ECONOMETRICA
JOURNAL OF THE ECONOMETRIC SOCIETY
An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics
http://www.econometricsociety.org/

Econometrica, Vol. 88, No. 1 (January, 2020), 1–32

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

ARUN G. CHANDRASEKHAR
Department of Economics, Stanford University, NBER, and J-PAL

HORACIO LARREGUY
Department of Govenement, Harvard University

JUAN PABLO XANDRI
Department of Economics, Princeton University

Econometrica, Vol. 88, No. 1 (January, 2020), 1–32

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

ARUN G. CHANDRASEKHAR
Department of Economics, Stanford University, NBER, and J-PAL

HORACIO LARREGUY
Department of Governement, Harvard University

JUAN PABLO XANDRI
Department of Economics, Princeton University

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. Agents are drawn from a mixture of learning types—Bayesian, who face incomplete information about others' types, and DeGroot, who average their neighbors' previous period guesses and follow the majority. We study (1) learning features of both types of agents 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 perform a reduced-form analysis and then structurally estimate the mixing parameter, finding the share of Bayesian agents to be 10% and 50% in the Indian-villager and Mexican-student samples, respectively.

KEYWORDS: Networks, social learning, Bayesian learning, DeGroot learning.

1. INTRODUCTION

INFORMATION AND OPINIONS ABOUT 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 to which they are able to assess how much independent information is contained among the social connections with whom they communicate, as well as in whether they account for the naivete of those connections.

In many settings, individuals cannot or do not transmit their beliefs to an arbitrarily fine degree, nor their information sets.[1] Individuals are constrained to processing coarser

Arun G. Chandrasekhar: arungc@stanford.edu
Horacio Larreguy: hlarreguy@fas.harvard.edu
Juan Pablo Xandri: jxandri@princeton.edu

We are grateful to Daron Acemoglu, Abhijit Banerjee, Esther Duflo, Ben Golub, Matthew O. Jackson, Markus Mobius, Omer Tamuz, Adam Szeidl, and Chris Udry 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 Zermeño, 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.

[1]In other cases, communicating very fine information may be more feasible. The extent to which this is feasible is surely context specific.

 https://doi.org/10.3982/ECTA14407

information from their network neighbors when engaging in learning. This constraint may be for many reasons, including but not limited to operating in settings where learning is through observations of others' actions or the costs of communicating very complex information are too high and therefore only summaries are transmitted.[2]

We study such a coarse communication environment. We focus on a setting in which 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.[3] 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 (see Gale and Kariv (2003) and Mossel and Tamuz (2010) for a myopic learning environment and Mossel, Sly, and Tamuz (2015) for a strategic learning environment). However, if all agents update their guess as the majority of their neighbors' prior guesses—as modeled by a *coarse DeGroot* model of learning, 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 guess. In practice, it might be that there is a mix of sophisticated (Bayesian) and naive (DeGroot) 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 coarse communication environment.

This paper develops an incomplete information model of social learning with coarse communication 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 all-Bayesian and all-DeGroot—and is a hybrid that serves as an empirically relevant benchmark as, a priori, agents are likely to be heterogeneous in the degree of sophistication when engaging in social learning.

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.

[2]Examples of coarse information include observing whether neighbors intend to join microfinance, the kind of crop or fertilizer neighbors use, or the political billboards neighbors post outside their houses.

[3]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, studied models where signals are endowed in the first period and then agents repeatedly pass on the averages of their neighbors' prior period averages to learn the state of the world. Under DeGroot learning, which we call continuous DeGroot because of the lack of coarseness, 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.

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 with coarse communication.[4] 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 composed 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 all-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 whether the majority of initial signals was 0 or 1 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 information all-Bayesian model (where there is coarse communication) and the oft-studied continuous DeGroot model (where agents can pass on their exact beliefs). For any sequence of growing networks with uniformly bounded degree, the standard all-Bayesian model with coarse communication 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, in the continuous DeGroot model where agents can pass on arbitrarily fine information, they too all learn the state of the world in the long run. In contrast, if the sequence has a non-vanishing share of finite clans in the limit, then the incomplete information Bayesian model with coarse communication 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 incomplete information Bayesian agents play strategically rather than myopically.

Third, we address whether realistic networks have clans, which lead to the failure of asymptotic learning in our incomplete information Bayesian model. We study a mixture of two canonical models—random geometric graphs and Erdős–Rényi 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, if local linking is bounded at some positive rate for any global linking rate tending to zero at the inverse of the number of nodes,[5] the share of clans is uniformly bounded from below and therefore learning under the incomplete information Bayesian model leads 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 time, subjects in the Indian experimental sample who are in a clan that comes to consensus remain stuck on the wrong guess when the all-Bayesian model would suggest that they change their guess. However, in the Mexican experimental sample, this number is 30.3%. Second, over 82.9% of the time, subjects in the Indian sample who have an information set

[4]Henceforth we omit the designation of coarse communication which is to be assumed unless noted otherwise.

[5]These technical choices are made to preserve the sparseness of the overall network.

that is dominated by a network neighbor fail to simply copy their neighbor, which is what Bayesians would do in an all-Bayesian environment. In contrast, this failure occurs 54.5% of the time in the Mexican data. Third, nearly 94.5% of the time, subjects in the Indian experiment respond to changes in their neighbors' guesses despite learning whether the majority of initial signals in the network were 0 or 1, which Bayesian agents would never do even in an incomplete information environment. This happens 60.0% of the time in the Mexican sample. Fourth, 93.1% of the time, subjects in the Indian sample respond to changes in the behavior of agents with a dominated information set, which again Bayesian agents (even in an incomplete information environment) 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 focus on 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 mixing parameter is 0.1 (with a standard error of 0.130) in the Indian network data, whereas the mixing parameter is 0.5 (with a standard error of 0.184) in the Mexican data. At the individual level, the mixing 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 likelihood that we maximize terminates once an agent hits a zero-probability event. Another way we could have in principle dealt with this issue would be by adding trembles and studying quantal response equilibria (QRE) as in Choi, Gale, and Kariv (2012). We show in our setting that this is computationally infeasible.[6] So, instead, we model agents that do not internalize mistakes and only consider the experimental data before agents reach zero-probability events, which we show delivers a consistent estimate of the mixing parameter.

Our results then, first, suggest that in settings where networks are not sufficiently expansive (i.e., clans are likely) and there is a significant number of individuals that are not sophisticated in their learning behavior, we should observe individuals exhibiting clustered, inefficient behavior. Returning to some of our motivating examples, we might then see groups of farmers that do not adopt a more efficient technology, or pockets of voters that continue to support candidates revealed to engage in malfeasance. Second, our results indicate that there might be complementarities between financial market completion and information aggregation. Due to limited access to financial institutions and contracting ability, to support informal exchanges, villagers may organize themselves into "quilts" which are composed of numerous clans (Jackson, Rodriguez-Barraquer, and Tan (2012)). To the extent that information links are built upon the same sorts of relationships, social learning might be vulnerable particularly when there are weak institutions. This then suggests that even partial financial market completion (e.g., the introduction of microcredit) might contribute to social learning through the dissolution of clans by reducing

[6]A back-of-the-envelope calculation shows that using QRE in our setting with seven agents would take 377,346,524 years in a case where the same computation with three agents would take 4.5 hours.

the need for quilts.[7] Third, our results motivate a simple set of policies to mitigate such type of behavior. Policy-makers should not only do a widespread diffusion of information about technologies, products, or incumbent politicians, but also encourage expansiveness by generating opportunities for engagement and conversations across groups of individuals, for example, via village meetings.

We contribute to active discussions in both theoretical and experimental literatures studying social learning on networks. The theoretical works closest to ours are Feldman et al. (2014), Mossel and Tamuz (2010), Mossel, Neeman, and Tamuz (2014a), and Mossel, Sly, and Tamuz (2015). Mossel and Tamuz (2010) considered a sequence of growing graphs and showed that, in our coarse communication setting, if all agents are Bayesian, then in the limit there will be consensus on the true state. Mossel, Sly, and Tamuz (2015) showed that the same result holds true if agents behave strategically, where agents follow Nash equilibrium strategies of a repeated game. Turning to the coarse DeGroot model, Feldman et al. (2014) studied a variation with asynchronous learning behavior, where in each period a single node is chosen uniformly at random to update, and showed that, as long as the network is sufficiently expansive, there will also be asymptotic learning. Mossel, Neeman, and Tamuz (2014a) focused on a synchronous setting and studied when networks are such that the majority reaches the right opinion, and studied unanimity in the special case of d-regular graphs, and, like Feldman et al. (2014), showed 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 has 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 studied 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, Hirshleifer, and Welch (1992), Smith and Sorensen (2000), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Lobel and Sadler (2015), Eyster and Rabin (2014)).[8] 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, Gale, and Kariv (2005) demonstrated that, in networks of three nodes, the data are consistent with Bayesian behavior. Similarly, Corazzini, Pavesi, Petrovich, and Stanca (2012) showed 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

[7] This dynamic of reduced triadic closure, even among information links and not just informal insurance links, due to the entry of microcredit is documented in Banerjee, Chandrasekhar, Duflo, and Jackson (2018).

[8] Eyster and Rabin (2014) showed 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."

could pass on real-numbers (arbitrarily fine information), the eventual averages reflected double counting, consistent with DeGroot-like behavior. Meanwhile, Mobius, Phan, and Szeidl (2015) conducted 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 (2012) also conducted lab experiments to look at Bayesian versus DeGroot-like learning. Mueller-Frank and Neri (2013) developed a general class of models where individuals perform what they called Quasi-Bayesian updating and showed 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 heterogeneous population in an incomplete information Bayesian environment. This framework nests all prior models that study cases with coarse DeGroot agents or only complete information Bayesian agents 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, some of which 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, by conducting our experiment in two contrasting settings, 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 in a coarse learning environment where every agent in the network is drawn to be either a Bayesian or a DeGroot type with independent and identically distributed probability (π). This nests the pure DeGroot model ($\pi = 0$), the complete information all-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 $N_i^\star := N_i \cup \{i\}$.

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 π of $\eta_i = B$. This describes how each agent processes information, either according to the DeGroot model or by using Bayesian updating. The process by which η_j are drawn is commonly known, as is the structure of the entire network.[9] Thus, it is 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. We can also interpret this as agents coarsely communicating by only stating their guess about the state of the world.[10]

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}, the historical set of guesses by all neighbors of i.

2.1.2. *Learning*

Consider a DeGroot agent ($\eta_i = D$). This agent, in a coarse 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.[11] Formally, for $t > 1$,

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

[9] 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.

[10] In Supplemental Material, Online Appendix F (Chandrasekhar, Larreguy, and Xandri (2020)), we extend the definition of the coarse DeGroot model to the case where the agents can communicate up to an m degree of granularity their beliefs about the state of the world rather than simply reporting one of $\{0, 1\}$. That is, agents can report beliefs $\{0, \frac{1}{m-1}, \dots, \frac{m-2}{m-1}, 1\}$ from rounding their averaging of their neighbors' and own prior-round coarse beliefs. We show that the results of our main theorem on the failure of asymptotic learning in Theorem 1 extend to this case with the appropriate generalizations.

[11] We designed the experiments in order to minimize the possibility of such ties. In particular, we selected Network 3 to have no possibility of ties under either the all-Bayesian model or the DeGroot model.

Next, consider a Bayesian agent ($\eta_i = \mathrm{B}$). Since there is incomplete information about the types of the other agents in the network, Bayesian individuals attempt to learn about the types η_j of all other agents in the network while learning about the underlying state of the world in order to make their most informed guess about it in every period.[12] 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 in the Supplemental Material (Chandrasekhar, Larreguy, and Xandri (2020)).

In what follows, when we say "incomplete information (Bayesian) model," we are always referring to the model with incomplete information composed of Bayesian and DeGroot agents that learn in a coarse communication environment. Bayesian agents use Bayes's rule as above and DeGroot agents average their neighbors' and own prior guesses, and both types of agents can only coarsely communicate their best guess in every period.

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 concepts of stuckness and a clan. We define 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). That is, i is stuck if there exists some $t > 0$ such that $a_{i,t+m} = \theta$ for all $m \in \mathbb{N}$.

Next, we define a clan as a set of nodes who are more connected among themselves than to those outside the group. The formal definition follows after some additional notation. First, 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.

Formally, define a group $C \subset V$ to be a *clan* if, for all $i \in C$, $d_i(C) \geq d_i(V \backslash C)$ and $|C| \geq 2$.[13] 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 that 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 $t \geq 1$ and some fixed a such that for all $i \in C$, $a_{it} = a$. Then $a_{i,t+\tau} = a$ for all $i \in C$ and $\tau > 0$.*

PROOF: See Appendix A. *Q.E.D.*

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 π and applies for $\pi < 1$ since DeGroot agents simply average their neighbors' past period behavior.[14] See Figure 1 for an illustration.

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

[12]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.

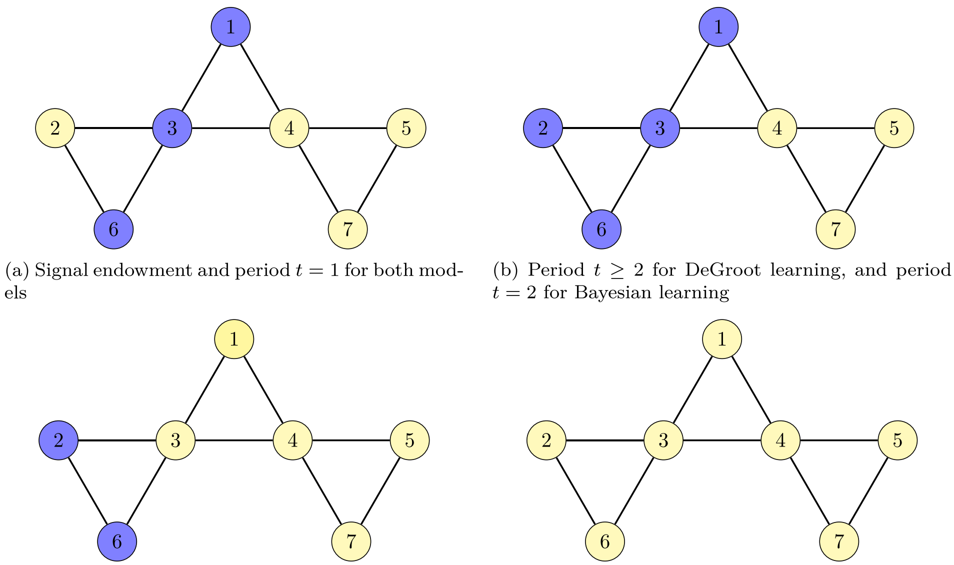

(a) Signal endowment and period $t = 1$ for both models

(b) Period $t \geq 2$ for DeGroot learning, and period $t = 2$ for Bayesian learning

(c) Period $t = 3$ for Bayesian learning

(d) Period $t \geq 4$ for Bayesian learning

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/light gray in this example).

nodes 3 and 2 in Panel C of Figure 2. Note that $N_2^\star \subset N_3^\star$. 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 $N_i^\star \subset N_j^\star$. Then $a_{i,t} = a_{j,t-1}$ for all $t > 2$.*

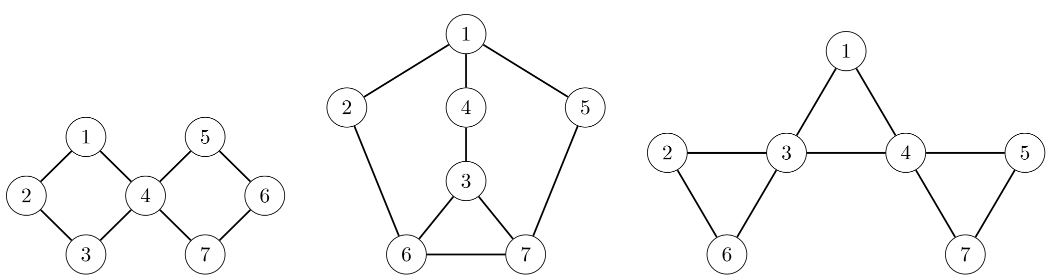

(a) Network 1

(b) Network 2

(c) Network 3

FIGURE 2.—Network structures chosen for the experiment.

In addition, the theoretical and experimental literature assumes away experimentation (see, e.g., Choi, Gale, and Kariv (2012), Mossel, Sly, and Tamuz (2015)).

[13]In Supplemental Material Appendix F, we extend the definition of a clan to the case where m coarse messages can be passed.

[14]With probability $(1 - \pi)^{|C|}$, a clan C consists of all-DeGroot agents.

PROOF: The proof is straightforward and therefore omitted. *Q.E.D.*

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 whether the majority of initial signals was 0 or 1, her guess should never change, irrespective of π. 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 η_j and π, 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* $\eta_i = \mathrm{B}$.

1. *If, at any period* τ, *agent* i *learns the majority of initial signals distributed* $\widehat{\theta} := \mathbf{1}\{\frac{1}{n}\sum_j s_j > \frac{1}{2}\}$, *then* $a_{i,t} = \widehat{\theta}$ *for all* $t \geq \tau$, *irrespective of any future sequence of actions by* i*'s neighbors.*
2. *If* $N_j^\star \subseteq N_i^\star$, *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,* $a_{i,t} = \mathrm{Function}(s_j, (A_k^{t-1})_{k \in N_i \setminus \{j\}})$.

PROOF: The proof is straightforward and therefore omitted. *Q.E.D.*

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 copy the actions of dominating agents, and agents who necessarily learn whether the majority of initial signals was 0 or 1, or reach a time when they necessarily dominate other agents' available information, ignore the actions of other agents going forward.

2.3. *Asymptotic Efficiency: Long-Run Behavior in Large Networks*

We have studied the differences in behavior between DeGroot and Bayesian agents in our environment, for $\pi \in [0, 1]$: all DeGroot, all Bayesian, and incomplete information intermediate cases. Now we assess how efficient learning is in the long run in a large network under these models. We show that non-trivial misinformation traps should occur if there are a non-trivial number of clans of bounded size in the network.

Then, as a point of contrast, we compare our results to the two most studied models in the literature. The first point of comparison is the complete information all-Bayesian case in this exact coarse environment ($\pi = 1$ in our notation). Prior work has shown that exactly in the same setup as ours, if all agents were Bayesian, then large, sparse networks should do quite well at information aggregation.

The second point of comparison is the continuous DeGroot model. Note that this goes outside of our setting in the sense that it replaces coarse communication with the ability of an agent to communicate arbitrarily fine numbers.[15] When studying naive learning

[15] It is also worth noting that the analog to this in a Bayesian environment would be Bayesians being able to pass on their entire posterior about the state of the world, which we omit for the obvious reasons.

models, the literature has often turned to this model to develop intuitions. The core result in this case is that in our setting, if naive agents can communicate with continuous bits of information, then large networks will aggregate information.[16]

Again, these points of comparisons are to illustrate that in our general environment with coarse learning, the intuitions from the literature are overturned. The main point that we illustrate is that, for a sensible sequence of networks, both the complete information all-Bayesian and continuous DeGroot 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 with a non-zero share of coarse DeGroot agents, there is a non-vanishing share of agents that remain 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 continuous DeGroot model, these $p_{i,0}$'s and their subsequent averages can be communicated directly. That is,

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

The guess is then $a_{i,t} = \mathbf{1}\{p_{i,t} > \frac{1}{2}\}$ and we denote the limit guess as $a_{i,\infty} = \lim_{t\to\infty} \mathbf{1}\{p_{i,t} > \frac{1}{2}\}$.

Meanwhile, in the Bayesian and coarse DeGroot 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 continuous DeGroot model and $a_{i,t}^{(n)}$ analogously be defined for each model. We define a sequence as *asymptotically efficient* if, for all $\epsilon > 0$,

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

Under very general conditions, we show that both the continuous DeGroot model and the complete information all-Bayesian model achieve asymptotic efficiency, but that the incomplete information model with coarse DeGroot agents 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*: $d_i(V_n) \le \bar{d}$ *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 continuous DeGroot model is asymptotically efficient*,
2. *the complete information all-Bayesian model is asymptotically efficient*, *and*
3. *the incomplete information Bayesian model with $\pi < 1$ Bayesian and $1 - \pi$ coarse DeGroot agents 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 model is not asymptotically efficient*.

[16]In Supplemental Material, Online Appendix F, we demonstrate that if DeGroot agents can transmit beliefs that are coarse but not necessarily binary, there is still a failure of asymptotic learning if there is a non-trivial share of the appropriate type of clan present.

PROOF: See Appendix A. *Q.E.D.*

Note that stuckness requires clan membership. In Supplemental Material Appendix G, we show that clans are essential for stuckness since a node must be in a clan that is itself entirely stuck to become stuck.

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, Sly, and Tamuz (2014b), also shown in Menager (2006)). Both results 2 and 3 remain true if the Bayesian agents play a Nash equilibrium in the normal form of the game instead of the assumption of myopic behavior by Bayesian agents. Specifically, Mossel, Sly, and Tamuz (2015) showed 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 do 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 π 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. In Supplemental Material, Online Appendix F, we extend this result to the case where coarse DeGroot agents can pass on m coarse estimates of beliefs $\{0, \frac{1}{m-1}, \dots, \frac{m-2}{m-1}, 1\}$ generated by rounding the prior round averages of their neighbors' and own coarse beliefs. For the appropriate extension of the notion of clans, again if there is a non-vanishing share of nodes in clans, the incomplete information model also fails asymptotic efficiency. This is prevented in the continuous DeGroot model by allowing $p_{i,t}$ to take any number in $[0, 1]$, so arbitrarily small changes in beliefs are effectively passed through to other agents.

2.4. *What Kinds of Networks Exhibit Stuckness?*

Having examined the properties of continuous DeGroot, complete information all-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 $n \to \infty$ 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.[17] 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

[17]See Supplemental Material Appendix C for simulations that demonstrate this on empirical network data.

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 checking 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 Erdős–Rényi 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 Erdős–Rényi (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 convinced of the wrong state of the world and are unable to shift away from this, if they behave in a coarse DeGroot manner.

Let $\Omega = [0, k]^2 \subset \mathbb{R}^2$ with k be a latent space.[18] We say an *RGG-ER mixture* is (α, β)-*mixed* if the network, where $(\alpha, \beta) \in [0, 1]^2$, is formed as follows. There exists a Poisson point process on Ω, 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 $A \subset \Omega$,

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

Given the draw of n nodes, the network forms as follows:

1. RGG component:
- If $d(i, j) \le r$, then i and j are linked with probability α.
- 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 β.

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

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

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

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

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

We assume $\alpha \ge \beta$ 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.

[18] The result is generalized to the case with $\Omega = [0, k]^h$. Here we take $h = 2$.

Formally, we are interested in a sequence of networks $(\Omega_k, \alpha_k, \beta_k)$ such that $\mathrm{vol}(\Omega_k) \to \infty$.[19] 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 and we have α_k either converging to a positive constant or fixed along the sequence. For the ER component, this requires $\beta_k = O(\frac{1}{\mathrm{vol}(\Omega_k)})$.

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 any two nodes in this group having a link is α_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 (α_k, β_k)-mixture of Random Geometric and Erdős–Rényi graphs, and a sequence $\{\Omega_k = [0, k]^2\}_{k\in\mathbb{N}}$ with $k \to \infty$ so $\mathrm{vol}(\Omega_k) \to \infty$. If $\beta_k = O(\frac{1}{\mathrm{vol}(\Omega_k)})$ and $\alpha_k \to \alpha > 0$, then the share of nodes belonging to local clans remains positive as $k \to \infty$. If $\alpha_k \to 0$, the share of such nodes vanishes as $k \to 0$.*

The proof is a direct corollary to Proposition A.2 in Appendix A.[20] 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 infeasible.[21]

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

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

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

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

PROOF: See Appendix A. *Q.E.D.*

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

[19] For a general h, we would have $\mathrm{vol}(\Omega) = k^h$ and we take $k \to \infty$.

[20] 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.

[21] Notice that in the particular case of a sparse Erdős–Rényi graph, the sequence would have $\alpha_k = \beta_k \to 0$, and so as noted at the end of the proposition, the share belonging to such local clans would vanish.

spectrum of the Laplacian of g, $L := I - D^{-1/2}GD^{-1/2}$, where $D = \text{diag}\{d_1(G), \dots, d_n(G)\}$, to bound $\phi(G)$ by the second-smallest eigenvalue of the Laplacian:

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

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.

In Appendix C, we study the expansiveness of the 75 Indian village networks collected in Banerjee, Chandrasekhar, Duflo, and Jackson (2019). We look at the household-level network focusing on the graph of links through which information is diffused. First, we find that no graph has an expansiveness above 0.4, let alone $\frac{\sqrt{2}}{2}$. Second, after simulating a DeGroot model with $p = 0.6$, we show that the share stuck is high, with a maximum of over 45% and an average of roughly 20–25%. Third, if we consider nodes that began with an incorrect signal, we find that there are nearly 60% of those nodes that get stuck in one village and on average there are around 30% stuck. Fourth, we find that the stuckness rate increases with the lack of expansiveness of the network. Altogether, this suggests that the structure of empirical networks in this setting has a considerable number of clans and may be susceptible to misinformation traps.

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, Chandrasekhar, Duflo, and Jackson (2013)), politics (Alt, Jensen, Larreguy, Lassen, and Marshall (2019), Cruz, Labonne, and Querubin (2018), Duarte, Finan, Larreguy, and Schechter (2019)), 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 of Mexico's most highly-ranked institutions of higher education and systematically places its undergraduate students studying economics or 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 seven subjects anonymously played the learning game on three different network structures,

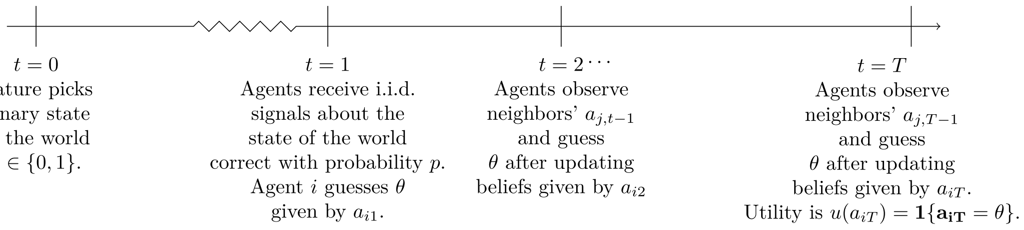

FIGURE 3.—Timeline.

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 seven 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. After each round, the game continued to the next round randomly and, on average, lasted six rounds.

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 at the time in this setting. In Mexico, students were paid $100 pesos if they guessed the state correctly and an additional $50 participation fee, which amount 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 seven 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 the 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 in both 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 six rounds. If the game continued to the second round, at the beginning of this round, each participant was shown the round-1 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 I presents the results for our experiment among Indian villagers and Table II 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 a consensus, all agents remain stuck. To assess the prevalence of this feature in the experimental data, Panel A presents the share of times that clans remain stuck despite 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 $t > 2$. Panel B shows 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, presents the share of times a Bayesian agent would have

TABLE I

REDUCED-FORM PATTERNS: INDIA[a]

Panel A: Stuckness

	(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)
	0.946 (0.0303)
Observations	74

Panel B: Information dominance

	(1) Share of times information dominated agent fails to copy dominating agent (Complete Information Bayesian predicts 0)
	0.829 (0.0380)
Observations	140

Panel C: Information revelation

	(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)
	0.931 (0.0203)	0.945 (0.0249)
Observations	159	73

[a]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.

learned whether the majority of initial signals was 0 or 1 and yet changes guess along the path in a manner consistent with DeGroot learning.[22]

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

[22]We compute this by enumerating all cases of the complete information Bayesian model in the networks we used and then calculating this share directly.

TABLE II

REDUCED-FORM PATTERNS: MEXICO[a]

Panel A: Stuckness

	(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)
	0.303 (0.144)
Observations	33

Panel B: Information dominance

	(1) Share of times information dominated agent fails to copy dominating agent (Complete Information Bayesian predicts 0)
	0.545 (0.0660)
Observations	112

Panel C: Information revelation

	(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)
	0.614 (0.0862)	0.600 (0.117)
Observations	57	35

[a] 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.

instead should restrict to only using the agent's 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.[23]

[23] Recall information dominance means that an agent's information set contains another's information set, so for instance, if node 3 information dominates nodes 2 and 6 in Network 3 from the experiment. We calculate this share by enumerating all such cases.

In the Indian village data, we find evidence that is consistent with DeGroot behavior and inconsistent with Bayesian behavior. In Panel A of Table I, 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% of the time, 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 whether the majority of initial signals was 0 or 1, the agent changed her opinion in the direction suggested by 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 incomplete information 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 the time. Column 1 of Panel C shows that, when agents learn whether the majority of initial signals was 0 or 1, 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 consistent with the (all) DeGroot model and inconsistently with any Bayesian model (or at least one with a π significantly different from 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 a pure DeGroot model, a pure Bayesian model, or an incomplete information 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 π 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 perform an *individual-level estimation*. In this case, the observational unit is each individual's action. In contrast with the network-level estimation, 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 we should think about the macro-dynamics of, and the model that governs, the social learning process. Nonetheless, both approaches yield similar results.

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

$$a_{it} = \begin{cases} a^{\star}_{it} & \text{with probability } 1-\epsilon, \\ 1-a^{\star}_{it} & \text{with probability } \epsilon, \end{cases}$$

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

The matrix $A^T_v = [a_{it,v}]$ is the data set 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 π, because they are independent of η: in period 1, every agent plays their signal, and in period 2, every agent plays the majority.

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

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

Taking logs, we can write the estimator as

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

In Appendix D, we prove the consistency of the estimator and show simulations that indicate that our estimators are consistent. Specifically, first, under each $\pi = \{0, 0.1, \ldots, 0.9, 1\}$, 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 seven nodes such that, at the parameter value, the likelihood of this configuration changes steeply if π changes minimally. While far from comprehensive at capturing all the experimental variation used to identify π, the reduced-form results of learning patterns that we describe above provide considerable contrast between Bayesian and DeGroot learning, thus contributing to its 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 as π 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 among 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.[24]

Before turning to our structural estimates of π, we estimate ϵ, which is common irrespective of network-level or individual-level estimation. Note that, for any node i in any graph v, both the Bayesian and DeGroot models, irrespective of π, prescribe the majority as an action in the second period. Therefore, recalling that $N_i^* = \{j : g_{ij} = 1\} \cup \{i\}$,

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

By standard arguments, this is a consistent and asymptotically normally distributed estimator since this is composed of a set of Bernoulli trials. Panel B of Table III shows that, in both cases, $\widehat{\epsilon}$ is similar: 0.1288 (standard error 0.007) in the Indian experimental sample and 0.134 (standard error 0.013) in the Mexican experimental 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 π.

5.1. *Main Structural Results: Network-Level Estimation*

Next, we turn to the estimation of π. 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 $a_{it}^{\star}(A^{t-1}; \eta, \pi) = a_{it}^{\star}((a_{i0})_{i=1}^n; \eta, \pi)$. This means that given a signal endowment, a type endowment, and a parameter π, 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 III presents the results of our structural estimation from the network-level perspective. In column 1, we present the results for the data from the experimental sample of Indian villagers and, in column 2, the data from the Mexican student experimental sample. In the Indian sample, we find that the share of Bayesian agents is low—$\widehat{\pi}$ is 0.1 (standard error 0.130) and we cannot reject that it is statistically different from zero. This is consistent with the reduced-form results in Table I, 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 π are consistent with our prior reduced-form results from Table II for the Mexican sample. There we find mixed behavior, where, for

[24]This procedure is analogous to clustering and, therefore, is conservative by exploiting only variation at the block level.

TABLE III

STRUCTURAL ESTIMATES[a]

Panel A: ϵ

(1) ϵ_{India}	(2) ϵ_{Mexico}
0.1288	0.134
(0.007)	(0.013)

Panel B: Network Level

(1) π_{India}	(2) π_{Mexico}
0.1	0.5
(0.130)	(0.184)

Panel C: Individual Level

(1) π_{India}	(2) π_{Mexico}
0.1	0.4
(0.113)	(0.268)

[a]Block-bootstrapped standard errors at the session level for π and at the agent level for ϵ are reported in parentheses.

example, agents that necessarily learned whether the majority of initial signals was 0 or 1 change their action consistent with DeGroot learning 60% of the time, and we estimate that $\widehat{\pi} = 0.5$ (standard error 0.184).

5.2. *Individual-Level Estimation*

Having looked at the estimates of π from a network-level approach, we turn our attention to the estimates from an individual-level approach. From this perspective, $a^{\star}_{it}(A^{t-1}; \eta, \pi) = a^{\star}_{it}(A^{t-1}_i; \eta_i, \pi)$ depends on the entire observed history A^{t-1}_i, the agent's type η_i, and the commonly known π.

Observe that because there is an ϵ 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. The terminal round $T^{\star}$, which is a function of the signal and type endowment as well as the sequence of shocks, is then 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 D 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 III. 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—$\widehat{\pi} = 0.1$ (standard error 0.113)—and again we cannot reject that it is different from zero. This is consistent with the reduced-form results in Table I. Similarly, the structural estimates of the share of Bayesian agents—$\widehat{\pi} = 0.4$ (standard error 0.268)—are consistent with our prior reduced-form results from Table II for the Mexican student data.

5.2.1. *Trembles and Quantal Response*

In our model, agents arrive at zero-probability events.[25] Since agents do not internalize that others can make errors, they may arrive 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.

An a priori natural alternative way to eliminate the zero-probability information set problem is to introduce disturbances (e.g., trembles as in the quantal-response equilibrium in Choi, Gale, and Kariv (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* $\Theta(T)$.[26] *Moreover, it is asymptotically tight; that is, any algorithm implementing Bayesian learning must have a running time of at least* $\Theta(T)$.

PROOF: See the computation in Appendix A. *Q.E.D.*

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. In Appendix A, we similarly show that the algorithm for the incomplete information model is $\Theta(n4^{2^n}T)$. 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* $\Theta(n4^{n(T-1)})$. *Moreover, the problem is NP-hard in* (n, T).

PROOF: See the computation in Appendix A. *Q.E.D.*

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 1.19×10^{16} more computations than the deterministic model. With the same T, the incomplete information model ($\pi \in (0, 1)$) involves 8.65×10^{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 trem-

[25] At $T = 4$, 58% of sessions arrive at a zero-probability event and this is considerably worse thereafter.

[26] 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 $\forall n > \underline{n}$, $c_1 \cdot |f_2(n)| < |f_1(n)| < c_2 \cdot |f_2(n)|$.

bles) takes 4 1/2 hours. The trembling hand $\pi = 1$ model, however, takes approximately 377,346,524 years.

6. CONCLUSION

We study a model of incomplete information of learning on social networks in which individuals aim to guess a binary state of the world after a series of rounds of repeated coarse communication in which they transmit their guesses to their network neighbors each period. Agents are heterogeneous in their sophistication of learning: they can be of either Bayesian or DeGroot type, where the former 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 a coarse learning environment with incomplete information. One such concept is that of clans—subgraphs of nodes who each have more links within the group as compared to outside the group. We show that realistic networks tend to have clans, which lead to the failure of asymptotic learning—non-trivial shares of agents will never learn the truth—under the incomplete information model with coarse communication. This result is 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 villager experimental sample, estimates indicate that there are approximately 10% Bayesian agents, whereas the data are best explained in the Mexican student experimental sample by a share of 50% Bayesian agents. These contrasting results point at the importance of contextual factors to understand how individuals 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, Rodriguez-Barraquer, and Tan (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, Rodriguez-Barraquer, and Tan (2012) consist of only node-adjacent clans, of which our Network 3 is an example. If agents only communicate with their favor-exchange networks,[27] 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.

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

[27] 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.

all $i \in C$. Let $T^u(i, t+\tau) = \frac{1}{d_i+1}[a_{i,t+\tau-1} + \sum_{j\in N_i} a_{j,t+\tau-1}]$ 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 $T^u(i, t+\tau) > 1/2$:

$$T^u(i, t+\tau) = \frac{\sum_{j\in\{N_i\cup\{i\}\}\cap C} a_{j,t+\tau-1}}{d_i+1} + \frac{\sum_{j\in N_i\cap C^c} a_{j,t+\tau-1}}{d_i+1}$$

$$\underbrace{=}_{\text{(i)}} \frac{d_i(C)+1}{d_i+1} + \frac{\sum_{j\in N_i\cap C^c} a_{j,t+\tau-1}}{d_i+1} \geq \frac{d_i(C)+1}{d_i+1},$$

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

PROOF OF THEOREM 1: For (1), Golub and Jackson (2010) studied a model where initial beliefs (or signals) $p_{i,t}^{(n)} \in [0, 1]$ are independently distributed, with some common mean $\mu = \mathbb{E}[p_{i,0}^{(n)}]$ and common finite variance, and agents update their beliefs according to a DeGroot model with weighting matrix $\mathbf{T}(n)$. They showed that, if $\mathbf{T}(n)$ corresponds to a uniform DeGroot weighting model (as the one we use), then $\operatorname{plim}_{n\to\infty} \max_{i\leq n} |p_{i,\infty}^{(n)} - \mu| = 0$, where $p_{i,\infty}^{(n)} = \lim_{t\to\infty} p_{i,t}^{(n)}$ (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 $\operatorname{plim}_{n\to\infty} \max_{i\leq n} |p_{i,\infty}^{(n)} - \mu| = 0$ (Golub and Jackson (2010)). 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 $\lim_{n\to\infty} \mathbb{P}\{\max_{i\leq n} |a_{i,\infty}^{(n)} - \theta| \leq \epsilon\} = 0$ 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, Sly, and Tamuz (2014b).

For Result (3), the assumption that $X_n \to x > 0$ implies that, for large enough n, there exist a number $h_n \in \mathbb{N}$ of clans of k members who are disjoint; that is, 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, $h_n \to \infty$ (since each $C_{j,n}$ has only $k < \infty$ members). Define $h = \liminf_{n\to} h_n/n$. In the incomplete information 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., $\eta_i = \mathrm{D}$ 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 α of all the disjoint clans satisfy this property, implying that at least a share αhk of agents chooses $a_{i,t} = 1 - \theta$ at every $t \in \mathbb{N}$, thus showing the desired result. *Q.E.D.*

PROOF OF PROPOSITION 4: Consider any set S such that $0 < \mathrm{vol}(S) \leq \frac{1}{2}\mathrm{vol}(V)$. Observe

$$\frac{\partial S}{\mathrm{vol}(S)} = \frac{\dfrac{\sum_{i\in S} d_i(V\backslash S)}{\sum_{i\in S} d_i(S)}}{1+\dfrac{\sum_{i\in S} d_i(V\backslash S)}{\sum_{i\in S} d_i(S)}}.$$

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

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

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. *Q.E.D.*

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. *Q.E.D.*

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

PROOF OF PROPOSITION 5: Let Ω_t be 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

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

and

$$\boldsymbol{\alpha}_t(i, \omega, t) := a_{i,t}(\omega)$$

to compute the objects $P_{i,t}(\omega)$ and $a_{i,t}(\omega)$ numerically, as in Appendix B.4. To compute them, we then have to loop across $\#(\Omega_t) \times \#(\Omega_t)$ states for each (i, t) to update $\boldsymbol{\sigma}_t$ to $\boldsymbol{\sigma}_{t+1}$ and $\#(\Omega_t)$ to update $\boldsymbol{\alpha}_t$. 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

$$\text{Computations} = nT(2^n)(1 + 2^n)k = \Theta(nT4^n).$$

Similarly, in the deterministic incomplete information model,

$$\text{Computations} = nT(4^n)(1 + 4^n)k = \Theta(nT16^n).$$

The ratio between the complete and incomplete information deterministic models is then

$$\frac{nT(4^n)(1 + 4^n)k}{nT(2^n)(1 + 2^n)k} = 2^n \frac{1 + 4^n}{1 + 2^n} \approx 4^n. \qquad \textit{Q.E.D.}$$

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 compute her beliefs about those agents' information sets. The partitional model presented in Appendix B.4 does not suffer this problem, by computing beliefs on all states, which we do here as well. Therefore, $\#(W_t) = 2^{n(t-1)}$ and then

$$\begin{aligned}
\text{Computations} &= k \sum_i \sum_{t=1}^{t-1} 2^{n(t-1)} \left(k + 2^{n(t-1)}\right) = k \sum_i \left(\sum_{t=1}^{T-1} 2^{n(t-1)} + \sum_{t=1}^{T-1} 2^{2n(t-1)} \right) \\
&= k \sum_i \left[\left(\frac{2^{nT} - 2^n}{2^n - 1} \right) + \frac{2^{2nT} - 2^{2n}}{2^n - 1} \right] \\
&= n \frac{\left(2^{nT} - 2^n\right)}{2^n - 1} \left(1 + \left(2^{nT} - 2^n\right) \frac{2^n - 1}{2^{2n} - 1} \right) = \Theta\left(n4^{n(T-1)}\right).
\end{aligned}$$

Therefore, the ratio between the complete information model with and without trembles is of order

$$\frac{\Theta\left(n4^{n(T-1)}\right)}{\Theta\left(nT4^n\right)} = \frac{1}{T} \times 4^{n(T-2)},$$

and for the incomplete information model, the equivalent ratio is

$$\frac{\Theta\left(n16^{n(T-1)}\right)}{\Theta\left(nT16^n\right)} = \frac{1}{T} \times 16^{n(T-2)}.$$

The NP-hardness result is shown in Hazla, Jadbabaie, Mossel, and Rahimian (2017). *Q.E.D.*

A.2. *Proof of Theorem 2*

A.2.1. *Bounds for Finite Graphs*

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

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 α. 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 $j \in C$, $d_j(L) \geq d_j(\Omega\backslash L)$. The goal of this section is to provide lower and upper bounds for the event

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

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

$$\begin{aligned}&\mathbb{P}\{g = (V, E) : C = V \cap L \textit{ is a local clan}\}\\ &\quad\geq \sum_{n=2}^{\infty}\left(\frac{\boldsymbol{\nu}}{4}\right)^n \frac{e^{-\boldsymbol{\nu}/4}}{n!} F^*(n-1)^n \times \alpha^{n(n-1)/2} > 0,\end{aligned} \tag{A.1}$$

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

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

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

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

PROOF: See Online Appendix E. *Q.E.D.*

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

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

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 $d_j = d_j(B(j,r)) + d_j(\Omega \backslash B(j,r))$ is a sum of independent Poisson random variables, so d_j is also Poisson, with expectation

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

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

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

which can only happen if $\beta_k = O(\omega_k^{-1})$; that is, $\beta_k/\omega_k \to \rho_\infty$. We also look only for sequences with $\alpha_k \to \alpha_\infty$, 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 (α_k, β_k), where $\omega_k \to \infty$, $\beta_k/\omega_k \to \rho_\infty > 0$, and $\alpha_k \to \alpha_\infty$, and let $L := B(i, \frac{3}{2}r) \subseteq \Omega_k$ for large enough k. Then*

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

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

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

so that

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

Since $F_\infty^*(1) = [1 + (2\alpha_\infty + \rho_\infty)\boldsymbol{\nu}]e^{-(2\alpha_\infty+\rho_\infty)\boldsymbol{\nu}} > 0$, 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 $k : \alpha_k < \boldsymbol{\nu}/4$ (which exists, since $\alpha_k \to 0$), then

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

REFERENCES

ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78, 1201–1236. [5]

ALT, J. E., A. JENSEN, H. LARREGUY, D. D. LASSEN, AND J. MARSHALL (2019): "Contagious Political Concerns: Identifying Unemployment Information Shock Transmission Using the Danish Population Network," Working Paper. [15]

AUMANN, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 1236–1239.

BANERJEE, A. (1992): "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 107, 797–817. [5]

BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): "The Difussion of Microfinance," *Science*, 341, 1236498. [15]

——— (2018): "Changes in Social Network Structure in Response to Exposure to Formal Credit Markets," Available at SSRN 3245656. [5]

——— (2019): "Using Gossips to Spread Information: Theory and Evidence From Two Randomized Controlled Trials," *The Review of Economic Studies*. [15]

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. [5]

CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2020): "Supplement to 'Testing Models of Social Learning on Networks: Evidence From Two Experiments'," *Econometrica Supplemental Material*, 88, https://doi.org/10.3982/ECTA14407. [7,8]

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. [5]

——— (2012): "Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data," *Review of Economic Design*, 16, 93–118. [4,9,24]

CHUNG, F. R. K. (1997): *Spectral Graph Theory, 92*. American Mathematical Society. [5,14]

CONLEY, T., AND C. UDRY (2010): "Learning About a New Technology: Pineapple in Ghana," *The American Economic Review*, 100, 35–69. [15]

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. [5]

CRUZ, C., J. LABONNE, AND P. QUERUBIN (2018): "Politician Family Networks and Electoral Outcomes: Evidence From the Philippines," *The American Economic Review*, 107, 3006–3037. [15]

DEGROOT, M. H. (1974): "Reaching a Consensus," *Journal of the American Statistical Association*, 69, 118–121. [2]

DEMARZO, P., D. VAYANOS, AND J. ZWIEBEL (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions," *The Quarterly Journal of Economics*, 118, 909–968. [2]

DUARTE, R., F. FINAN, H. LARREGUY, AND L. SCHECHTER (2019): "Networks, Information and Vote Buying," Working Paper. [15]

ERDŐS, P., AND A. RÉNYI (1959): "On Random Graphs, I," *Publicationes Mathematicae (Debrecen)*, 6, 290–297. [3,13]

EYSTER, E., AND M. RABIN (2014): "Extensive Imitation Is Irrational and Harmful," *The Quarterly Journal of Economics*, 129, 1861–1898. [5]

FAFCHAMPS, M., AND F. GUBERT (2007): "The Formation of Risk Sharing Networks," *Journal of Development Economics*, 83, 326–350. [12]

FELDMAN, M., N. IMMORLICA, B. LUCIER, AND S. M. WEINBERG (2014): "Reaching Consensus via Non-Bayesian Asynchronous Learning in Social Networks," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 192. [2,5]

GALE, D., AND S. KARIV (2003): "Bayesian Learning in Social Networks," *Games and Economic Behavior*, 45, 329–346. [2]

GEANAKOPLOS, J. (1994): "Common Knowledge," *Handbook of Game Theory with Economic Applications*, 2, 1437–1496.

GOLUB, B., AND M. JACKSON (2010): "Naive Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2, 112–149. [2,26]

GRAHAM, B. S. (2017): "An Econometric Model of Network Formation With Degree Heterogeneity," *Econometrica*, 85, 1033–1063. [12]

HAZLA, J., A. JADBABAIE, E. MOSSEL, AND M. A. RAHIMIAN (2017): "Bayesian Decision Making in Groups Is Hard," arXiv:1705.04770. [28]

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. [4,25]

JADBABAIE, A., P. MOLAVI, A. SANDRONI, AND A. TAHBAZ-SALEHI (2012): "Non-Bayesian Social Learning," *Games and Economic Behavior*, 76, 210–225. [2]

LIGGETT, T. M. (1985): *Interacting Particle Systems*. New York: Springer-Verlag. [2]

LOBEL, I., AND E. D. SADLER (2015): "Information Diffusion in Networks Through Social Learning," *Theoretical Economics*, 10, 807–851. [5]

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. [12,13]

MENAGER, L. (2006): "Consensus, Communication and Knowledge: An Extension With Bayesian Agents," *Mathematical Social Sciences*, 51, 274–279. [12]

MENGEL, F., AND V. GRIMM (2012): "An Experiment on Learning in a Multiple Games Environment," *Journal of Economic Theory*, 147, 2220–2259. [6]

MOBIUS, M., T. PHAN, AND A. SZEIDL (2015): "Treasure Hunt: Social Learning in the Field". [6]

MORRIS, S. (2000): "Contagion," *The Review of Economic Studies*, 67, 57–78.

MOSSEL, E., AND O. TAMUZ (2010): "Effcient Bayesian Learning in Social Networks With Gaussian Estimators," arXiv:1002.0747. [2,5]

MOSSEL, E., J. NEEMAN, AND O. TAMUZ (2014a): "Majority Dynamics and Aggregation of Information in Social Networks," *Autonomous Agents and Multi-Agent Systems*, 28, 408–429. [5]

MOSSEL, E., A. SLY, AND O. TAMUZ (2014b): "Asymptotic Learning on Bayesian Social Networks," *Probability Theory and Related Fields*, 158, 127–157. [12,26]

——— (2015): "Strategic Learning and the Topology of Social Networks," *Econometrica*, 83, 1755–1794. [2,5,9,12]

MUELLER-FRANK, M., AND C. NERI (2013): *Social Learning in Networks: Theory and Experiments*. [6]

OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. Cambrdige, MA: MIT Press.

PENROSE, M. (2003): *Random Heometric Graphs, 5*. Oxford University Press. [3,13]

SMITH, L., AND P. SORENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398. [5]

Co-editor Fabrizio Zilibotti handled this manuscript.

Manuscript received 16 May, 2016; final version accepted 13 May, 2019; available online 14 August, 2019.