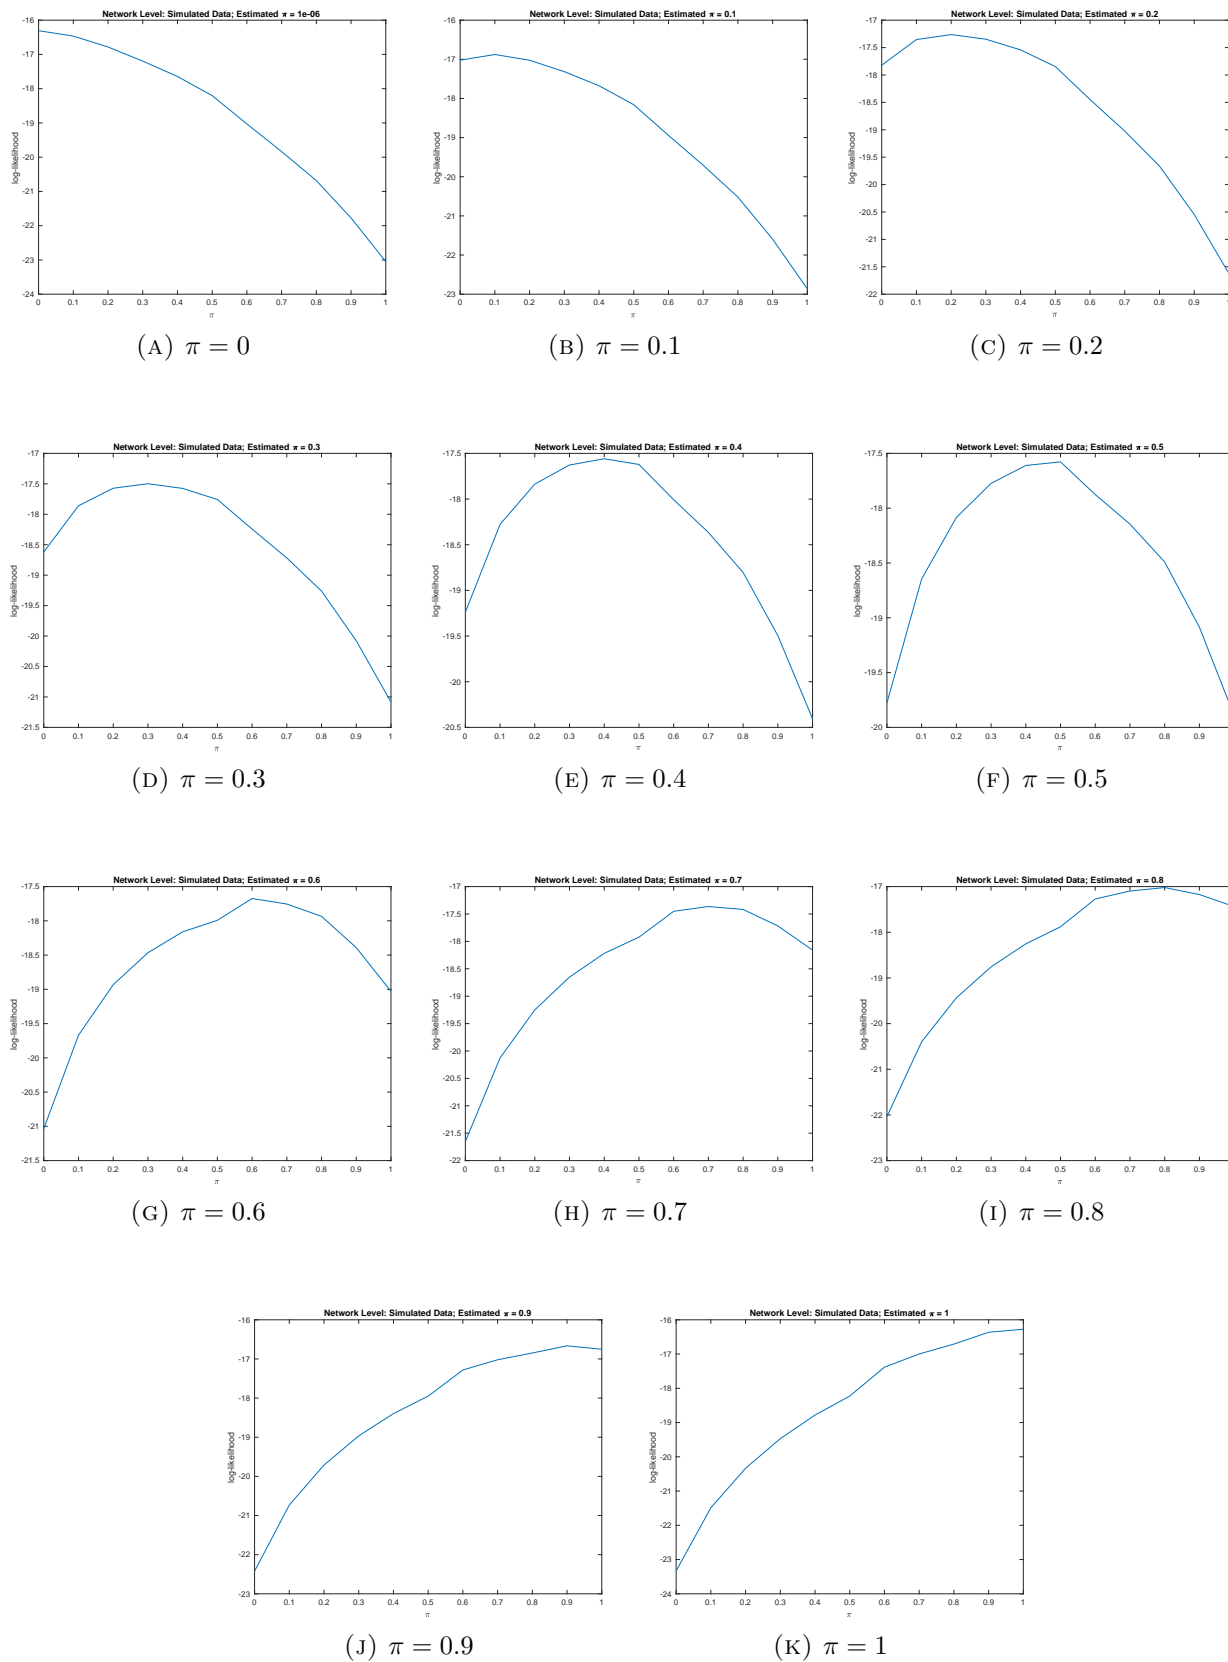


FIGURE 4. Objective functions for MLEs of  $\pi$  for simulated data generated at various  $\pi$  (Network-level estimation).

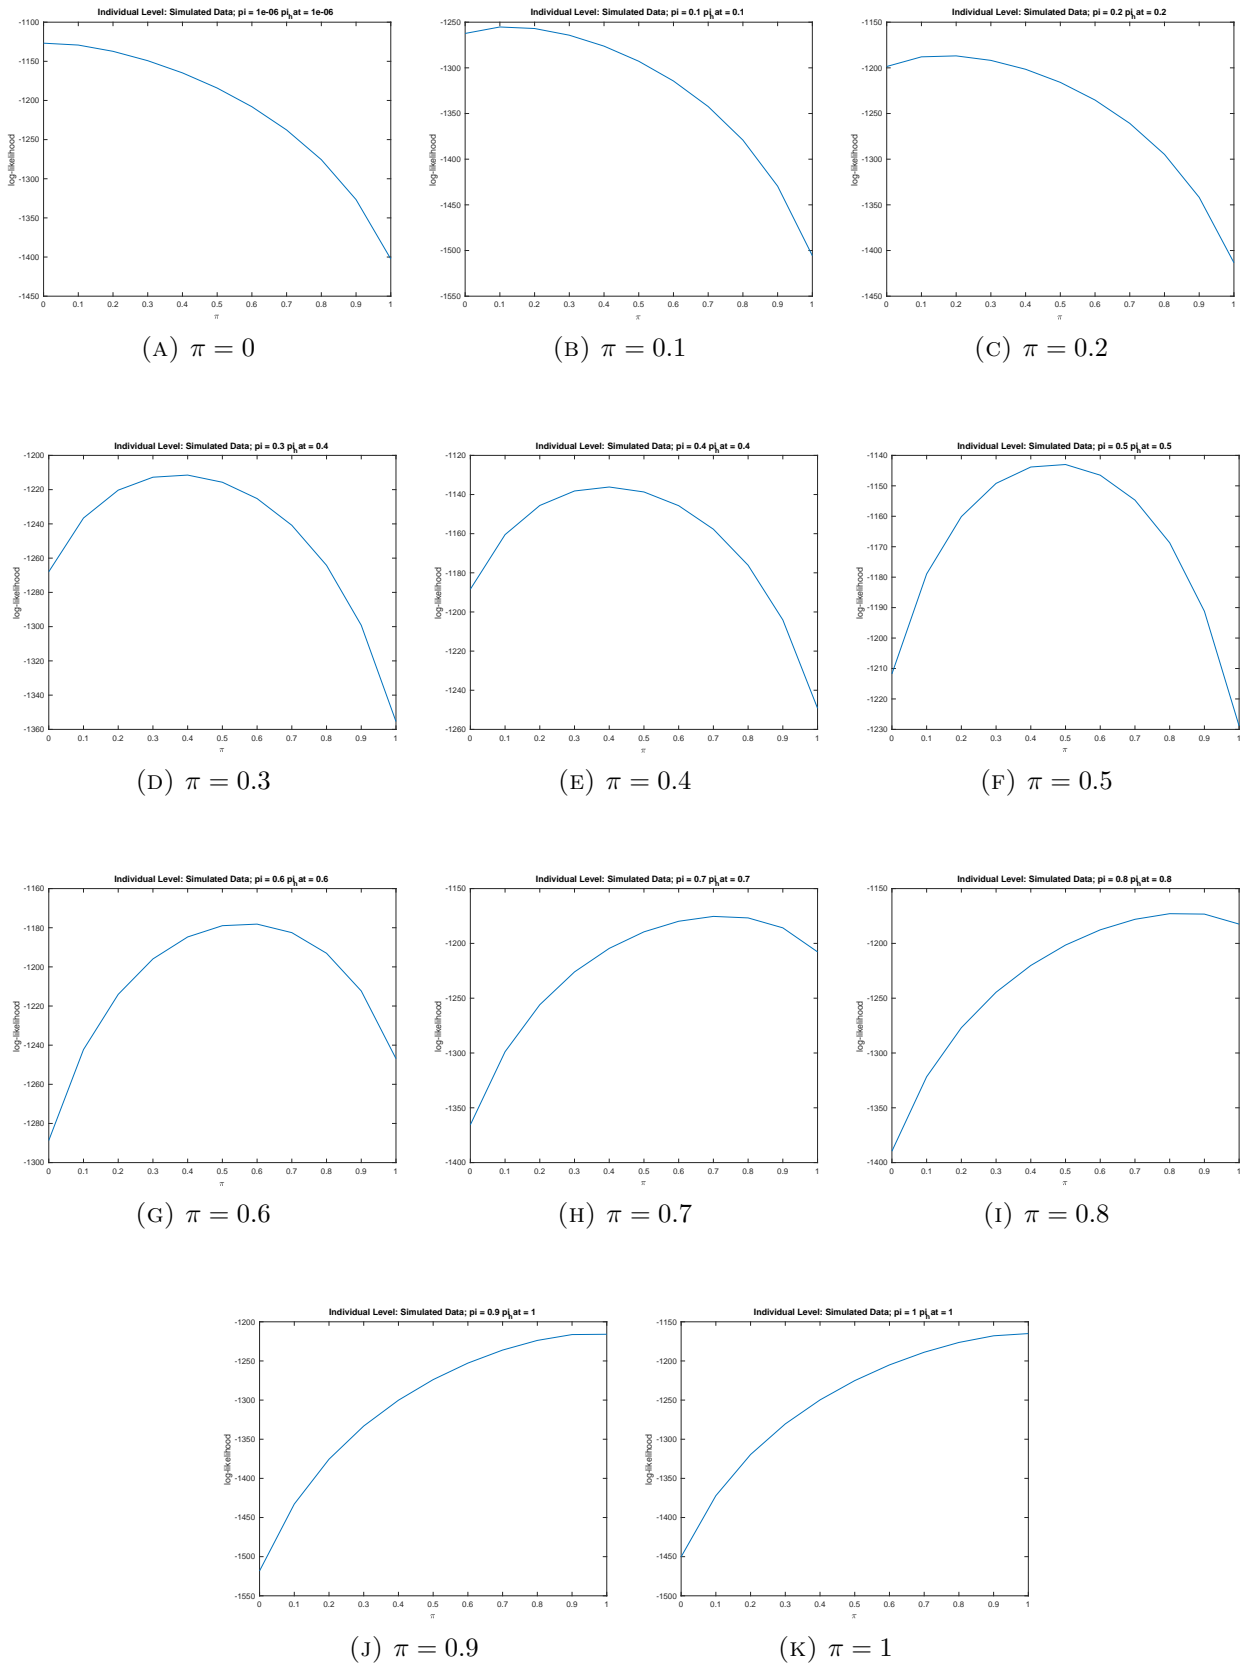


FIGURE 5. Objective functions for MLEs of  $\pi$  for simulated data generated at various  $\pi$  (Individual-level estimation).



ONLINE APPENDIX:  
NOT FOR PUBLICATION

APPENDIX D. RGG-ER MIXTURES

For the given metric space  $(\Omega, d)$ , we denote  $B(i, r)$  to be the open ball centered at  $i \in \Omega$  with radius  $r > 0$ . The model is *Euclidean* if  $\Omega \subseteq \mathbb{R}^h$  is an open set and  $d(i, j) := \sqrt{\sum_{i=1}^k (x_i - y_i)^2}$ . The results in this section uses an Euclidean model with  $h = 2$  and uniform Poisson intensity;  $f(i) = 1$  for all  $i \in \Omega$ . However, all results are easily generalizable for any intensity function  $f$ , and non-Euclidean models (we clarify this below) with higher dimensions. For any measurable  $A \subseteq \Omega$  define the random variable  $n_A = \{\text{number of nodes } i \in A\}$ . The Poisson point process assumption implies that  $n_A \sim \text{Poisson}(\lambda\mu(A))$ , where  $\mu(\cdot)$  is the Borel measure over  $\mathbb{R}^h$ . For any potential node  $j \in \Omega$ , define  $d_j(A) := \{\text{number of links } j \text{ has with nodes } i \in A\}$ .  $d_j = d_j(\Omega)$  denotes the total number of links  $j$  has (i.e., its degree).

Define  $\nu := \mathbb{E}\{n_A\}$  with  $A = B(i, r)$  as the expected number of nodes in a “local neighborhood,” which is  $\nu := \lambda\pi r^2$  in the Euclidean model with  $h = 2$ .<sup>28</sup> Define also the volume of  $\Omega$  simply as its measure; i.e.,  $\text{vol}(\Omega) := \mu(\Omega)$ . It is also useful to define  $\omega := \text{vol}(\Omega)/\nu$ , so that the expected number of nodes on the graph can be expressed as  $\mathbb{E}[n_\Omega] = \lambda\text{vol}(\Omega) = \nu \times \omega$ .

A *local clan* is a non-trivial clan  $C \subset V$  (i.e., with  $\#C \geq 2$ ) where the probability of a link forming between any pair  $\{i, j\} \in C$  is  $\alpha$ . A necessary condition for  $C$  to be a local clan is that  $C \subset L := B(i, \frac{r}{2})$  for some  $i \in \Omega$ . With the above definitions,  $C \subseteq L$  is a local clan if  $\#C \geq 2$  and, for all  $j \in C$ ,  $d_j(L) \geq d_j(\Omega \setminus L)$ .

For a given  $j \in L$ , and a given number of nodes  $N_L$  in  $L$ , the number of links  $j$  has to other nodes in  $L$  is distributed Binomial( $\alpha, N_L$ ). Since  $N_L \sim \text{Poisson}(\lambda\mu(L))$  then  $d_j(L) \sim \text{Poisson}(\alpha \times \nu(L))$ , where  $\nu(L) = \lambda\mu[B(i, \frac{r}{2})] = \lambda\pi r^2/4$ .<sup>29</sup>

Define  $H_j := B(j, r) \setminus L$  and  $K_j := \Omega \setminus H_j$ . Because of the assumptions on the mixed RGG model,  $\mathbb{P}(j \text{ has link with given } i \in L) = \mathbb{P}(j \text{ has link with } i \in H_j) = \alpha$ , and  $\mathbb{P}(j \text{ has link with given } i \in K_j) = \beta$ . Also, because these sets are disjoint and  $\Omega = L \cup H_j \cup K_j$ ,  $d_j = d_j(L) + d_j(H_j) + d_j(K_j)$ , and  $d_j(L)$ ,  $d_j(H_j)$  and  $d_j(K_j)$  are independent Poisson random variables, with distributions

$$(D.1) \quad d_j(L) =_d d(L) \sim \text{Poisson}\left(\alpha \times \frac{\nu}{4}\right),$$

$$(D.2) \quad d_j(H_j) =_d d(H_j) \sim \text{Poisson}\left(\alpha \times \frac{3}{4}\nu\right), \text{ and}$$

<sup>28</sup>If  $h > 2$ ,  $\nu := \lambda \times (R\sqrt{\pi})^h / \Gamma(1 + h/2)$

<sup>29</sup>If  $h \geq 2$ , we have that in the euclidean model,  $\lambda\mu(L) = 2^{-h}\nu$ .

$$(D.3) \quad d_j(K_j) =_d d(K_j) \sim \text{Poisson}(\beta \times (\omega - 1) \nu),$$

where “ $=_d$ ” stands for equality in distribution, using that **(a)** uniform intensity implies  $d_j(A) = d(A)$  for any  $A$ , **(b)**  $N_{H_j} \sim \text{Poisson}(\lambda\mu(H_j))$  and  $N_{K_j} \sim \text{Poisson}(\lambda\mu(K_j))$ , where  $\mu(H_j) = \mu(B(j, r)) - \mu(B(i, r/2)) = (3/4)\pi r^2$  and  $\mu(K_j) = \text{vol}(\Omega) - [\mu(H_j) + \mu(L)] = (\omega - 1)\pi r^2$ . It is useful to work with the random variable

$$(D.4) \quad d_j(\Omega \setminus L) := d_j(H_j) + d_j(K_j) =_d d(\Omega \setminus L) \sim \text{Poisson} \left[ \nu \left( \frac{3}{4}\alpha + (\omega - 1)\beta \right) \right].$$

The goal of this section is to provide lower and upper bounds for the event

$$(D.5) \quad \mathcal{B}_L := \{g = (V, E) : C = V \cap L \text{ is a local clan}\} = \left\{ g = (V, E) : \#C \geq 2 \text{ and } \bigwedge_{j \in C} \{d_j(L) \geq d_j(\Omega \setminus L)\} \right\}$$

the problem of course being that, even though  $d_j(L)$  and  $d_j(\Omega \setminus L)$  are independent, the same variables across agents  $j \in C$  may very well not be (and usually will not be).

Given  $i \in \Omega$  and  $\hat{r} > r$ , an *annulus*  $\mathbf{An}(i, r, \hat{r}) \subset \Omega$  is the ring between the outer ball for radius  $r'$  and the inner ball with radius  $r$ , i.e.,

$$\mathbf{An}(i, r, \hat{r}) := \{j \in \Omega : r \leq d(i, j) < r'\} = B(i, \hat{r}) \setminus B(i, r).$$

The most important fact to keep in mind for Proposition A.1 is that, in the Euclidean model, the distributions of  $d_j(L)$ ,  $d_j(G_j)$  and  $d_j(H_j)$  given by equations D.1, D.2 and D.3 **do not depend on the chosen node  $j$** . This is the key property that allows us to obtain bounds on the probability of the existence of clans that do not depend on the particular nodes drawn in  $L$ .

**Proof of Proposition A.1**. We develop some notation. We denote  $d_j(A) \mid V$  as the number of nodes  $j$  has with nodes  $i \in V \cap A$ , conditional on a realized, finite set of nodes  $V \subseteq \Omega$ . Also, if  $X, Y$  are random variables, we use “ $X \supseteq Y$ ” to denote first order stochastic dominance of  $X$ . Let  $H^* = \mathbf{An}(i, \frac{r}{2}, \frac{3}{2}r)$  and  $K^* := \Omega \setminus \{L \cup H^*\}$ . Conditional on a realization of  $V$ , define  $d^* \mid V$  as the the number of links that a potential node would have if it has a probability  $\alpha$  of forming links with nodes in  $H^*$  (i.i.d across nodes in  $H^*$ ), and a probability  $\beta$  of forming links with nodes in  $K^*$  (again, i.i.d across nodes in  $K^*$ ). This can be summarized as

$$(D.6) \quad d^* \mid V = \text{Binomial}(\alpha, n_{H^*}) + \text{Binomial}(\beta, n_{K^*}),$$

where  $n_A := \#\{V \cap A\}$  is the number of realized nodes in set  $A \subseteq \Omega$ .<sup>30</sup> Equation D.6 also implies that, integrating over  $V$ , we get that  $d^* \sim \text{Poisson}(\alpha\nu(H^*) + \beta\nu(K^*))$ . This

<sup>30</sup>We use the convention that, if  $n_A = 0$ , then  $\text{Binomial}(\gamma, n_A) = 0$  with probability 1, for any  $\gamma \in [0, 1]$ .




FIGURE 6.  $H_J = B(J, r) \setminus B(i, \frac{r}{2})$ ;  $H^* = B(i, \frac{3r}{2}) \setminus B(i, \frac{r}{2})$ ;  $L = B(i, \frac{r}{2})$

implies

$$\begin{aligned} \mathbb{E}(d^*) &= \alpha\nu(H^*) + \beta\nu(K^*) = \lambda \left\{ \alpha\mu\left(\mathbf{A}\left(i, \frac{r}{2}, \frac{3}{2}r\right)\right) + \beta\left(\mu[\Omega] - \mu(L) + \mu\left(\mathbf{A}\left(i, \frac{r}{2}, \frac{3}{2}r\right)\right)\right) \right\} = \\ &= \lambda \left\{ \alpha\left(\frac{9}{4} - \frac{1}{4}\right)\pi r^2 + \beta\left(\omega\pi r^2 - \frac{1}{4}\pi r^2 - \left(\frac{9}{4} - \frac{1}{4}\right)\pi r^2\right) \right\} = \\ &= \lambda\pi r^2 \left[ 2\alpha + \beta\left(\omega - \frac{9}{4}\right) \right] = \left[ 2\alpha + \beta\left(\omega - \frac{9}{4}\right) \right] \times \nu \end{aligned}$$

using that  $\mu[\mathbf{A}(i, r, \hat{r})] = (\hat{r}^2 - r^2)\pi$  and  $\mu(\Omega) = \omega\lambda\pi r^2$  by the definition of  $\omega$ , as we have seen above.

We first show the lower bound A.1 in 5 steps.

**Step 1:** For any  $j \in L$ ,  $H_j \subseteq H^*$ .

To show this, first we show that  $B(j, r) \subseteq B(i, \frac{3}{2}r)$ . Take  $x \in B(j, r)$ , so that  $d(j, x) < r$ . Then  $d(x, i) < d(x, j) + d(j, i) < r + \frac{1}{2}r = \frac{3}{2}r$  using that  $j \in B(i, \frac{r}{2})$ . Then  $H_j = B(j, r) \setminus B(i, r/2) \subseteq B(i, \frac{3}{2}r) \setminus B(i, r/2) = A^*$ , as we wanted to show.

**Step 2:** For any realization of  $V$ , have  $d^*|_V \supseteq d_j(\Omega \setminus L)|_V$  for all  $j \in C = V \cap L$ .

We provide a more heuristic proof for this statement. Define  $K^* = \Omega \setminus \{L \cup H^*\}$ . Because for all  $j \in C$ ,  $H_j \subseteq H^*$  and also  $K_j \supseteq K^*$ . Defining  $Z_j := H^* \setminus H_j$  we can decompose  $\Omega \setminus L$

as

$$\Omega \setminus L = H_j \cup Z_j \cup K^*,$$

which are disjoint sets. Now, according to the RGG model, a node  $j$  has a probability  $\alpha$  to make a link with any node in  $H_j$  since  $H_j = B(j, r) \setminus L$ , but has probability  $\beta \leq \alpha$  to make a link with nodes in  $Z_j \cup K^*$ . Therefore, conditional on  $V$ ,

$$d_j(\Omega \setminus L) \mid V \sim \text{Binomial}(\alpha, n_{H_j}) + \text{Binomial}(\beta, n_{Z_j}) + \text{Binomial}(\beta, n_{K^*}),$$

where we use the fact that  $Z_j \cup K^* = K_j$  and hence  $j$  has probability  $\beta$  of making successful links there. Meanwhile, for  $d^*$ ,

$$d^* \mid V \sim \text{Binomial}(\alpha, n_{H_j}) + \text{Binomial}(\alpha, n_{H_j}) + \text{Binomial}(\beta, n_{K^*}),$$

since  $Z_j \subseteq H^*$ . Therefore, since  $\alpha \geq \beta$ ,  $d^* \mid V \supseteq d_j(\Omega \setminus L) \mid V$

**Step 3:** Suppose we condition on the realized subgraph  $g_C = (C, E_C)$ . Then,

$$\mathbb{P}(\mathcal{B}_L \mid g_C) \geq \prod_{j \in C} F^*[d_j(L)],$$

where  $F^*(\cdot)$  is the cdf of  $d^*$ .

Given  $g_C$ , the in-degrees  $d_j(L)$  are known values.  $y_j := d_j(\Omega \setminus L)$  are independent random variables, conditional on the realization of  $C$ , since **(a)** they are independent conditional on  $V$  and **(b)** the realization of nodes in  $\Omega \setminus L$  is independent of  $g_C$ . Therefore, conditioning on both  $g_C = (C, E_C)$  and  $V \setminus C$  (i.e., taking expectations over the links with nodes in  $V \setminus C$ ),

$$\begin{aligned} \text{(D.7)} \quad \mathbb{P} \left\{ g : \bigwedge_{j \in C} \{d_j(L) \geq d_j(\Omega \setminus L)\} \mid g_C, V \setminus C \right\} &= \prod_{j \in C} \mathbb{P}[g : d_j(\Omega \setminus L) \leq d_j(L) \mid V \setminus C, g_C] \geq \\ &\geq \prod_{j \in C} \mathbb{P}[d^* \leq d_j(L) \mid V], \end{aligned}$$

where we use the fact that  $d^* \mid V \supseteq d_j(\Omega \setminus L) \mid V$  for all  $j \in C$ . Moreover,  $(n_{H^*}, n_{K^*})$  are sufficient statistics for the conditional distribution of  $d^*$ , and hence  $d^* \mid V$  is also independent of  $C$  and of  $d_j(L)$  for all  $j \in C$  (and hence  $d^* \mid V = d^* \mid (V \setminus C)$ ). Therefore, using D.7 and taking expectations over  $V \setminus C$ ,

$$\mathbb{P}(\mathcal{B}_L \mid g_C) = \mathbb{E}_{V \setminus C} \left\{ \mathbb{P} \left\{ g : \bigwedge_{j \in C} \{d_j(L) \geq d_j(\Omega \setminus L)\} \mid g_C, V \setminus C \right\} \right\} \geq$$

$$\begin{aligned} \mathbb{E}_{V \setminus C} \left\{ \prod_{j \in C} \mathbb{P}[g : d_j(\Omega \setminus L) \leq d_j(L) \mid V \setminus C, g_C] \right\} &= \prod_{j \in C} \mathbb{E}_{V \setminus C} \{ \mathbb{P}[g : d_j(\Omega \setminus L) \leq d_j(L) \mid V \setminus C, g_C] \} \geq \\ \mathbb{E}_{V \setminus C} \left\{ \prod_{j \in C} \mathbb{P}[d^* \leq d_j(L) \mid V \setminus C] \right\} &= \prod_{j \in C} \mathbb{E}_{(n_{H^*}, n_{K^*})} [\mathbb{P}(d^* \leq d_j(L) \mid n_{H^*}, n_{K^*})] = \\ &= \prod_{j \in C} F^*[d_j(L)], \end{aligned}$$

where we use the independence of  $y_j$  conditional on  $g_C$ , the fact that  $(n_{H^*}, n_{K^*})$  are sufficient statistics for  $d^* \mid (V \setminus C)$ , and that  $F^*$  is the cdf of the unconditional Poisson distribution of  $d^*$  that we derived above.

**Step 4:** Given  $n_L = \#\{C\} \geq 2$ , we have that

$$\mathbb{P}(\mathcal{B}_L \mid n_L) \geq F^*(n_L - 1)^{n_L} \times \alpha^{n_L(n_L - 1)/2}$$

Given  $n_L$ , we want to get a lower bound on the probability that, for  $n_L$  independent random draws  $d_j^* \sim \text{Poisson} \left[ \left( 2\alpha + \left( \omega - \frac{9}{4} \right) \beta \right) \times \nu \right]$ , we have that  $d_j(L) \geq d_j^*$  for all  $j \in C$ . One of these potential subgraphs is a clique, where  $d_j(C) = n_L - 1$  for all  $j \in C$ . Since  $g_C \mid n_L$  is an Erdos-Renyi graph with parameter  $p = \alpha$ , we know that the probability that

$g_C$  is a clique is  $\alpha \binom{n_L}{2} = \alpha^{n_L(n_L - 1)/2}$ . Therefore,

$$\begin{aligned} \mathbb{P} \left\{ g : \bigwedge_{j \in C} \{d_j(L) \geq d_j(\Omega \setminus L)\} \mid n_L \right\} &= \sum_{g_C: \#C = n_L} \mathbb{P} \left\{ g : \bigwedge_{j \in C} \{d_j(L) \geq d_j(\Omega \setminus L)\} \mid g_C \right\} \times \mathbb{P}(g_C \mid n_L) \geq \\ &= \sum_{g_C: \#C = n_L} \prod_{j \in C} F^*[d_j(L)] \times \mathbb{P}(g_C \mid n_L) \geq \left[ \prod_{j \in C} F^*(n_L - 1) \right] \times \alpha^{n_L(n_L - 1)/2} = \\ &= F^*(n_L - 1)^{n_L} \times \alpha^{n_L(n_L - 1)/2} \end{aligned}$$

**Step 5:**  $\mathbb{P}(\mathcal{B}_L) \geq \sum_{n=2}^{\infty} \left( \frac{\nu}{4} \right)^n \frac{e^{-\nu/4}}{n!} F^*(n - 1)^n \times \alpha^{n(n-1)/2}$ .

The previous result implies  $\mathbb{P}(\mathcal{B}_L) \geq \mathbb{P}(n_L \geq 2) \times \mathbb{E}_{n_L} \left\{ F^*(n_L - 1)^{n_L} \times \alpha^{n_L(n_L - 1)/2} \mid n_L \geq 1 \right\}$ . The fact that  $n_L \sim \text{Poisson}(\nu/4)$  gives us the desired expression.

**Step 6:**  $\mathbb{P}(\mathcal{B}_L) \leq \sum_{d=1}^{\infty} \left(\alpha \frac{\nu}{4}\right)^d \frac{e^{-\alpha\nu/4}}{d!} \hat{F}(d).$

The upper bound comes simply from the fact that the event  $\mathcal{B}_L \subseteq \{d_j(L) \geq d_j(\Omega \setminus L)\}$  for a particular  $j \in C$ . Since  $d_j(A) =_d d(A)$  for any  $j \in \Omega$  and any measurable  $A \subseteq \Omega$  we obtain the upper bound using D.4. For it to be a local clan, we need  $n_L \geq 2$  which necessarily implies that  $d_j(L) \geq 1$  for any  $j \in C$ . Moreover, since  $d_j(\Omega \setminus L)$  is independent of both  $n_L$  and  $d_j(L)$ , we get that

$$\begin{aligned} \mathbb{P}(\mathcal{B}_L) &\leq \mathbb{P}\{d_j(\Omega \setminus L) \leq d_j(L) \mid d_j(L) \geq 1\} = \mathbb{P}\left[\hat{F}(d_j(L)) \mid d_j(L) \geq 1\right] = \\ &\sum_{d=1}^{\infty} \left(\alpha \frac{\nu}{4}\right)^d \frac{e^{-\alpha\nu/4}}{d!} \hat{F}(d) \end{aligned}$$

where  $\hat{F}$  is the cdf of  $d_j(\Omega \setminus L) \sim \text{Poisson}\left[\left(\frac{3}{4}\alpha + (\omega - 1)\beta\right) \times \nu\right]$ , and using the fact that  $d_j(L) \sim \text{Poisson}(\alpha\nu/4)$ .  $\square$

An obvious corollary of this Proposition is that, for  $\alpha \in (0, 1)$ , the probability of finding a local clan is strictly positive in any local neighborhood over  $\Omega$ .

## APPENDIX E. 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.

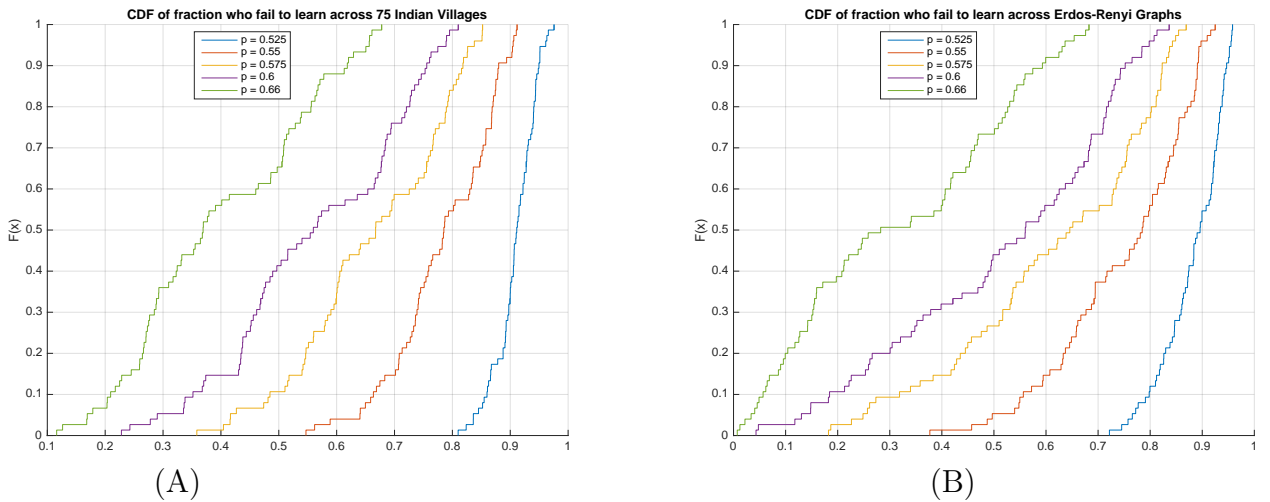


FIGURE 7. Both panels (A) and (B) present CDFs of the fraction of nodes that initially received the signal  $1 - \theta$  that became stuck at the wrong belief for various levels of  $p$ . Panel (A) presents results where we conduct simulations using the 75 Indian village networks from [Banerjee et al. \(2013\)](#). 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.