Games and Economic Behavior 129 (2021) 32–54

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Cognitively-constrained learning from neighbors [☆]

Wei Li [a,*], Xu Tan [b]

[a] University of British Columbia, Canada
[b] University of Washington, United States of America

ARTICLE INFO

Article history:
Received 25 March 2020
Available online 27 May 2021

JEL classification:
D03
D83
D85

Keywords:
Depths of reasoning
(Mis)learning in networks
Heterogeneous cognitive ability
Iterated learning procedure

ABSTRACT

We present a new framework in which agents with limited and heterogeneous cognitive ability—modeled as finite depths of reasoning—learn from their neighbors in social networks. Each agent tracks old information using Bayes-like formulas, and uses a shortcut when reasoning on behalf of multiple neighbors exceeds her cognitive ability. Surprisingly, agents with moderate cognitive ability are capable of partialing out old information and learn correctly in *social quilts*, a tree-like union of cliques (fully-connected subnetworks). Agents with low cognitive ability may fail to learn in any network, even when they receive a large number of signals. We also identify a critical cutoff level of cognitive ability, determined by the network structure, above which an agent's learning outcome remains the same even when her cognitive ability increases.

1. Introduction

Learning from social networks, a common experience for many, is cognitively demanding. In order to learn correctly, we need to identify old information and incorporate new, and only new information, into our beliefs. Mounting experimental evidence has shown that learning errors are common, and there is large heterogeneity among subjects' learning outcomes. As a very recent example, Chandrasekhar et al. (2020) show in two experiments that 90% of subjects in an Indian village, and 50% of subjects in a Mexican college, fail to learn in simple networks. They suggest that the large heterogeneity among subjects may be due to their levels of education, a proxy of their cognitive ability. We present a new framework in which agents with limited and heterogeneous levels of cognitive ability—modeled as finite and different depths of reasoning—learn from their neighbors. This framework embeds classic learning rules such as naive learning, but it is rich enough to allow agents to make sophisticated inferences about what is new information from their neighbors. We show how an agent's learning outcomes depend on her cognitive ability and characterize the associated patterns of learning errors, which are potentially testable in the lab and in the field.

☆ We thank Andreas Blume, Jeff Ely, Ben Golub, Vitor Farinha Luz, Drew Fudenberg, Fahad Khalil, Jacques Lawarrée, Stephen Morris, Marit Rehavi, Tanya Rosenblat, Edward Schlee, Kyungchul Song, Wing Suen, Joel Watson, Quan Wen, and especially Matt Jackson, Li, Hao, Mike Peters, the editor and two anonymous referees for insightful comments and extensive discussions. We also thank the seminar participants at UBC, UW, UNC, MSU, WashU, PKU, UIBE, SHUFE, HKU, Emory University, SFU, UNSW, UTS, the Second Annual Conference on Network Science and Economics, the 2016 CETC, the 2016 North American Summer Meeting of the Econometric Society, the 2017 ASSA annual meeting, and the 2017 NSF/CEME Decentralization Conference for great feedback. Wei Li thanks SSHRC Insight Grant 435-2018-0265 and the Hampton Established Scholar Grant 12R74425 for financial support.

* Corresponding author.
E-mail addresses: wei.li@ubc.ca (W. Li), tanxu@uw.edu (X. Tan).

https://doi.org/10.1016/j.geb.2021.05.004
0899-8256/

In this model, agents want to learn the true state, for instance, whether the MMR vaccine leads to autism in children, a belief contributing to severe measle outbreaks globally. An agent learns by forming and updating her estimates about the state distribution, such as the probability that MMR leads to autism.[1] In period 1, the agent first forms an estimate using her initial signal. She then simultaneously reports her estimate to each neighbor and hears their reports. Next, she gets another signal from nature and period 1 ends. In each ensuing period, the same process repeats: she updates her estimates and exchanges reports with her neighbors before receiving a signal from nature.

The main innovation of our model is that agents not only form estimates about the state distribution, they also form *estimates about their neighbors' estimates*, as a store of that neighbor's old information. Agents of all cognitive abilities try to identify old information and extract new information from their neighbors' reports. An agent with a higher cognitive ability identifies old information more effectively by reasoning on behalf of her neighbors. Consider Ann with a neighbor Bob. In each period, a naive Ann attributes Bob's current report to a new signal Bob received from nature. She does not identify any old information. A slightly more sophisticated Ann—a low cognitive ability type—considers Bob's previous report as old information, and only attributes the *change* in his report as new information. That is, she believes Bob's previous report reflects his prior belief about the state, and his current report reflects his posterior belief. Then by Bayes' rule, she extracts the difference as a signal (she believes) that Bob received from nature. A yet more sophisticated Ann—a moderate or higher cognitive ability type—can estimate what Bob should report given their previous reports. For example, Ann's estimate about Bob's report should include what Bob has learned from herself and from their common neighbors. If Bob's actual report differs from her estimate, she attributes this difference, the *unexpected change*, to a new signal Bob received from nature. Agents of all cognitive abilities then update their estimates using these newly extracted signals from neighbors and their own signals.

Our first finding is that agents with moderate cognitive ability or above can learn in *social quilts*, networks in which every two agents who belong to the same circle are connected.[2] Intuitively, a social quilt is a tree of cliques (fully-connected subnetworks). It includes well-known networks such as trees, stars, and some of the core-periphery networks (see Fig. 2 for an example). Because every pair of agents in a clique is connected, when information reaches one member of the clique, all other members who have moderate cognitive ability are able to identify it as new information. More importantly, everyone correctly estimates that all other members in the clique have learned this information from the same agent. Thus, they avoid over counting the same signal by mistake. The overall tree structure of a social quilt also ensures that each signal travels through the network once and only once. As a result, all information is correctly aggregated despite agents' moderate cognitive abilities.

Second, we show that agents with low cognitive abilities may believe in the wrong state even when the network receives an arbitrarily large number of signals. The reason is that agents fail to identify old information among their neighbors. For example, a new signal of Ann is learned by Bob. But because Bob changed his report, Ann believes that Bob has learned a new signal and thus learns her own signal back from Bob. This back-and-forth process continues, and as time goes on, they believe in a fast-growing number of this signal. We show that faulty early signals may grow so fast that they dominate all the correct signals later. We hasten to add, though, even rudimentary ability to identify old information improves learning outcomes. Agents with low cognitive ability make far fewer learning errors per period than naive agents who don't have such ability, which is consistent with experimental evidence such as Grimm and Mengel (2020). Moreover, if agents with higher cognitive abilities know who have low cognitive abilities, their learning outcomes on a network level may improve significantly.

Our third finding establishes an important cutoff property for agents' learning. We show that there exists a cutoff level of cognitive ability for each agent, which is uniquely determined by the network structure. Moreover, these cutoff levels are often not high in real data, as we illustrate using the high school friendship network of Fig. 3. We show that if Ann has this cutoff level of cognitive ability, she can achieve *her best learning outcome*—her learning outcome when she has an arbitrarily high cognitive ability. But if Ann's cognitive ability is below this cutoff, then as it increases, she makes fewer (if any) learning mistakes. Consequently, her neighbors are less likely to be misled by her learning mistakes. In the end, all agents in the network may also make fewer (if any) learning mistakes.

Literature review

We model cognitive ability as depths of reasoning. As such, it shares some similarities with the level-k model adapted to the social learning context (see Kubler and Weizsacker (2004) and Penczynski (2017)). Naive agents believe that every report reflects a new private signal. This is the same as an agent with inferential naivete in Eyster and Rabin (2010) who believes that each predecessor's action only reflects his private signal. A low cognitive ability agent in our model treats the difference in each neighbor's report as new information. It is as if she treats every difference in two predecessors' actions as

[1] Her estimate is the log-likelihood ratio of the probabilities she assigns to the possible states, formally defined in Section 2. It reflects, for example, the log-likelihood ratio of her belief that MMR causes autism over her belief that MMR does not cause autism.

[2] A *path* is an ordered sequence of distinct agents in which each pair of adjacent agents is connected. A *circle* is a path in which the first and last agent are connected. A *tree* is a network in which there is no circle.

a new signal. A moderate cognitive ability agent i can be viewed as one who understands that her predecessor $i-1$ should learn from his predecessor $i-2$, and consider only the unexpected change as coming from a new signal, and so on. But our learning environment is more complex because it involves repeated exchanges of information—all agents must update their beliefs in each period—instead of each agent takes one action based on her belief and exits.

While novel, our way of modeling heterogeneity in cognitive ability is motivated and supported by the experimental literature on *strategic sophistication*. This literature studies whether subjects can perform iterative and higher-order reasoning to form the correct (higher-order) beliefs.[3] It often finds limited strategic sophistication: subjects have bounded and heterogeneous levels of sophistication, and they have difficulty in eventually forming correct beliefs (see Nagel (1995), Ho et al. (1998), Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006), Shields and Xin (2012), among many others). Nagel (1995) shows in a guessing-the-average game that about 50% of the subjects exhibit some limited ability in forming first-order and second-order beliefs over the behaviors of other participants. Costa-Gomes et al. (2001) show that 70% of the subjects have some strategic sophistication, for instance, they can best respond to the naive type's behavior or to use strict dominance arguments, but they fall far short of the sophistication required for equilibrium play. While we focus on learning instead of strategic interactions, each agent needs to form cognitively-constrained beliefs about each neighbor's beliefs given the observed reports, which is akin to forming second-order beliefs. Moreover, even in an individual decision problem, Charness and Levin (2009) show that subjects suffer from the winner's curse because they have difficulty with contingent reasoning, and they do worse as the environment becomes more complex. Collectively, this body of evidence supports modeling cognitive constraints as having limited depths of reasoning.

The vast theoretical literature on learning in networks primarily focused on two well-known benchmarks (see Golub and Sadler (2017) for a recent survey). First, Bayesian learning models show that agents with infinite cognitive ability can learn if the network is common knowledge (see Gale and Kariv (2003), Acemoglu et al. (2011), Mueller-Frank (2013), Mossel et al. (2015)).[4] A related literature on herding and social learning shows that Bayesian agents may mislearn when it is rational to ignore their private information (see Bikhchandani et al. (1992) and Banerjee (1992) among many others). While agents in this model vary in cognitive abilities and cannot form Bayesian beliefs, they still use Bayes-like formulas to form cognitively-constrained beliefs. As a result, we show that they learn just like Bayesian agents in networks such as social quilts.

Second, many existing papers study agents who follow reasonable rules of thumb to learn because of the cognitive burden of Bayesian learning.[5] DeGroot (1974) studies the simplest rule of thumb in which agents update their opinions by taking a weighted average of their neighbors' most recent reports.[6] Many papers have studied boundedly-rational learning in networks. In Bala and Goyal (1998), each agent updates her belief about the optimal action based on the choices made by her neighbors, but she does not try to infer why her neighbors choose these actions. In more recent quasi-Bayesian learning models such as Molavi et al. (2018), Levy and Razin (2018) and Mueller-Frank and Neri (2021), agents have the same depth of reasoning and treat the most recent report from each neighbor as new information. They update beliefs using a Bayes-like formula. In general, these agents fail to learn asymptotically. Our model is closer to the quasi-Bayesian learning models in which agents try to incorporate new information from their neighbors given their cognitive abilities; they do not update by a mechanical rule of thumb. Moreover, our model builds a bridge between Bayesian learning and DeGroot learning in that we allow for a wide range of cognitive abilities. A few papers have studied agents with different cognitive abilities, but to our knowledge, this model is the first to study how agents with finite and heterogeneous cognitive abilities learn in a network.[7]

More closely related to this model, Li and Tan (2020) study *locally Bayesian learning*: how Bayesian agents learn under the (possibly misspecified) beliefs that their local networks are the entire network. They identify an iterative learning procedure to implement locally Bayesian learning and relate specific features of the network to the learning outcomes. Their locally Bayesian agents can perform highly sophisticated reasoning on behalf of any long chain of neighbors. One can think of these locally Bayesian agent as one with *an infinitely high* level of cognitive ability. Here we present a tractable way of modeling agents who cannot reason in such a sophisticated way and have to use shortcuts due to their finite cognitive abilities. We then show that agents need only a cutoff level of cognitive ability to achieve their best, that is, their locally Bayesian learning outcomes.

Section 2 sets up a general model to illustrate our learning procedure and Section 3 analyze agents' learning outcomes. In Section 4, we characterize an agent's cutoff level of cognitive ability and studies its implications on her learning outcomes. Section 5 discusses our model and concludes.

[3] Many experiments have studied learning and information aggregation. See Anderson and Holt (1997), Celen and Kariv (2004), Alevy et al. (2007), Cai et al. (2009), Bai et al. (2015), Mobius et al. (2015), among others. Canen et al. (2020) provides an empirical method of testing local interactions in networks.

[4] In particular, Bayesian agents learn correctly if the network is common knowledge, and the message space is sufficiently fine (tagging the source to each signal, for instance). Even with coarse message spaces, they still learn asymptotically as long as each agent has an infinitesimally small influence.

[5] See Ellison and Fudenberg (1993, 1995), DeMarzo et al. (2003), Golub and Jackson (2010), Jadbabaie et al. (2012), among many others.

[6] DeGroot agents may learn asymptotically if each agent in the network has a vanishing social influence, in which case the Law of Large Number holds.

[7] Mueller-Frank (2014) and Chandrasekhar et al. (2020) both consider agents with two types of cognitive abilities: they may be either naive or Bayesian. Mueller-Frank (2014) shows that even one Bayesian agent can significantly improve all agents' learning outcomes.

2. Model

Network and information

Consider a network (g, G): $g = \{1, 2, \ldots, I\}$ represents a finite set of agents, and G represents the set of the links among them, $ij \in G$ if i and j are linked. Agent i (she) is a generic agent. Let $\mathrm{N}_i = \{j : ij \in G\}$ be the set of agent i's neighbors and agent j (he) be a generic neighbor of agent i. The network is undirected, so information flows both ways: $ij \in G$ if and only if $ji \in G$. It is also path-connected: for any $i, h \in g$, there is a *path* $(i_0 i_1 \ldots i_l)$ such that all agents are distinct, $i_0 = i$, $i_l = h$, and $i_k i_{k+1} \in G$ for all $k < l$.[8]

Agent i observes only her *local network*, consisting of herself, all her neighbors, and all the links among them. Denote her local network as (g_i, G_i), where $g_i = \mathrm{N}_i \cup \{i\}$ and $G_i = \{hj : h, j \in g_i \text{ and } hj \in G\}$. Similarly, the subnetwork that agent i and her neighbor j both observe is their *shared local network*, consisting of themselves, their common neighbors, and all the links among them. Denote their shared local network as $\left(g_{ij}, G_{ij}\right)$, where $g_{ij} = g_i \cap g_j$ and $G_{ij} = G_i \cap G_j$. Recall that in a clique, every pair of agents are connected. Then, for any clique $\{i, i_1, \ldots, i_{k-1}\}$, the *shared local network* of these k agents consists of themselves, their common neighbors, and all the links among them. Denote the shared local network of these k agents as $(g_{ii_1 \ldots i_{k-1}}, G_{ii_1 \ldots i_{k-1}})$, where $g_{ii_1 \ldots i_{k-1}} = g_i \cap g_{i_1} \cap \ldots \cap g_{i_{k-1}}$, and $G_{ii_1 \ldots i_{k-1}} = G_i \cap G_{i_1} \cap \ldots \cap G_{i_{k-1}}$. For instance, in a 4-agent clique, $g = \{1, 2, 3, 4\}$ and $G = \{12, 13, 14, 23, 24, 34\}$. The shared local network of any subset of these agents is the clique itself: $g_1 = g_{12} = g_{123} = g_{1234} = g$ and $G_1 = G_{12} = G_{123} = G_{1234} = G$.

There are two possible states: $S = \{s_1, s_2\}$. The two states are a priori equally likely: $\Pr(s_1) = \Pr(s_2) = \frac{1}{2}$.[9] The true state is realized before learning begins. Agents receive signals from nature about the state. The support of agent i's signals is finite: $X^i = \{x^{\emptyset}, x^{i,1}, \ldots, x^{i,M_i}\}$, where $M_i \geq 2$. Agent i receives the uninformative signal $x^{\emptyset}$ with probability $\phi^i_{\emptyset} \in (0, 1)$. For each informative signal, let $\phi^i_{mn} \in (0, 1)$ be agent i's conditional probability of receiving signal $x^{i,m}$ if the state is s_n, $n = 1, 2$.

Time is discrete: $t = 0, 1, \ldots$. In each period up to T, agent i gets a realized private signal x^i_t according to her information structure. The signals are independent across agents and time conditional on the state. No informative signal arrives at or after period T, which is randomly drawn from a geometric distribution over $\mathbb{N} \cup \{\infty\}$.[10] The common knowledge among agents includes the prior over S, the distribution of T, and the fact that their signals are independent across agents and time conditional on the state.

To focus on agents with limited cognitive abilities, we assume that agents only learn from their local networks and do not make inferences about the entire network. Moreover, we assume that agent i only knows her own information structure, namely, the number of her informative signals, the probability of getting $x^{\emptyset}$, and the conditional probability of getting informative signals given the state. She does not know any neighbor j's information structure except that each agent's information structure has full support. That is, $M_j \geq 2$, $\phi^j_{\emptyset} \in (0, 1)$, and $\phi^j_{mn} \in (0, 1)$ for all $m \leq M_j$ and $n = 1, 2$. We further discuss the reason underlying our assumption on the agents' information structures in Remark 3 of Section 2.2.4.

Reports and timeline

Agent i updates her beliefs about the state distribution in each period. Because these beliefs depend on her cognitive ability c_i (to be specified below), we denote her belief in period t as $(p^{i,c_i}_t(s_1), p^{i,c_i}_t(s_2))$.[11] Each agent can report her belief directly, but to simplify notations, we let her report the log-likelihood ratio of the probabilities of the two states. That is, agent i reports her *own estimate* at time t[12]:

$$b^{i,c_i}_t = \log p^{i,c_i}_t(s_1) - \log p^{i,c_i}_t(s_2).$$

Similarly, for each signal x^i_t from nature, the log-likelihood ratio of the two states is

$$\alpha^{ii}_t = \log \Pr\left(s_1 \mid x^i_t\right) - \log \Pr\left(s_2 \mid x^i_t\right).$$

[8] From now on, we use $(ij \ldots k)$ to denote a sequence of (possibly repeat) agents in which the order matters, and $\{i, j, \ldots, k\}$ to denote a set of distinct agents whose order does not matter.

[9] Our model easily extends to any finite number of states as long as the agents have a (possibly not uniform) common prior about the state distribution (see the extended model in the Appendix and the proofs). We use two symmetrically distributed states here because the simpler notations allow us to focus on the learning rule.

[10] More formally, let $\Pr(T = \infty) = \delta$ and $\Pr(T = k) = (1 - \delta)^k \delta$ for some $\delta \in (0, 1)$. If $T = \infty$, agents can receive an infinite number of signals, and if $T = 1$, agents receive their initial signal only. The latter is the focus of many existing models, while we allow the signals to arrive over time.

[11] Agents are cognitively constrained and cannot form beliefs as a Bayesian agent does. Therefore these probability distributions are not their Bayesian beliefs. They are agent i's *cognitively-constrained beliefs*, and agents with different cognitive abilities may form different beliefs based on the same set of reports. Throughout this paper, we refer to the cognitively-constrained beliefs as beliefs to ease exposition.

[12] Her estimate is only one number because it summarizes her beliefs given the binary states.

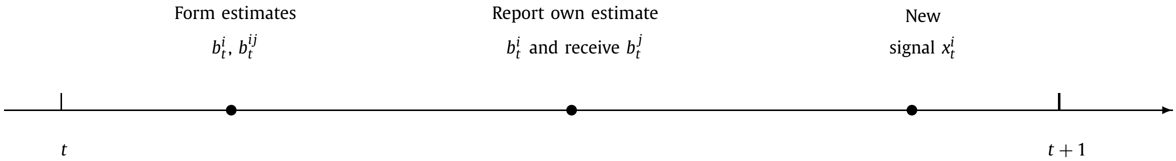

Fig. 1. Timeline.

In each period, agent i also forms beliefs about each neighbor j's belief about the state distribution (to be specified below).[13] Formally, $(p_t^{ij,c_i}(s_1), p_t^{ij,c_i}(s_2))$ is what agent i believes about agent j's belief in period t. Again, we use the log-likelihood ratios of these beliefs. Let agent i's *estimate about agent j's estimate* at period t be

$$b_t^{ij,c_i} = \log p_t^{ij,c_i}(s_1) - \log p_t^{ij,c_i}(s_2).$$

To simplify notations, we suppress c_i and use b_t^i (and b_t^{ij}) when the relevant c_i is clear. For instance, when a formula applies for any cognitive ability, or when all agents in the network have the same cognitive ability. When we compare agent i's learning outcomes under two cognitive abilities, c_i and c_i', we explicitly express this dependence by using b_t^{i,c_i} and $b_t^{i,c_i'}$.

The timeline in Fig. 1 summarizes how agent i learns in each period. Note that she reports in each period before she observes any new private signal.

2.1. Cognitively-constrained learning rule

We model an agent's cognitive ability as her ability to identify old information depending on her depth of reasoning. To operationalize it, we assume that each agent has cognitive ability $c_i \in [0, \infty)$. For instance, a naive agent i ($c_i = 0$) does not identify any old information. With low cognitive ability ($c_i = 1$), agent i identifies each neighbor j's previous report as old information. With higher cognitive abilities ($c_i \geq 2$), agent i takes past reports among all her neighbors into consideration.

Agents of all cognitive abilities use the following two *Bayes-like formulas* to learn.[14] First, agent i extracts a signal x_t^{ij} in period t from agent j, which captures what she believes to be the new information agent j received from nature at the end of period $t-1$. Specifically, agent i extracts this signal x_t^{ij} by using her estimate about agent j (b_t^{ij}) as the prior, and agent j's actual report (b_t^j) as the posterior. Using a Bayes-like formula,

$$b_t^j = \log\left(\frac{p_t^j(s_1)}{p_t^j(s_2)}\right) = \log\left(\frac{p_t^{ij}(s_1)}{p_t^{ij}(s_2)} \cdot \frac{\Pr(x_t^{ij} \mid s_1)}{\Pr(x_t^{ij} \mid s_2)}\right) = b_t^{ij} + \log\left(\frac{\Pr(x_t^{ij} \mid s_1)}{\Pr(x_t^{ij} \mid s_2)}\right).$$

Let α_t^{ij} be the log-likelihood ratio of the probabilities of the extracted signal x_t^{ij} given the states, which is the last term of the above equation. Rearrange terms, we have agent i's *signal extraction rule*:

$$\alpha_t^{ij} = b_t^j - b_t^{ij}. \tag{1}$$

Intuitively, agent i removes what she believes to be agent j's old information from his most recent report. Next, she updates her own belief using a Bayes-like formula. Specifically, her belief in period $t+1$ uses her belief in period t as the prior, and incorporates her own private signal and the new signals she extracted from her neighbors by formula (1). That is,

$$b_{t+1}^i = b_t^i + \alpha_t^{ii} + \sum_{j \in N_i} \alpha_t^{ij}. \tag{2}$$

Formula (2) is agent i's *updating rule*.[15] Key to both formulas is agent i's estimate about agent j's estimate b_t^{ij}: it stores all the old information agent i believes agent j has. We now describe how cognitive ability determines the way agent i forms b_t^{ij}.

First, agents with lower cognitive abilities learn mechanically: if $c_i = 0$, agent i does not identify any old information: $b_t^{ij,0} = 0$ for all t. By formula (1), agent i believes every report from agent j reflects a new signal: $\alpha_i^{ij} = b_t^j$. Next, agent i with $c_i = 1$ identifies agent j's previous report as old information:

$$b_1^{ij,1} = 0 \text{ and } b_t^{ij,1} = b_{t-1}^j \text{ for all } t \geq 2.$$

[13] Similar to the previous footnote, these beliefs are not agents' second-order beliefs because they are cognitively-constrained. We refer to them as beliefs to ease exposition.

[14] Because these agents are not Bayesian, they cannot form posterior beliefs as Bayesians. Throughout this paper, we say they use a Bayes-like formula to reflect that they apply Bayes' rule mechanically to *what they believe* to be new and independent signals.

[15] To see this, note that agent i applies Bayes' rule to signals x_t^i and x_t^{ij} such that for $n = 1, 2$, $p_{t+1}^i(s_n) \propto p_t^i(s_n)\Pr(x_t^i \mid s_n)\prod_j \Pr(x_{t-1}^{ij} \mid s_n)$. Take the log-likelihood ratio, we have equation (2).

By formula (1), $\alpha_t^{ij} = b_t^j - b_{t-1}^j$ for all $t \geq 2$. Intuitively, she treats the entire change in agent j's report from period $t-1$ to period t as the new signal agent j has learned. Section 2.2 provides an example of how agents with lower cognitive abilities learn.

In contrast, agent i with $c_i \geq 2$ can form her estimates involving up to c_i agents using a Bayes-like formula. To fix ideas, consider a triangle network with three agents $\{1, 2, 3\}$. Suppose that agent 1 has $c_1 = 2$, so that she can form b_t^{12} using a Bayes-like formula. In order to do so, she needs to compute what agent 2 learns from neighbor 3's report, which requires agent 1 to estimate what agent 2 believes neighbor 3 believes about the state distribution. Agent 1's sophistication is reflected in her ability to change perspective: she can reason on behalf of agent 2. She uses b_t^{123}, her estimate of agent 2's estimate of agent 3's estimate, to store agent 2's old information about agent 3 (in agent 1's opinion). Only then can agent 1 compute the signal (she believes) agent 2 extracts from report b_t^3. In the same way, she estimates what agent 2 learns from her own report. She then forms estimate b_t^{12} by following Bayes-like formulas (4) specified below. The limitation of $c_1 = 2$ is reflected in that she cannot form estimate b_t^{123} by Bayes-like formulas. Instead, she forms degenerate estimates b_t^{123} using a shortcut specified in formulas (5) and (6) below.

We now specify the learning procedure when $c_i \geq 2$. For each clique $\{i, i_1, \ldots, i_{k-1}\}$ with $k \leq c_i$, agent i forms and updates an estimate $b_t^{ii_1 \ldots i_{k-1}}$ about agent i_1's estimate about ... agent i_{k-1}'s estimate as follows.[16]

Initial values: at $t = 1$, $b_1^i = \alpha_0^{ii}$, $b_1^{ii_1} = \ldots = b_1^{i,i_1,\ldots,i_{k-1}} = 0$.

In each period $t+1$, agent i:

Step 1: updates own estimate by applying signal extraction rule (1) and updating rule (2).

Step 2: updates her estimate about the estimate of each neighbor $h \in g_{ii_1 \ldots i_{k-1}}$.

- **Extract signals from all neighbors:** extracts a new signal about what she believes that agent i_1 believes ... agent i_{k-1} has learned from agent h similar to formula (1).

$$\alpha_t^{ii_1 \ldots i_{k-1}h} = b_t^h - b_t^{ii_1 \ldots i_{k-1}h}. \tag{3}$$

- **Update all estimates:** incorporates the new signals from (3) and forms her estimate about agent i_1's estimate about ... about agent i_{k-1}'s estimate similar to formula (2).

$$b_{t+1}^{ii_1 \ldots i_{k-1}} = b_t^{ii_1 \ldots i_{k-1}} + \sum_{h \in g_{ii_1 \ldots i_{k-1}}} \alpha_t^{ii_1 \ldots i_{k-1}h}. \tag{4}$$

Step 3: sets degenerate estimates.

- **Type 1 degenerate estimate: binding cognitive constraint.** In any clique $\{i, i_1, \ldots, i_{c_i}\}$, agent i sets her (c_i+1)-agent estimate to be her c_i-agent estimate without agent i_1:

$$b_{t+1}^{ii_1 i_2 \ldots i_{c_i}} = b_{t+1}^{ii_2 \ldots i_{c_i}}. \tag{5}$$

- **Type 2 degenerate estimate: last agent is a repeat agent.** For any $b_{t+1}^{ii_1 \ldots i_{k-1}l}$ where $l \in \{i, i_1, \ldots, i_{k-1}\}$ is a repeat agent, agent i sets her multi-agent estimate equal to the one without the repeat agent:

$$b_{t+1}^{ii_1 \ldots i_{k-1}l} = b_{t+1}^{ii_1 \ldots i_{k-1}}. \tag{6}$$

Note that the last agent is the only place where repeat agent is allowed. Agent i does not form estimates involving (c_i+2) or more agents. ‖

2.2. A simple two-agent example

We use a running example with binary symmetric *informative* signals to illustrate the learning rule above. We assume that agents have identical information structures, and drop the agent index from the set of possible signals.

Example 1. The network has two agents and one link: $g = \{1, 2\}$ and $G = \{12\}$. The set of possible signals is $X = \{x^{\emptyset}, x^1, x^2\}$, where $x^{\emptyset}$ is uninformative. Agent 1, Ann, gets the only informative signal: $x_0^1 = x^1$. Agent 2, Bob, gets no informative signal. Let the Bayesian posterior given the signal be $\Pr(s_1 \mid x_0^1) = \phi$, where $\phi \in (\frac{1}{2}, 1)$, and the log-likelihood ratio be $\varphi \equiv \log \phi - \log(1-\phi)$.

At $t = 0$, Ann observes x_0^1. At $t = 1$, Ann reports her estimate based on x_0^1: $b_1^1 = \alpha_0^{ii} = \varphi$. Bob reports the symmetric prior since he has no signal: $b_1^2 = 0$.

[16] These are again the log-likelihood ratios of what agent i believes about ... about agent i_{k-1}'s belief about the state distribution. Also, notice that agent $i, i_1, \ldots, i_{k-1}$ must know each other to form an estimate involving each other. That is, they must be in the same shared network, or they form a clique.

Table 1
A two-agent example with $c_1 = c_2 = 0$.

	Ann			Bob		
	b_t^1	b_t^{12}	α_{t-1}^{12}	b_t^2	b_t^{21}	α_{t-1}^{21}
$t=1$	φ	0	n/a	0	0	n/a
$t=2$	φ	0	0	φ	0	φ
$t=3$	2φ	0	φ	2φ	0	φ
$t \geq 4$	$2^{t-2}\varphi$	0	$2^{t-3}\varphi$	$2^{t-2}\varphi$	0	$2^{t-3}\varphi$

Table 2
A two-agent example with $c_1 = c_2 = 1$.

	Ann			Bob		
	b_t^1	b_t^{12}	α_{t-1}^{12}	b_t^2	b_t^{21}	α_{t-1}^{21}
$t=1$	φ	0	n/a	0	0	n/a
$t=2$	φ	0	0	φ	φ	φ
$t=3$	2φ	φ	φ	φ	φ	0
$t=4$	2φ	φ	0	2φ	2φ	φ

Table 3
A two-agent example with $c_1 = c_2 \geq 2$.

	Ann			Bob		
	b_t^1	b_t^{12}	α_{t-1}^{12}	b_t^2	b_t^{21}	α_{t-1}^{21}
$t=1$	φ	0	n/a	0	0	n/a
$t=2$	φ	φ	0	φ	φ	φ
$t \geq 3$	φ	φ	0	φ	φ	0

2.2.1. Naive learning agent: each report reflects a new signal

Observation*: agents overcount exponentially when $c_1 = c_2 = 0$.* At each period, Ann treats Bob's previous report as new information and adds it to her report. Thus $b_2^1 = b_2^2 = \varphi$, $b_3^1 = b_3^2 = 2\varphi$, and so on. As time goes on, they believe in an exponentially increasing number of the initial signal. Their learning outcomes are summarized in Table 1.

2.2.2. Agents with low cognitive ability: any change reflects a new signal

Observation*: agents' overcount linearly when $c_1 = c_2 = 1$.* At $t = 2$, Ann and Bob report $b_2^1 = \varphi$ and $b_2^2 = \varphi$. Because Ann treats the change in Bob's report as new information (even though the change is due to the signal Bob just learned from her), she extracts a signal $\alpha_2^{12} = b_2^2 - b_1^2 = \varphi$ and incorporates it at $t = 3$. Similarly, at $t = 4$, Bob treats the change in Ann's report as a new signal and incorporates it. In each odd period, Ann extracts a new copy of x_0^1, and thus $b_t^1 = (t+1)/2\varphi$. In each even period, Bob extracts a new copy of x_0^1, and thus $b_t^2 = t/2\varphi$. Their learning outcomes are summarized in Table 2.

2.2.3. Agents with higher cognitive abilities: unexpected change reflects a new signal

An agent with higher cognitive ability ($c_i \geq 2$) understands that a neighbor learns from herself and his neighbors. When $c_1 = c_2 \geq 2$, their ability to reason on behalf of a neighbor pays off. Ann and Bob learn correctly, unlike their counterparts with lower cognitive abilities.

Observation*: agents' learning outcomes are correct if $c_1 = c_2 \geq 2$.*

At $t = 2$, Ann notices that Bob's period-1 estimate agrees with what she expects him to report: $b_1^2 = b_1^{12} = 0$. Ann learns nothing from Bob: $\alpha_1^{12} = 0$, and her report does not change: $b_2^1 = \varphi$. But Bob notices that Ann's period-1 estimate is not the prior: $b_1^1 \neq b_1^{21}$, and extracts a signal $\alpha_1^{21} = \varphi$. Bob then adds this signal to his estimate: $b_2^2 = b_1^2 + \alpha_1^{21} = \varphi$. Next, by equation (3), $\alpha_1^{121} = b_1^1 - b_1^{12} = \varphi$, and by equation (4), $b_2^{12} = b_1^2 + \alpha_1^{121} = \varphi$. These equations reflect that Ann expects Bob to learn her initial signal x_0^1, and thus she correctly forms b_2^{12}. Bob expects Ann to learn nothing from him, so he updates to $b_2^{21} = \varphi$. For all $t \geq 3$, the agents' estimates and their estimates of each other's estimate agree. There is no new information: $\alpha_t^{12} = \alpha_t^{21} = 0$, and their estimates remain unchanged. Both believe the true state is s_1 with probability ϕ, which is the Bayesian posterior. Their learning outcomes are summarized in Table 3.

2.2.4. Remarks on the main assumptions

Remark 1: key component–agent i's estimate about j's estimate. In our model, b_t^{ij} contains all the old information that agent i believes agent j has learned from all reports in $t-1$. Note that b_t^{ij} excludes agent j's private signal in period $t-1$, which arrives at the end of period $t-1$. This way of storing old information is simple to implement.[17] Agent i iteratively

[17] Alternatively, we can let b_t^{ij} be agent i's updated belief about agent j's beliefs given all information including her own private signal x_{t-1}^i, but it is more complex and not necessary. Consider the following example. At $t = 1$, given her initial signal, agent i can update her belief about agent j's belief because

computes her period-$(t+1)$ estimate from her period-t estimate. She begins by extracting new signals from her neighbors and treating all of them as independent. Then she uses her period-t estimate as her prior, applies the updating rule (2) to form her period-$(t+1)$ estimate using the extracted signals. Thus, her period-$(t+1)$ estimate reflects her most up-to-date estimate about the state distribution given her cognitive ability.

Remark 2: why do we use degenerate estimates? Each agent faces a cognitive constraint beyond which she can not process information in a Bayesian way. For example, when $c_i = 2$, agent i understands that neighbor j should learn from a common neighbor k. She may want to compute—given the reports both i and j can observe—what agent j believes agent k believes about the state distribution.[18] But agent i with $c_i = 2$ cannot do it and must use degenerate estimates, which can lead to learning errors when agent i and k observe a different set of reports from those observed by agent i, j and k (see Example 3 below). We show, however, when the degenerate estimates are without loss in Section 4.

Remark 3: knowledge about the network. Inherited from boundedly rational learning models, our agents face an implicit constraint, namely, they do not know, nor do they learn about the network. We make this assumption because learning about the entire network from neighbors' reports in addition to learning about the state distribution is cognitively and computationally intractable.[19] Instead, each agent behaves *as if* her local network is the entire network, which imposes a subtle restriction on an agent's belief about neighbors' information structure. Under this belief, any signal agent i extracts from a neighbor by equation (1), must be this neighbor's signal from nature. We assume that the agents' beliefs about other's information structure have full support, so that an agent can rationalize any extracted signal from a neighbor as one coming from nature.[20]

Remark 4: extension–one's learning also depends on neighbors' cognitive abilities. We assume, as a behavioral heuristic, that agent i's learning rule depends only on her own cognitive ability. She does not know or form beliefs over her neighbors' cognitive abilities, nor does she take into account her neighbors' cognitive ability when she estimates their estimates. The reason is that these agents don't have enough cognitive ability to form such complex beliefs. It is conceptually simple (but notationally cumbersome) to extend our model to allow each agent to know her neighbors' cognitive abilities. In this case, an agent's estimate of a neighbor's estimate must depend on that neighbor's cognitive ability, so she uses a set of neighbor-specific learning rules. We illustrate this case in Example 4 below.

3. Learning outcomes

Can agents learn correctly given their cognitive constraints? To answer this question, we first lay out our notions of correct learning. Let X_t^i be the set of signals agent i receives from nature up to and including period t. Let X_t be the union of X_t^i for all $i \in g$. Recall that T is the period at or after which the agents receive no informative signal, and thus X_T contains all the realized signals the network receives.

Definition 1. Agent i's learning outcomes are **correct** if her estimates converge to the log-likelihood ratios of the Bayesian posterior for all sequences of realized signals X_T: $\lim_{t\to\infty} b_t^i = \log \Pr(s_1 \mid X_T) - \log \Pr(s_2 \mid X_T)$. They are **asymptotically correct** if she learns the true state almost surely when agents receive infinitely many signals ($T = \infty$).

With a finite number of signals ($T < \infty$), the Bayesian posterior is bounded away from 0 and 1. In this case, the notion of correct learning requires each agent's long-run estimate to match the log-likelihood ratio of the Bayesian posterior. With infinitely many signals, asymptotic learning requires that agents learn the true state with probability 1.

3.1. Correct learning with moderate cognitive ability

With moderate cognitive ability ($c_i = 2$), agents can reason on behalf of each neighbor but no further. Despite this constraint, we now show that they (and all agents with higher cognitive abilities) learn correctly in the following type of network.

their signals are correlated unconditional on the state. But doing so requires agent i to know the information structure of agent j, which is unrealistic in many settings. More importantly, such computations are not needed. After agent i hears agent j's report b_1^j, she can easily back out agent j's initial signal using the signal extraction rule (1).

[18] This is akin to forming third-order belief and its computational difficulty. In simple dominant solvable games, Kneeland (2015) show many players don't have high enough level of rationality to form such beliefs.

[19] To do so requires an agent to form beliefs about the number of agents in the network and about all the possible paths through which information can travel to her. She also needs to update all these beliefs each time she receives new information. Computationally demanding even in a three-agent network (see Gale and Kariv (2003)), such complex updating in a typical network is beyond the capacity of modern computers. For instance, Valiant (1979) show that counting the number of $s-t$ paths in a graph is #P-complete, and Roberts and Kroese (2007) consider algorithms to improve counting in known networks.

[20] If this richness assumption does not hold, say, if agents are assumed to have identical information structures, then agent i may detect certain reports from neighbor j cannot come from j's signal from nature. If i thinks that the report is due to information from neighbors of agent j unknown to i, then agent i must update her beliefs about agent j's local network and also revise her beliefs about the state given her updated belief about the network, which is beyond her cognitive ability.

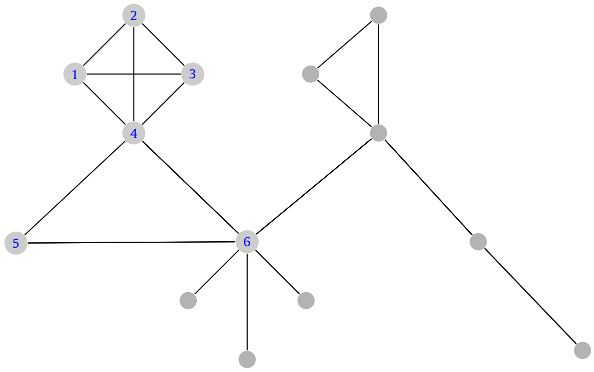

Fig. 2. A social quilt.

Definition 2. A network (g, G) is **a social quilt** if any agents i and j who belong to the same circle are connected: $ij \in G$.

Intuitively, it is a tree-like union of cliques—each node of the tree consists of a set of *fully-connected* agents. Fig. 2 illustrates a social quilt, which include well-known subnetworks such as trees, cliques, stars, lines, and some of the core-periphery networks.

Proposition 1. *In a social quilt, agents learn correctly if and only if $c_i \geq 2$ for all agents.*

Agents with moderate cognitive ability learn correctly in social quilts, because the shortcuts they use due to their cognitive constraints—the degenerate estimates in expression (5) and (6)—turn out to be correct. For instance, if agent i has $c_i = 2$, agent i uses her estimate about agent k's estimate as what she believes to be agent j's estimate about agent k's estimate ($b_t^{ijk} = b_t^{ik}$). Similarly, she uses her estimate about agent j's estimate to be what she believes to be agent j's estimate about her own estimate ($b_t^{iji} = b_t^{ij}$). These degenerate estimates are correct because of a special feature of social quilts: the shared local network of every pair of agents is a clique. To see this, suppose that agent i and j has two common neighbors, then the four of them form a circle. By the definition of a social quilt, these two common neighbors must be connected as well, and thus agent i and j's shared local network is a four-agent clique. Clearly, within a clique, everyone observes the same set of reports. Consider the clique $\{1, 2, 3, 4\}$ of Fig. 2. As $g_{41} = g_{42} = \{1, 2, 3, 4\}$, agent 4 thinks that agent 1 hears the reports from the same set of people as agent 2. If agent 4 were able to use Bayes' rule to estimate agent 1's estimates of agent 2's estimates, b_t^{412}, she would use the reports from agents $\{1, 2, 3, 4\}$ just like her own estimates about agent 2's estimates. Thus, $b_t^{412} = b_t^{42}$, and the degenerate estimates she uses are in fact correct.

Because of this special feature, each agent with moderate cognitive ability makes no local learning errors in a social quilt. For agent i, forming estimates about each neighbor's estimates correctly is sufficient. She would not learn any better even if she could form sophisticated estimates involving more neighbors (see Proposition 4 in Section 4). In addition, agents make no learning errors globally because there is no chance for old information to reach an agent repeatedly though the unobserved part of the network. This is because a social quilt contains no *simple circle*—a circle of more than three agents in which each agent has exactly two links to others in the circle. Thus new information reaches each agent once and only once, who learns it correctly.

For necessity, we show that even if just one agent has low cognitive ability (say $c_1 = 1$), all agents' learning could be wrong. Suppose that agent 1 receives the only informative signal at $t = 0$ and reports it at $t = 1$. Then each neighbor $j = 2, 3, 4$ learns it and incorporates it into his estimate at $t = 2$. Because $c_1 = 1$, agent 1 cannot recognize that the change in j's report at $t = 2$ reflects what j has learned from herself. Instead, she treats this change as a new signal and incorporates it at $t = 3$. In doing so, she overcounts her initial signal by three copies. These duplicated copies are in turn treated as new signals by 1's neighbors at $t = 4$. In this way, all agents learn more and more copies of the initial signal.[21] In the limit, everyone believes in the state that is most likely given the initial signal. But the Bayesian posterior is interior: it cannot rule out any state.

The main takeaway from Proposition 1 is that from a planner's perspective, social quilts are conducive toward agents' learning even if they are not very sophisticated. For instance, to facilitate learning about a new vaccine, a planner may want to have each local village to be fully connected, and the village leaders connected in a tree, such as the islands-connections network in Jackson and Rogers (2005). In this way, the community can learn whether the vaccine is effective without being misled by a few faulty initial trials.

In light of Proposition 1, a natural question is whether there are other networks such that agents can learn correctly. This question is complex because in general, a network often contains simple circles (such as the network in Example 2)

[21] The network's learning outcomes improve significantly if all cognitive abilities are known and agents can account for that (see Example 4) in Section 5.2.

and diamonds with a link (networks that are neither cliques nor simple circles, such as the network in Example 3). Li and Tan (2020) show that in networks with both features, agents' learning outcomes may exhibit divergent oscillations. Thus, the long-run learning outcomes may not converge and the learning dynamics is difficult to characterize. To gain tractability, we consider a stronger notion of correct learning: Agent i's learning outcomes are *strongly correct* if her estimate in every period is the log-likelihood ratio of the Bayesian posteriors given all information that has reached her in that period.[22] The next result follows from Proposition 4 of this paper and Proposition 1 of Li and Tan (2020).

Corollary 1. *When $c_i \geq 2$ for all agents, agents' learning outcomes are strongly correct if and only if the network is a social quilt.*

Intuitively, this corollary states that whenever a network is not a social quilt, agent i will make learning errors in incorporating information from her neighbors even if she has arbitrarily high cognitive ability. She either repeatedly counts old information as new, or her opinions swing back and forth. Fundamentally, she makes such errors because she cannot account for how information may reach her from the unobserved parts of the network.

3.2. Learning outcomes with low cognitive ability

Several papers have studied the learning outcomes of agents with $c_i = 0$ and show that these agents cannot learn asymptotically (see Levy and Razin (2018), Molavi et al. (2018) and Mueller-Frank and Neri (2021)).[23] Here we focus on the learning outcomes of agent i with $c_i = 1$, who treat each neighbor's previous report as his old information. We will show that she believes in duplicate copies of true signals as in Example 1 and fails to learn asymptotically. Such severe learning errors can contribute to the spread of fake news and misinformation.[24] We also compare their learning patterns with those of the naive-learning agents. These results can be used in experiments to estimate the subjects' heterogeneous cognitive abilities.

For simplicity, we consider the case of initial signals only. Let $\mathbf{b}_1$ be the vector of agents' estimates at period 1.[25] Its i-th element is agent i's estimate $b_1^i = \alpha_0^{ii}$, the log-likelihood ratio of the two states given agent i's initial signal. Similarly, $\mathbf{b}_t$ is the vector of the agents' estimates at period t, with its i-th element being agent i's estimate b_t^i. Let $\mathbf{A}$ be the adjacency matrix of the network and its il-th entry be $\mathbf{A}_{il}$. That is, $\mathbf{A}_{il} = 1$ if $il \in G$ and $\mathbf{A}_{il} = 0$ otherwise. Also, $\mathbf{A}_{ii} = 0$, and $\mathbf{I}$ is the identity matrix.

If all agents have low cognitive ability, they count each change in a neighbor's reports as a new signal. Observe that $\mathbf{A}_{il}$ measures whether agent l can learn agent i's initial signal in one period, since it equals 1 if and only if $l \in \mathrm{N}_i$. Similarly, $(\mathbf{A}^k)_{il}$, the il-th entry of matrix $\mathbf{A}^k$, measures how many walks of length k the initial signal can travel from i to l.[26] It precisely measures how agents with $c_i = 1$ learn. Agents' estimates in period t are:

$$\mathbf{b}_t = \left(\sum_{k=0}^{t-1} \mathbf{A}^k \right) \mathbf{b}_1. \tag{7}$$

Proposition 2. *Consider a network with at least one circle of an odd length. Suppose all agents have low cognitive ability; then there exists a vector* $\mathbf{w}$ *such that if $\sum_i w_i \alpha_0^{ii} \neq 0$, agents' estimates converge as $t \to \infty$.*

Agents' estimates are given by (7). As time goes on, the new information extracted by each agent with low cognitive ability contains an increasing number of copies of all agents' initial signals. The number of duplicated copies grow at the same exponential rate. While agents may disagree in the beginning as their initial signals travel through the network, their eventual estimates are dominated by the same weighted average of their initial estimates.[27] The network structure

[22] Formally, let $d(il)$ be the *distance*, or the length of the shortest path, between agent i and agent $l \in g$, with $d(ii) = 0$. One period is required for agent l to incorporate a private signal into his report, and $d(il)$ periods are required for the signal to travel from l to i. Therefore, at the beginning of period t, agent l's signals that can reach agent i are those signals agent l received up to and including period $t - d(il) - 1$.

[23] In particular, Molavi et al. (2018) examine a general class of naive learning procedures in which every agent uses the same function to aggregate her neighbors' reports (in the form of the log-likelihood ratios). They show that these agents learn asymptotically if this function is homogeneous of degree one. Intuitively, this condition translates into unanimity: an agent stops learning if she and all her neighbors agree with each other. In the $c_i = 0$ case, each agent treats her neighbors' reports as new information in every period. They keep learning from their neighbors even when they are in agreement (as in Table 1), and thus they may fail to learn asymptotically. In the DeGroot model, however, agents repeatedly take the weighted average of their neighbors' opinions, and their opinions do not change if they agree with each other. Thus, it is possible for DeGroot agents to learn asymptotically, as shown by Golub and Jackson (2010).

[24] For example, after the Brexit referendum, only 29% of the public correctly believe that European immigrants pay £4.7b more in taxes than what they receive in benefits and services. These wrong beliefs also vary greatly among different groups. See "*Brexit Misperceptions*," the October 2018 study by the Policy Institute at King's College (https://www.kcl.ac.uk/policy-institute/research-analysis/the-publics-brexit-misperceptions).

[25] Throughout this paper, we use bold font for vectors.

[26] A walk from i to l is a sequence of (possibly repeated) agents $(i_0 i_1 \ldots i_k)$ such that $i_0 = i$, $i_k = l$ and $i_h i_{h+1} \in G$ for all $h < k$.

[27] We rule out the indeterminate knife-edge case when $\sum_i w_i \alpha_0^{ii} = 0$. Thus agents believe in the most likely state given a weighted average of their estimates, where vector $\mathbf{w}$ is proportional to the Perron-Frobenius left eigenvector of $\mathbf{A}$.

determines the weights agents use. For instance, in a clique, the weights are equal and thus all agents believe in the state most likely given the initial signals.

Proposition 2 immediately implies that agents with low cognitive ability may all believe in the wrong state. This occurs if an influential agent—whose estimate every agent places a large weight on—gets a wrong signal. One may attribute these learning errors to the fact that the number of initial signals is finite. But this intuition is incomplete: the Law of Large Numbers (LLN) may fail when all agents have low cognitive ability.

Proposition 3. *In any network with three or more agents, if $c_i = 1$ for all agents, their learning outcomes are not asymptotically correct with a positive probability.*

We prove this result by construction: when some agents receive wrong initial signals, all agents may believe in the wrong state even if they receive an arbitrarily large number of correct signals later. Intuitively, the duplicate copies of the initial wrong signals are repeatedly learned by the agents because of their low cognitive ability. These duplicate copies then travel through the network, growing at an exponential rate and dominating all the later correct signals. Such failure of LLN could also occur when only a subset of agents have low cognitive ability.

We use a 4-agent circle to show that agents with moderate cognitive ability can learn asymptotically correctly, but they cannot do so with low cognitive ability.

Example 2. Four agents are connected in a circle $C = (1234)$. Recall from Example 1 that $S = \{s_1, s_2\}$, $X = \{x^{\emptyset}, x^1, x^2\}$, and $\log\left(\Pr(s_1 \mid x^1)/\Pr(s_2 \mid x^1)\right) = \varphi > 0$. The true state is s_1. Consider the following sequence of realized signals X_T: $x_0^i = x^2$ and $x_t^i = x^1$ for all $t \geq 1$ and $i \in \{1, 2, 3, 4\}$. The correct Bayesian posterior is $\Pr(s_1 | X_T) = 1$.

Clearly, x^1 (resp. x^2) is more indicative of state s_1 (resp. s_2). Under X_T, everyone receives a wrong signal initially, but receives infinitely many correct signals x^1 from period 1 onward. At $t = 1$, $b_1^i(s_1) = -\varphi$ for all four agents.

Observation*: asymptotically correct learning when $c_i \geq 2$ for all agents.* At $t = 2$, each agent learns two x^2, one from each neighbor, and one x^1 from nature, so $b_2^i = -2\varphi$. With $c_i = 2$, agent i expects each neighbor to learn one copy of x^2 from her and to add it to his initial signal of x^2, thus $b_2^{ij} = -2\varphi$. At $t = 3$, agent i learns nothing from her neighbors because $b_2^{ij} = b_2^j$, and one x^1 from nature, and thus $b_3^i = -\varphi$. At $t = 4$, agent i learns two x^1, one from each neighbor, and an additional x^1 from nature, and thus $b_4^i = 2\varphi$. For all $t \geq 5$, agents learn (multiple copies of) x^1 from each neighbor and one x^1 from nature. In the limit, they believe correctly that the true state is s_1.

Observation*: wrong learning outcomes when $c_i = 1$ for all agents.* At $t = 2$, agent i's estimate remains $b_2^i = -2\varphi$. But when $c_i = 1$, agent i believes that each neighbors' previous report represents his old information, and thus $b_2^{ij} = -\varphi$. Then, she learns one copy of x^2 from each neighbor in the beginning of $t = 3$. Similarly, in each period $t \geq 3$, each agent learns one x^2 from each neighbor, and one x^1 from nature, and thus $b_t^i = -t\varphi$. In the limit, they believe incorrectly that the true state is s_2. ◇

While agents with low cognitive ability make many learning errors, even this rudimentary ability to identify old information can reduce learning errors significantly. To see this, we compare the learning outcomes of these agents with the naive learning agents. Recall that naive agents treat every report as a new signal. Then agent i's estimate in period t is formed by adding up all the period-$(t-1)$ estimates in her local network:

$$\mathbf{b}_t = (\mathbf{A} + \mathbf{I})\mathbf{b}_{t-1} \ \Rightarrow\ \mathbf{b}_t = (\mathbf{A} + \mathbf{I})^{t-1}\mathbf{b}_1. \tag{8}$$

In particular, the il-th entry of $(\mathbf{A} + \mathbf{I})^{t-1}$ reflects the number of copies of signal x_0^i that agent l extracted by period $t-1$.[28] Given how agents' estimates evolve when they are naive and when they have low cognitive ability (expression (8) and (7)), we have:

Corollary 2. *Let x_0^i be an informative signal. Then as $t \to \infty$, the difference between the number of copies of x_0^i each agent believes in when they are naive and that when they have low cognitive ability ($c_i = 1$) goes to infinity.*

The difference in the number of copies of x_0^i each agent l believes in when agents are naive and when they have low cognitive ability is the il-th entry of $\sum_{k=0}^{t-1}\left(\binom{t-1}{k} - 1\right)\mathbf{A}^k$ by the Binomial Theorem. This difference clearly increases in t and grows to infinity as $t \to \infty$. This result suggests in actual experiments, which typically last a few rounds, agents will low cognitive ability make far fewer errors than naive-learning agents.

[28] Levy and Razin (2018) show that as $t \to \infty$, all agents believe in the state that is most likely given a weighted average of the initial beliefs, where the weights are the Perron-Frobenius eigenvector of $\mathbf{A} + \mathbf{I}$.

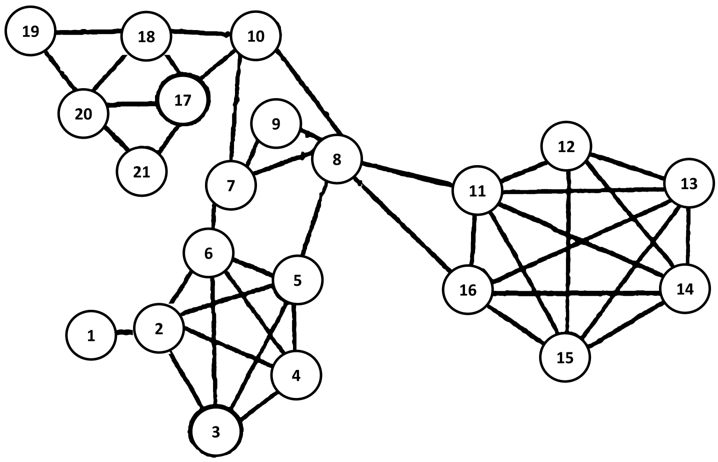

Fig. 3. One component of the high-school friendship network.

4. Main property: a cutoff level of cognitive ability

We now explore how an agent's learning outcomes depend on her cognitive ability more generally. First, as an agent's cognitive ability increases, is there a limit as to how well she can learn? Second, how does this limit depend on the network structure? We begin with defining a cutoff level of cognitive ability for each agent. It is worth noting that each agent's cutoff is *uniquely determined by her (local) network structure*.

Definition 3. Let $\hat{c}_i$ be the **cutoff level of cognitive ability** of agent i such that:

(1) $\hat{c}_i = 2$: if agent i and her neighbor j's shared local network $\left(g_{ij}, G_{ij}\right)$ is a clique for all $j \in \mathrm{N}_i$; otherwise,
(2) $\hat{c}_i = k$: if k is the smallest integer such that the shared local network of agent i and any of her $k-1$ neighbors $\left(g_{ii_1 \ldots i_{k-1}}, G_{ii_1 \ldots i_{k-1}}\right)$ is a clique.

We illustrate this definition using Fig. 3, which contains one component of the high-school friendship network from Coleman (1961).[29]

Begin with agent 1: her cutoff level is $\hat{c}_1 = 2$ because she only has one neighbor and their shared local network includes only themselves, which is a clique. Agent 2's cutoff is also $\hat{c}_2 = 2$ because his shared local network with each neighbor is a clique: (g_{12}, G_{12}) is a line of two agents, and each (g_{2h}, G_{2h}) for $h \in \{3, 4, 5, 6\}$ is a clique of five agents. Similarly, in a social quilt, every agent has a cutoff level of 2 because every pair of agents' shared local network is a clique.

When does an agent's cutoff get higher than 2? Consider agent 7 in Fig. 3. Her shared local network with neighbor 8 contains agents $\{7, 8, 9, 10\}$. It is not a clique because 9 and 10 are not connected. Clearly, the shared local network of $\{7, 8, 9\}$ (and also $\{7, 8, 10\}$) is a clique. Then by part (2) of Definition 3, $\hat{c}_7 = 3$. Interestingly, among the 146 students in this high-school friendship network, 119 have a cutoff of 2, 23 have a cutoff of 3, and only 4 of them have a cutoff of 5. These four students are $\{11, 13, 14, 15\}$ in Fig. 3.[30]

We now show that agent i's learning outcomes are the same, *ceteris paribus*, for all cognitive abilities above this cutoff.

Proposition 4. *(1) Agent i's estimates are the same when her cognitive ability is above her cutoff:* $\mathbf{b}_t^{i,c_i} = \mathbf{b}_t^{i,\hat{c}_i}$ *for all* $c_i \geq \hat{c}_i$ *and all t, holding all other agents' cognitive abilities* ($\mathbf{c}_{-i}$) *and sequence of realized signals* X_T *constant. (2) Agent i's cutoff is tight: for any* $c_i < \hat{c}_i$, $\mathbf{b}_t^{i,c_i} \neq \mathbf{b}_t^{i,\hat{c}_i}$ *for some* $\mathbf{c}_{-i}$*, t, and* X_T.

Proposition 4 is useful because it shows that in real networks, an agent may not need a high cognitive ability to obtain her *best learning outcomes*—her learning outcomes under an arbitrarily high level of cognitive ability. On one hand, if agent i's best learning outcomes are correct, then she learns correctly as long as her cognitive ability is above her cutoff level. But if she is unable to learn correctly at the cutoff level, then her learning outcomes will not improve. On the other hand, if an agent's cognitive ability falls short of her cutoff level ($c_i < \hat{c}_i$), she will make learning errors due to her cognitive constraint. She makes fewer learning errors if she has higher cognitive ability (up to her cutoff level), which in turn will improve

[29] The full network of 146 students as cited by Feld (1991) is reproduced in Fig. A1 in the appendix.

[30] To see why, notice that the shared local network of these four agents include agent 16 and 12, who are not connected themselves. But the shared networks of these four agents together with either agent 12 or agent 16 are cliques. Therefore the cutoff is 5 for these four agents.

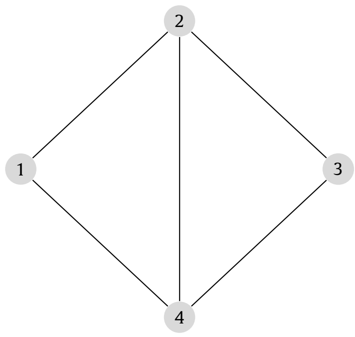

Fig. 4. A diamond with a link.

other agents' learning outcomes. A more subtle implication of Proposition 4 is that agent i does not need to form estimates involving more than $\hat{c}_i$ agents. Therefore our learning rule is even less demanding on agent i's cognitive ability than it first appears.

We prove this proposition through Lemma A1 and A2 in the appendix. Lemma A1 shows why agent i's learning outcomes stay the same when her cognitive ability is above $\hat{c}_i$, holding everything else constant. Since agent i's neighbors are finite, and she only learns from her neighbors, the set of reports she can observe and make use of is finite as well. Thus, it is not surprising that she needs only a finite level of cognitive ability to store old information in her local network. But this intuition is incomplete. In fact, agent i's cutoff does not depend on (and can be much lower than) her *degree*, the number of neighbors she has. Rather, it depends on the structure of her local network. Recall from Definition 3, the shared local network of i and any of her $\hat{c}_i - 1$ neighbors must be a clique. Within a clique, agents observe the same reports and thus form the same estimates. Specifically, any $\hat{c}_i$-agent estimate is the same as any $(\hat{c}_i + 1)$-agent estimate in the clique. Therefore, their degenerate estimates of using the former to replace the latter lead to no additional learning errors.

Lemma A1 explains why Proposition 1 apply to all agents with cognitive ability $c_i \geq 2$. Every agent in a social quilt has a cutoff of $\hat{c}_i = 2$, which is enough for every agent to learn correctly. But whether agent i's learning outcomes are correct when her cognitive ability is at (or above) her cutoff level depends on the entire network structure. This is because each agent only knows and learns from her local network; she cannot avoid learning errors due to the network structure unknown to her. For instance, in the 4-agent circle of Example 2, each agent also has a cutoff level of $\hat{c}_i = 2$. But they make learning errors even with moderate (or higher) cognitive abilities, because they don't know old information can travel around the circle and reach them again. In other words, if every agent is at (or above) her cutoff level of cognitive ability, their learning errors (if any) are caused by their network knowledge constraint. Such errors persist even if agents are more sophisticated.

Next, Lemma A2 shows that the cutoff is tight. If her cognitive ability falls short of her cutoff, agent i's degenerate estimates can be based on a wrong set of shared reports. Specifically, agent i has to use degenerate estimate as in expression (5) to set her $(c_i + 1)$-agent estimate, which is done by removing an agent from the $(c_i + 1)$ agents. When $c_i < \hat{c}_i$, the shared local network of the remaining c_i agents may not be a clique, and thus their shared reports may differ from those shared by the original $(c_i + 1)$ agents. Then, her multi-agent estimates would differ from the ones she would have formed if she is not cognitively constrained. Once her $(c_i + 1)$-agent estimate is wrong, it affects her c_i-agent, ..., 2-agent estimates iteratively and leads to wrong estimates. To illustrate, consider the network in Fig. 4, the same as the subgraph of agents $\{7, 8, 9, 10\}$ in Fig. 3.

Recall that $\hat{c}_1 = \hat{c}_3 = 2$ while $\hat{c}_2 = \hat{c}_4 = 3$. If $c_2 = 2$, then agent 2 cannot form her 3-agent estimate b_t^{234}. Instead, she uses degenerate estimate $b_t^{234} = b_t^{24}$. But the reports seen by agent 2 and 4 are from agents $\{1, 2, 3, 4\}$, while the reports seen by $\{2, 3, 4\}$ are from the three of them. If agent 1 gets an informative signal and reports it, these two sets of reports differ. Consequently, agent 2 makes learning errors. We now contrast the learning outcomes when $c_i \geq 3$ and when $c_i = 2$ for all the agents in Fig. 4.

Example 3. Consider the diamond with a link network in Fig. 4. Let $x_0^1 = x_0^3 = x^1$ be the only informative signals. As before, $\log\left(\Pr(s_1 \mid x^1)/\Pr(s_2 \mid x^1)\right) = \varphi > 0$. The log-likelihood ratio of the state distribution based on η copies of x^1 is $\eta\varphi$. When $\eta < 0$, it means that the agent believes in $|\eta|$ copies of x^2 instead.

At $t = 1$, $b_1^1 = b_1^3 = \varphi$, and $b_1^2 = b_1^4 = 0$. At $t = 2$, $b_2^1 = b_2^3 = \varphi$, and $b_2^2 = b_2^4 = 2\varphi$. Then their learning outcomes depend on their cognitive abilities.

Observation*: correct learning for agent* 2 *and* 4 *when* $c_i \geq 3$. Agent 2 and 4 observe the entire network and have sufficient cognitive abilities since $\hat{c}_2 = \hat{c}_4 \geq 3$. They know their reports $b_2^2 = b_2^4 = 2\varphi$ are based on signals from agent 1 and 3. They do not learn from each other and their estimates remain at 2φ for all $t \geq 2$. Agent 2 and 4 also know agent 1 and 3 do not know each other, and expect agent 1 and 3 to make the following learning errors.

Table 4
Learning in a diamond with a link with $c_i = 3$.

	$b^1_t = b^3_t$	$b^2_t = b^4_t$	$\alpha^{12}_{t-1} = \alpha^{14}_{t-1}$
$t = 1$	φ	0	n/a
$t = 2$	φ	2φ	0
$t = 2\tau + 1, \tau \in \mathbb{N}$	3φ	2φ	φ
$t = 2\tau + 2, \tau \in \mathbb{N}$	φ	2φ	$-\varphi$

Since agent 1 and 3 are symmetric, we focus on agent 1. At $t = 3$, agent 1 learns two copies of x^1, one from each neighbor, so $b^1_3 = 3\varphi$. She also believes that agent 2 should learn from agent 4 just like she learns from agent 4, and thus $b^{12}_3 = b^{14}_3 = 3\varphi$. At $t = 4$, she learns two copies of x^2 (the opposite of x^1), one from each neighbor. This is because $b^2_3 = 2\varphi$ while $b^{12}_3 = 3\varphi$, and thus $\alpha^{12}_3 = b^2_3 - b^{12}_3 = -\varphi$. That is, she rationalizes the fact that agent 2 and 4 report two, instead of three, copies of x^1 by believing each of them has received a copy of x^2. So $b^1_4 = \varphi$. She also believes agent 2 should learn from agent 4 just like she learns from 4, and thus $b^{12}_4 = b^{14}_4 = \varphi$. For all $t \geq 3$, agent 1 and 3's estimates oscillate between 3φ in odd periods and φ in even periods. Their learning outcomes are summarized in Table 4.

***Observation**: diverging opinion swings when $c_i = 2$ for all i.* We keep track of only agent 1 and 2 due to symmetry. At $t = 3$, $b^1_3 = 3\varphi$ and $b^2_3 = 2\varphi$ as above. But now agent 2 does not have the ability to reason on behalf of agent 1 for agent 4. Instead, she sets $b^{214}_2 = b^{24}_2 = 2\varphi$ using her degenerate estimate. Then, agent 2 believes agent 1 does not learn from agent 4 just like she does not learn from 4, and thus $b^{21}_3 = 2\varphi$. At $t = 4$, agent 2 learns two copies of x^1, one from agent 1 and one from agent 3: $b^2_4 = 4\varphi$. Agent 2 learns x^1 from agent 1 because she thought agent 1 should report 2φ but he reports 3φ: $b^1_3 = 3\varphi$ while $b^{21}_3 = 2\varphi$. Their opinions oscillate and diverge over time (derivations are in the appendix):

$$(b^1_5, b^2_5) = (7\varphi, 0), (b^1_6, b^2_6) = (-7\varphi, 10\varphi), (b^1_7, b^2_7) = (27\varphi, -19\varphi), \ldots$$

As time goes on, *all* agents' estimates oscillate and move away from the log-likelihood ratio of the Bayesian posterior, which is 2φ. Not only agent 2 and 4 fail to account for the oscillations of agent 1 and 3, they exacerbate such errors. $\diamond$

Local improvement in an agent's cognitive ability could thus be valuable for the whole network. That is, if an agent falls short of her cutoff level of cognitive ability, then as her cognitive ability increases, she makes fewer learning errors. Moreover, these learning errors do not spread to her neighbors, who in turn do not mislead their neighbors. As a result, all agents make fewer learning errors. In the example above, all agents exhibit diverging belief oscillations when $c_2 = c_4 = 2$. But when $c_2 = c_4 \geq 3$, only agent 1 and 3 have belief oscillations which are limited in scope.

5. Discussions and concluding remarks

5.1. Decisive evidence

Can agents learn correctly regardless of their cognitive abilities and the network structure? The answer is yes, if every agent's signal provides *decisive evidence*, which occurs when all agents' information structures are partitional over S, the set of possible states.[31] Intuitively, decisive evidence requires that each informative initial signal rules out some states completely. Let $S = \{s_1, \ldots, s_n\}$ in this extension and consider initial signals only. Then as in Aumann (1976), agent i's information structure can be represented by a mapping $\mathcal{P}^i : S \to 2^S \setminus \emptyset$, associating each state s_n with a non-empty element $\mathcal{P}^i(s_n)$. At s_n, agent i considers $\mathcal{P}^i(s_n)$ to be the set of possible states. Moreover, $\mathcal{P}^i$ induces an information partition over S: for any $s_n \in S$, $s_n \in \mathcal{P}^i(s_n)$, and for any $s_n, s_{n'} \in S$, $s_{n'} \in \mathcal{P}^i(s_n)$ implies $\mathcal{P}^i(s_{n'}) = \mathcal{P}^i(s_n)$.

Proposition 5. *If every agent has decisive evidence, then agents of all cognitive abilities learn correctly in any network.*

Proposition 5 holds due to a special feature of decisive evidence: agents may make mistakes identifying old information, but they still learn correctly. All that an agent needs is the ability to use Bayes' rule in a rudimentary manner—to remove impossible states given her signal and her neighbors' reports. Thus, Proposition 5 applies to agents with low cognitive ability as well as naive-learning agents. Specifically, whenever agent i extracts a new signal from a neighbor, she eliminates the states deemed impossible according to the new signal. Even if agent i extracts the same signal from multiple neighbors unknowingly, she removes the same set of states as she would have given one such signal. Thus, the agents' learning outcomes are correct: they believe in the states in the intersection of all possible states indicated by their initial signals.[32]

[31] In the classic Bayesian framework, agents have a common prior and update their beliefs over an *expanded space* containing each agent's beliefs about S, her beliefs about the network, and all her higher-order beliefs about their neighbors' beliefs. In general, this expanded space is also partitional. To emphasize that our partition is only over S, we call it decisive evidence instead of partitional information.

[32] This is also related to the literature on knowledge and consensus which studies under which conditions and reporting protocols, repeated communication among a finite set of individuals leads to consensus. See Aumann (1976), Geanakoplos and Polemarchakis (1982), Parikh and Krasucki (1990), Mueller-Frank (2013), among many others.

The same logic holds more generally: if any agent's signal rules out a state, all agents learn it correctly. This result suggests another way to reduce learning errors in a complex network. Instead of targeting influential agents as in the optimal seeding problem, sending a few agents with decisive evidence may be more effective.

5.2. When agents know their neighbors' cognitive abilities

So far, agent i learns according to her cognitive ability only. If agents know their neighbors' cognitive abilities, then they can apply our learning procedure and account for their neighbors' cognitive abilities explicitly. Similar to the level-k learning literature, we assume that an agent with a lower cognitive ability cannot understand how her neighbor with equal or higher cognitive ability learns. So she just follows the learning rule for her cognitive ability like before. But an agent with a higher cognitive ability knows how each neighbor with a lower cognitive ability learns, and accounts for his learning errors. This additional ability to detect learning errors can improve agents' learning outcomes, as illustrated below.

Example 4. Consider the two agents in Example 1. Recall that x_0^1 is the only informative signal. Suppose that $c_1 = 1$, $c_2 = 2$, and their cognitive levels are common knowledge.

At $t = 1$, Ann reports $b_1^1 = \alpha_0^{ii} = \varphi$ and Bob reports $b_1^2 = 0$. At $t = 2$, Ann learns nothing from Bob and reports $b_2^1 = \varphi$, while Bob learns a signal from Ann and reports $b_2^2 = \varphi$. At $t = 3$, Ann with $c_1 = 1$ treats the change in Bob's report as a new signal, and reports $b_3^1 = 2\varphi$, while Bob remains at $b_3^2 = \varphi$. At $t = 4$, there is no change in Bob's report, so Ann remains at $b_4^1 = 2\varphi$. The key is that Bob knows $c_1 = 1$ and expects Ann to treat his change from b_1^2 to b_2^2 as a new signal in period 3, and thus expects $b_3^1 = 2\varphi$. So Bob learns nothing and reports $b_4^2 = \varphi$. For $t \geq 5$, $b_t^1 = 2\varphi$ and $b_t^2 = \varphi$. Bob learns correctly despite Ann's (single) learning error. ⋄

The intuition of this example also extends to Proposition 1: if cognitive abilities are known, one low cognitive ability agent needs not mislead the entire network.

Corollary 3. *Suppose that agents can account for neighbors' cognitive abilities. All agents with $c_l \geq 2$ learn correctly in a social quilt if agents with $c_i < 2$ are isolated. That is, each agent with $c_i < 2$ has only one neighbor, and that neighbor j has $c_j \geq 2$.*

While an agent with $c_i = 1$ still makes mistakes by counting all changes in her neighbor j's report as new, her errors no longer affect the rest of the network if she is isolated. Her report reaches her neighbor j, who understands and only passes on the part of agent i's report that is truly due to new signals. Moreover, because agents with moderate ability (or above) form consensus eventually, errors made by agents with low cognitive ability are bounded. Intuitively, if agents who pass on information from one subnetwork to another have sufficient cognitive ability, they can account for errors made by the fringe.

5.3. Conclusion

Learning in networks naturally requires iterative and sophisticated reasoning because each agent needs to account for what a neighbor has learned from another in order to identify old information. Our novel framework equates cognitive ability with depth of reasoning, namely, whether an agent can estimate what one neighbor has learned from another neighbor given their shared history of reports. We show that agents with moderate cognitive ability may learn in some networks, but low cognitive ability causes and reinforces learning errors in general. In real networks, however, agents may not need high cognitive ability to achieve their best learning outcomes.

Theoretically, we fill a gap between models with locally Bayesian agents who have arbitrarily high cognitive ability and models with naive learning agents. In terms of methodology, agents in our model remain true to Bayesian learning in spirit: they try to extract new information to the best of their cognitive ability using Bayes-like formulas. One may alternatively model an agent's limited cognitive ability as having a finite memory state space for any history of reports, which is a topic for further research (see for instance Hellman and Cover (1970), Piccione and Rubinstein (1997) and Wilson (2014)).

Appendix A. An extension and proofs

A.1. An extended learning rule with repeat agents

In the cognitively-constrained learning rule described in Section 2.1, agents only form (non-degenerate) estimates involving distinct agents. Recall that agent i forms all the estimates involving up to c_i agents using a Bayes-like formula, that is, $\mathbf{b}_t^{ii_1\ldots i_{k-1}}$ with $k \leq c_i$. She then uses two types of degenerate estimates. The first type is used because she cannot form estimates involving more than c_i agents due to her cognitive ability. That is, agent i sets $\mathbf{b}_t^{ii_1 i_2\ldots i_{c_i}} = \mathbf{b}_t^{ii_2\ldots i_{c_i}}$. The second type involves a repeat agent as the last agent, and agent i sets $\mathbf{b}_t^{ii_1\ldots i_{k-1}h} = \mathbf{b}_t^{ii_1\ldots i_{k-1}}$ for any $h \in \{i, i_1, \ldots, i_{k-1}\}$.

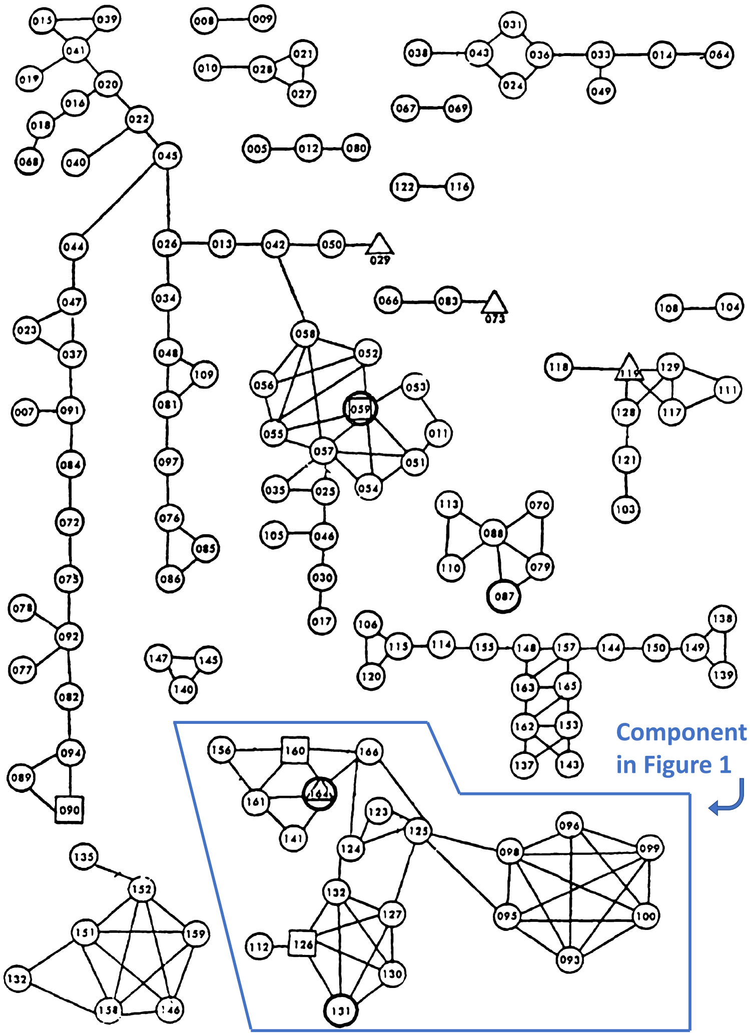

Fig. A1. High school friendship network in Coleman (1961).

The first type of degenerate estimate is due to her cognitive constraint. One may wonder why we include the second type of degenerate estimate and whether it affects agents' learning outcomes. We now present an extended learning rule allowing for estimates of *any* sequence of agents, including repeat agents. We will show that as long as $c_i \geq \hat{c}_i$, agent i's learning outcomes using this extended learning rule are the same as those using the cognitively-constrained learning rule in Section 2.1 involving only distinct agents. In this case, our cognitively-constrained learning rule is without loss. If $c_i < \hat{c}_i$ instead, Proposition 4 shows that the first type of degenerate estimate can lead to additional learning errors, which remains true in this extended learning rule.

Formally, suppose there are multiple possible states $S = \{s_1, \ldots, s_n\}$, and let $\boldsymbol{\beta}^i_t = \left(\beta^i_t(s_1), \ldots, \beta^i_t(s_N)\right)$ be agent i's estimate under this extended learning rule. Similar to the agent's estimate in the text, it is a vector of the log-likelihood ratios of the state probabilities: $\beta^i_t(s_n) = \log p^i_t(s_n) - \log p^i_t(s_N)$ for all s_n.[33] Let a sequence of agents be *fully-connected* if it contains at least two distinct agents and every pair of distinct agents is connected. Agent i's multi-agent estimates $\boldsymbol{\beta}^{ij}_t, \ldots, \boldsymbol{\beta}^{ij\ldots l}_t$ for each fully-connected sequence $(ij \ldots l)$ are similarly defined. As in formula (3), agent i extracts new information from her neighbors using Bayes-like formulas. For example, agent i extracts the following from neighbor j,

$$\boldsymbol{\alpha}^{ij}_{t-1} = \boldsymbol{\beta}^{j}_{t-1} - \boldsymbol{\beta}^{ij}_{t-1}.$$

Moreover, let the *length* of a sequence of agents $(ij \ldots l)$ be the total number of agents, including repeat agents, in the sequence. Let $distinct(ij \ldots l) \subseteq g_i$ be the set of all the distinct agents in this sequence.

[33] We abuse notations slightly by continue to use $\mathbf{p}$ for beliefs and $\boldsymbol{\alpha}$ for extracted new signals in this extended learning rule. They simplify exposition and should be clear in context.

The learning rule for $c_i = 0, 1$ is the same as that in Section 2.1. If $c_i \geq 2$, the extended learning rule is as follows.

Initial values. At $t = 0$, agent i receives signal x_0^i. At $t = 1$, agent i learns as in Section 2.1. In addition, all her estimates involving two or more agents are $\boldsymbol{\beta}_1^{ij...l} = \mathbf{0}$.

For all $t \geq 2$. If $k \leq c_i$, agent i forms all her period-t estimates $\boldsymbol{\beta}_t^i$ and k-agent estimates $\boldsymbol{\beta}_t^{ii_1...i_{k-1}}$ involving k distinct agents in the same way using formulas (1), (2), (3) and (4). Moreover, for any sequence of fully-connected agents $(ij\ldots l)$ with a length up to c_i, she learns in the following two steps similar to those in Section 2.1.

Signal extraction rule-extended: she extracts new information similar to formula (3): for each $h \in g_{ij...l}$,

$$\boldsymbol{\alpha}_{t-1}^{ij...lh} = \boldsymbol{\beta}_{t-1}^{h} - \boldsymbol{\beta}_{t-1}^{ij...lh}. \tag{9}$$

Updating rule-extended: she incorporates the information from above similar to (4):

$$\boldsymbol{\beta}_t^{ij...l} = \boldsymbol{\beta}_{t-1}^{ij...l} + \sum_{h \in g_{ij...l}} \boldsymbol{\alpha}_{t-1}^{ij...lh}. \tag{10}$$

Type 1A degenerate estimate: binding cognitive constraint. In any sequence $(ij\ldots lh)$ of $(c_i + 1)$ fully-connected agents, agent i sets her $(c_i + 1)$-agent estimate equal to her c_i-agent estimate as follows. First, if there is a repeat agent in the sequence, then agent i sets $\boldsymbol{\beta}_t^{ij...lh}$ to be her c_i-agent estimate of the sequence with the last repeat agent removed. Otherwise, agent i uses the same degenerate estimate as in equation (5), that is, removing the second agent (agent j) from the sequence:

$$\boldsymbol{\beta}_t^{ij...lh} = \boldsymbol{\beta}_t^{i...lh}. \tag{11}$$

Similarly, this can be interpreted as agent i behaves as if she and agent j have the same estimate about what the other agents in the sequence believe. Agent i does not form estimates involving $(c_i + 2)$ or more agents. ‖

Because this extended learning rule allows for repeat agents explicitly, agent i only needs to form degenerate estimates if her cognitive constraint is binding. We now show that neither the repeat agents nor the order of agents in the sequence matters in agent i's learning outcomes as long as her cognitive ability is above her cutoff. To change the order of agents, let $\sigma(\cdot)$ be the permutation operator such that $\{i_{\sigma(1)}, \ldots, i_{\sigma(k-1)}\}$ is a permutation of agents $\{i_1, \ldots, i_{k-1}\}$. Then we have:

Proposition A1. *Suppose that $c_i \geq \hat{c}_i$. Then for all t, X_T and $\mathbf{c}_{-i}$,*

1. *repeat agents do not matter: for any sequence of fully-connected agents $(ij\ldots l)$ with a length up to c_i, if distinct$(ij\ldots l) = \{i, i_1, \ldots, i_{k-1}\}$, then $\boldsymbol{\beta}_t^{ij...l} = \boldsymbol{\beta}_t^{ii_1...i_{k-1}}$;*
2. *order of agents does not matter: for any clique $\{i, i_1, \ldots, i_{k-1}\}$ with $k \geq \hat{c}_i$ and any permutation $\sigma(\cdot)$ of these agents in the clique, let $k' = \min(k, c_i)$, then*

$$\boldsymbol{\beta}_t^{ii_1...i_{\hat{c}_i-1}} = \boldsymbol{\beta}_t^{ii_{\sigma(1)}...i_{\sigma(\hat{c}_i-1)}} = \boldsymbol{\beta}_t^{ii_{\sigma(1)}...i_{\sigma(\hat{c}_i)}} = \ldots = \boldsymbol{\beta}_t^{ii_{\sigma(1)}...i_{\sigma(k'-1)}}.$$

Part 1 of Proposition A1 shows that repeat agents do not affect agent i's estimates as long as her cognitive ability is above the cutoff level. To be concrete, recall agent 2 in Fig. 3 with $\hat{c}_2 = 2$, and the clique $\{2, 3, 4, 5, 6\}$ she belongs to. Let her cognitive ability be $c_2 = 5$. Then, part 1 implies that for instance, $\boldsymbol{\beta}_t^{23432} = \boldsymbol{\beta}_t^{234}$. This result motivates us to remove the repeat agents from the learning rule in the text because it is simpler. Part 2 of Proposition A1 shows that forming estimates involving more than $\hat{c}_i$ agents does not change agent i's learning outcomes because they are identical to her corresponding $\hat{c}_i$-agent estimates. For example, agent 2 above has $\boldsymbol{\beta}_t^{23} = \boldsymbol{\beta}_t^{24} = \boldsymbol{\beta}_t^{256} = \boldsymbol{\beta}_t^{2654} = \boldsymbol{\beta}_t^{23456}$. Intuitively, each subset of g_i involving more than $\hat{c}_i$ agents is fully connected, and thus her $\hat{c}_i$-agent estimates already capture all the information in this subset. Together, this result shows that agent i needs to form only up to her $\hat{c}_i$-agent estimates involving distinct agents if $c_i \geq \hat{c}_i$.

The proposition also shows that the order of agents in estimates involving $\hat{c}_i$ or more agents does not matter. Thus as long as $c_i \geq \hat{c}_i$, agent i's learning outcomes are unaffected by how she sets the first type of degenerate estimates, that is, which agent to remove. On the other hand, if $c_i < \hat{c}_i$, the binding cognitive constraint will lead agent i to make more learning errors because she has to remove one agent from her $(c_i + 1)$-agent estimates. Regardless of which agent she removes, this affects her degenerate estimates because each subset of agents may have access to reports that differ from the reports observed by all the $(c_i + 1)$ agents. Since the qualitative nature of the learning error is similar, we remove the second agent (agent i_1) in the learning rule in the Section 2.1.

Before the proof, we highlight an immediate corollary of this proposition.

Corollary A1. *Suppose that $c_i \geq \hat{c}_i$. Then for all t, X_T and $\mathbf{c}_{-i}$, agent i's estimates are the same under both learning rules: $\boldsymbol{\beta}_t^i = \mathbf{b}_t^i$.*

By part 1 and 2 of Proposition A1, neither type of degenerate estimate affects i's learning outcomes as long as $c_i \geq \hat{c}_i$. Then, her multi-agent estimates involving only distinct agents are the same under either learning rule: $\boldsymbol{\beta}_t^{ii_1\ldots i_{k-1}} = \mathbf{b}_t^{ii_1\ldots i_{k-1}}$ for $k \leq c_i$. Her estimates are then iteratively derived from these multi-agent estimates. Therefore, her estimates under these two learning rules are the same when $c_i \geq \hat{c}_i$.

Proof of Proposition A1. First, consider any sequence of agents $(ij\ldots l)$ with a length smaller or equal to c_i. The shared local network of $(ij\ldots l)$ is equal to $(g_{ii_1\ldots i_{k-1}}, G_{ii_1\ldots i_{k-1}})$ because they have the same distinct agents. Second, by definition, the shared local networks of all cliques of $\hat{c}_i$ agents or above are fully connected. Within agent i's local network, we have in each clique of k agents ($\hat{c}_i \leq k$), $g_{ii_1\ldots i_{\hat{c}_i-1}} = g_{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i-1)}} = g_{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i)}} = g_{ii_{\sigma(1)}\ldots i_{\sigma(k-1)}}$. We prove both parts of this proposition by induction on time t.

Clearly, the induction hypothesis is true at $t = 1$. By the assumption of a symmetric prior, all agent i's estimates involving two or more agents are $\mathbf{0}$. Next, suppose both parts of the induction hypothesis hold at period t. Then at $t+1$, for some agent $h \in g_{ii_1\ldots i_{k-1}}$, agent i believes that agent j believes $\ldots$ the signal agent l extracted from agent h is:

$$\boldsymbol{\alpha}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^{h} - \boldsymbol{\beta}_t^{ij\ldots lh}.$$

Similarly, agent i believes that agent i_1 believes $\ldots$ that the signal agent i_{k-1} extracted from agent h is:

$$\boldsymbol{\alpha}_t^{ii_1\ldots i_{k-1}h} = \boldsymbol{\beta}_t^{h} - \boldsymbol{\beta}_t^{ii_1\ldots i_{k-1}h}.$$

We now prove that $\boldsymbol{\beta}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^{ii_1\ldots i_{k-1}h}$, and thus the two extracted signals are identical.

There are two cases. First, if the length of $(ij\ldots lh)$ is smaller or equal to c_i, then $\boldsymbol{\beta}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^{ii_1\ldots i_{k-1}h}$ by our induction hypothesis since both sequences contain the same distinct agents. If the length of the sequence $(ij\ldots lh)$ is equal to $c_i + 1$, then agent i sets the (c_i+1)-agent estimates equal to her c_i-agent estimates as in the learning rule. In particular, if there is a repeat agent in the sequence $(ij\ldots lh)$, then agent i sets $\boldsymbol{\beta}_{t+1}^{ij\ldots lh}$ to be her c_i-agent estimates of the sequence with the last repeat agent removed. These c_i-agent estimates are equal to $\boldsymbol{\beta}_t^{ii_1\ldots i_{k-1}h}$, again by our induction hypothesis since both sequences contain the same distinct agents. Otherwise, all c_i+1 agents in the sequence $(ij\ldots lh)$ must be distinct, that is $distinct(ij\ldots lh) = \{i, i_1, \ldots, i_{k-1}, h\}$. Then to set her degenerate estimates, agent i uses her corresponding c_i-agent estimates by removing the second agent in the sequence. We claim that $\boldsymbol{\beta}_t^{i\ldots lh} = \boldsymbol{\beta}_t^{i\ldots i_{k-1}h}$. To prove this claim, note that if $j = i_1$, then the remaining sequences contain the same set of distinct agents and thus are the same by the induction hypothesis. If $j \neq i_1$, we can first permute the sequence $(ii_1\ldots i_{k-1}h)$ to match the sequence $(ij\ldots lh)$, so that $(ii_{\sigma(1)}\ldots i_{\sigma(k-1)}h)$ equals $(ij\ldots lh)$. Moreover, $(i\ldots i_{\sigma(k-1)}h)$ equals $(i\ldots lh)$ when $i_{\sigma(1)} = j$ is removed. Since there are $c_i \geq \hat{c}_i$ agents in these two sequences after j is removed, then $\boldsymbol{\beta}_t^{i\ldots lh} = \boldsymbol{\beta}_t^{i\ldots i_{\sigma(k-1)}h} = \boldsymbol{\beta}_t^{i\ldots i_{k-1}h}$ by the second part of the induction hypothesis. From expression (10), we have

$$\boldsymbol{\beta}_{t+1}^{ij\ldots l} = \boldsymbol{\beta}_t^{ij\ldots l} + \sum_{h \in g_{ij\ldots l}} \boldsymbol{\alpha}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^{ii_1\ldots i_{k-1}} + \sum_{h \in g_{ii_1\ldots i_{k-1}}} \boldsymbol{\alpha}_t^{ii_1\ldots i_{k-1}h} = \boldsymbol{\beta}_{t+1}^{ii_1\ldots i_{k-1}}. \tag{12}$$

This proves the first part of the proposition.

For the second part, consider any clique of g_i, $\{i, i_1, \ldots, i_{k-1}\}$, with $k \geq \hat{c}_i$ agents. For any $h \in g_{ii_1\ldots i_{\hat{c}_i-1}}$,

$$\boldsymbol{\alpha}_t^{ii_1\ldots i_{\hat{c}_i-1}h} = \boldsymbol{\beta}_t^{h} - \boldsymbol{\beta}_t^{ii_1\ldots i_{\hat{c}_i-1}h},$$

which we claim is equal to

$$\boldsymbol{\alpha}_t^{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i-1)}h} = \boldsymbol{\beta}_t^{h} - \boldsymbol{\beta}_t^{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i-1)}h}.$$

To prove the claim, we show that $\boldsymbol{\beta}_t^{ii_1\ldots i_{\hat{c}_i-1}h} = \boldsymbol{\beta}_t^{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i-1)}h}$. Again consider three cases. First, if $c_i > \hat{c}_i$ and $h \in \{i, i_1, \ldots, i_{\hat{c}_i-1}\} \cup \{i, i_{\sigma(1)}, \ldots, i_{\sigma(\hat{c}_i-1)}\}$, then by the first part of the induction hypotheses, the repeat agent h does not matter. Without h, these estimates are equal by the second part of the induction hypotheses. Second, if $c_i > \hat{c}_i$ and $h \notin \{i, i_1, \ldots, i_{\hat{c}_i-1}\} \cup \{i, i_{\sigma(1)}, \ldots, i_{\sigma(\hat{c}_i-1)}\}$, the estimates are the same by the second part of the induction hypothesis. Third, if $c_i = \hat{c}_i$, then both estimates above involve $c_i + 1$ agents and thus agent i uses the degenerate estimates, that is, the corresponding c_i-agent estimates defined in the learning rule. Using both parts of the induction hypothesis, we can show that these two degenerate estimates are equal regardless of whether agent h is a repeat agent or not. The counterpart of equation (12) shows that $\boldsymbol{\beta}_{t+1}^{ii_1\ldots i_{\hat{c}_i-1}} = \boldsymbol{\beta}_{t+1}^{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i-1)}}$.

In the same way, we can show that estimates with sequences of agents longer than $\hat{c}_i$ but smaller than c_i are the same as the $\hat{c}_i$-agent estimates, $\boldsymbol{\beta}_{t+1}^{ii_{\beta(1)}\ldots i_{\sigma(\hat{c}_i-1)}} = \boldsymbol{\beta}_{t+1}^{ii_{\sigma(1)}\ldots i_{\sigma(\hat{c}_i)}} = \ldots = \boldsymbol{\beta}_{t+1}^{ii_{\sigma(1)}\ldots i_{\sigma(k'-1)}}$. $\square$

A.2. Proofs

We present the proofs in order of appearance, but we use Proposition 4 to prove Proposition 1. Moreover, we prove all propositions for a more general environment with N a priori equally likely states: $\{s_1, \dots, s_N\}$. Thus, the estimates below are all vectors of log-likelihood ratios:

$$\mathbf{b}_t^i = (b_t^i(s_1), \dots, b_t^i(s_N)), \text{ where } b_t^i(s_n) = \log p_t^i(s_n) - \log p_t^i(s_N).$$

Proof of Proposition 1. For sufficiency, we first show that in a social quilt, the cutoff cognitive ability is $\hat{c}_i = 2$ for all agents. There are a few cases. First, for any neighbor $j \in \mathrm{N}_i$ such that $\mathrm{N}_i \cap \mathrm{N}_j = \emptyset$, $g_{ij} = \{i, j\}$ must be a clique. Second, for any neighbor $j \in \mathrm{N}_i$ such that $\mathrm{N}_i \cap \mathrm{N}_j = \{k\}$, $g_{ij} = \{i, j, k\}$ also must be a clique. Third, each pair of agent i and j's common neighbors k and l are connected. Otherwise, there exists a circle $(ikjl)$ such that two agents are not connected, which violates the definition of a social quilt. Therefore $\hat{c}_i = 2$. By Proposition 1 of Li and Tan (2020), agents learn correctly in a social quilt when they can form estimates over any sequence of neighbors in their local networks, that is, when they have arbitrarily high cognitive abilities. Then, our Proposition 4 shows that the agents' learning outcomes are the same if $c_i \geq \hat{c}_i$ for all agents. Their result and our Proposition 4 together imply that all agents learn correctly if $c_i \geq 2$ for all agents.

For necessity, suppose that at least one agent, say agent i, has $c_i = 1$. Let agent i receive the only informative signal x_0^i. She reports it at $t = 1$. Since the network is path-connected, she has at least one neighbor, agent j. At $t = 2$, agent j learns the signal: $\mathbf{b}_2^j = \mathbf{b}_1^i$. Since $c_i = 1$, agent i treats agent j's change in report as a new signal and incorporates it at $t = 3$, and she does the same for all her neighbors: $\mathbf{b}_3^i = (1 + |\mathrm{N}_i|)\mathbf{b}_1^i$. The change in agent i's report in turn is treated as a new signal by each neighbor at $t = 4$ regardless of that neighbor's cognitive ability. In this way, in each even period from $t = 2$ onward, each neighbor of i extracts at least one more copy of the initial signal, and passes it on to their neighbors. In addition, more copies of the initial signal are extracted if more agents have a cognitive ability of 1. As $t \to \infty$, everyone believes in the state(s) most likely given the initial signal with probability 1. But the Bayesian posterior given the signal is strictly between 0 and 1. □

Proof of Proposition 2. Recall that $\mathbf{A}$ is the adjacency matrix of the network. Moreover, $\mathbf{b}_t$ is the vector of the agents' estimates at period t, with its i-th row being agent i's estimates $\mathbf{b}_t^i$. From the text, $(\mathbf{A}^k)_{il}$, the il-th entry of matrix $\mathbf{A}^k$, measures how many walks of length k the initial signal can travel from i to l.

We now show by induction that the number of copies of x_0^i agent l learns at time t is the il-th entry of the matrix $\sum_{k=0}^{t-1} \mathbf{A}^k$. This is clearly true when $t = 1, 2$. At $t = 1$, only agent i has learned the signal: $(\mathbf{A}^0)_{ii} = 1$, $(\mathbf{A}^0)_{ih} = 0$ for all $h \neq i$. At $t = 2$, each neighbor of agent i extracts one copy of x_0^i. Now suppose the claim holds at time $t \geq 2$. At the beginning of time $t + 1$, the total number of copies agent l newly extracts from all her neighbors in period $t + 1$ is the sum of copies of signals l's neighbors have learned in time t: $\sum_{h \in \mathrm{N}_l} (\mathbf{A}^{t-1})_{ih}$, which is exactly $(\mathbf{A}^t)_{il}$.

Because agents learn from any change in a neighbor's report, there is a directed circle of length 2 among any pair of agents. If there is at least one circle of an odd length, then the common divisor of these circles is 1, and thus the adjacency matrix $\mathbf{A}$ is aperiodic. It is also irreducible because the network is connected by assumption. By the Perron-Frobenius theorem, $\mathbf{A}$ has both a positive right and a positive left eigenvector, $\mathbf{v}$ and $\mathbf{w}$, corresponding to $\lambda > 1$, the largest eigenvalue of $\mathbf{A}$. We can normalize it so that $\mathbf{w}^T\mathbf{v} = 1$, where $\mathbf{w}^T$ is the transpose of $\mathbf{w}$. Moreover, $\lim_{k \to \infty} \mathbf{A}^k/\lambda^k = \mathbf{v}\mathbf{w}^T$.

Recall that $\mathbf{b}_t = \left(\sum_{k=0}^{t-1} \mathbf{A}^k\right)\mathbf{b}_1$. First, observe that as $k \to \infty$, $\mathbf{A}^k\mathbf{b}_1 = \lambda^k(\mathbf{v}\mathbf{w}^T)\mathbf{b}_1$. That is, the new information agent i learns in period $t = k$ is a weighted average of agents' initial estimates, where the weight agent i gives to j's signal is $\lambda^k v_i w_j$. In the case of binary states, as $k \to \infty$, if $\mathbf{w}^T\mathbf{b}_1 > 0$, all agents believe the new information at $t = k$ suggests state s_1 is more likely than s_2, and vice versa. Next, agents' estimates in period t is the sum of all the information they have learned in the past. Take M sufficiently large that A^M is sufficiently close to $\mathbf{A}^\infty$. In period $t > M$, $\mathbf{b}_t$ is sufficiently close to $\left(\sum_{k=0}^{M} \mathbf{A}^k + \sum_{k=M+1}^{t} \lambda^k(\mathbf{v}\mathbf{w}^T)\right)\mathbf{b}_1$. The second term dominates the first term as $t \to \infty$, and thus if $\mathbf{w}^T\mathbf{b}_1 > (<)0$, all agents believe in s_1 (s_2). Their beliefs are indeterminate if $\mathbf{w}^T\mathbf{b}_1 = 0$, which is ruled out by assumption. □

Proof of Proposition 3. We show by construction that agents' learning outcomes are wrong with a positive probability even with an infinite number of informative signals. Let the true state be $s = s^*$. There exists a possible signal $x^{i,m}$ belonging to some agent i such that for some other state $s' \neq s^*$, $s' \in \arg\max_{s_n} \Pr(s_n \mid x^{i,m})$. We focus on the case that s' is the unique mostly likely state given $x^{i,m}$ (the case of multiple most likely states is analogous). Denote signal $x^{i,m}$ as x'. Clearly, $\Pr(s' \mid x') > \Pr(s^* \mid x')$. Consider the following sequence of finite signals: first, nature sends x' to agent i at $t = 0$ and $t = 1$. Next, no informative signals reach any agents for the next $D + \tau$ periods, that is, from $t = 2$ to $t = t^* \equiv 2 + D + \tau$. This interval first allows both signals to travel to every other agent, which takes D periods, and then it allows agents to repeatedly learn copies of these signals in the next τ periods. Two steps to determine τ. Recall that the set of all possible signals that agents can receive from nature is $X = \cup_i X^i$. In the first step, we identify the integer k' such that

$$\frac{\Pr(s' \mid k' \text{ copies of } x')}{\Pr(s^* \mid k' \text{ copies of } x')} \geq \frac{\Pr(s^* \mid x^*)}{\Pr(s' \mid x^*)}, \text{ where } x^* \equiv \arg\max_{x \in X} \frac{\Pr(s^* \mid x)}{\Pr(s' \mid x)}. \tag{13}$$

To avoid carrying this likelihood ratio for the rest of the proof, for any signal x (or set of signals), we denote it as the following:

$$b(s', s^* \mid x) = \log \Pr(s' \mid x) - \log \Pr(s^* \mid x).$$

In the second step, we require that in each period from period $t = D + \tau$, the repetition must be strong enough such that every agent l extracts at least $2k'$ copies of x' (excluding other later exogenous signals). That is, for all $t \geq D + \tau$, $(\mathbf{A}^{t-1} + \mathbf{A}^{t-2})_{il} \geq 2k'$.

Next, we claim that regardless of the signals agents receive from nature after period t^*, all agents believe s' is increasingly more likely than s^* over time. That is, $\lim_{t\to\infty} b_t^h(s') - b_t^h(s^*) = \infty$ for all $h \in g$. This is because starting from period $t = t^* + 1$, agent l extracts at least $4k'$ copies of x' from any neighbor ($2k'$ copies are from herself to the neighbor and back to her, and the other $2k'$ copies are from a neighbor of the neighbor), plus a signal from nature. The new information to him is

$$\left(b_{t+1}^l(s') - b_{t+1}^l(s^*)\right) - \left(b_t^l(s') - b_t^l(s^*)\right) \geq b\left(s', s^* \mid 3k' \text{ copies of } x'\right). \tag{14}$$

That is, the signal l extracts should favor s' over s^* at least as much as $3k'$ copies of x' since period t^*. This shows that the initial condition that each agent learns at least $2k'$ copies of x' persists regardless of the exogenous signals reaching the network after period t^*. Therefore the process described above lasts: in each period every agent believes more strongly that state s' is more likely than s^*, and thus all agents believe s^* is not the true state with probability arbitrarily close to 1 as $t \to \infty$.

Lastly, for any state $\tilde{s} \neq s'$, we can repeat the same process above replacing s^* with $\tilde{s}$. As a result, we can show that all agents believe in s' with probability arbitrarily close to 1 as $t \to \infty$. Because both the initial sequence of signal x' and the number of periods up to t^* are finite, and we do not restrict the signals starting from period $t^* + 1$, agents believe in the wrong state with a positive probability. □

Proof of Proposition 4. We prove this proposition via the following two lemmas. First, as long as $c_i \geq \hat{c}_i$, agent i's $\hat{c}_i$-agent estimates $\mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i-1},c_i}$ are independent of c_i.

Lemma A1. *When agent i's cognitive ability is above her cutoff, for each clique of agents, her $\hat{c}_i$-agent estimate remains the same regardless of her c_i. That is, $\mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i-1},c_i} = \mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i-1},\hat{c}_i}$ for all cliques $\{i, i_1, \ldots, i_{\hat{c}_i-1}\}$, $c_i \geq \hat{c}_i$, $\mathbf{c}_{-i}$, t, and X_T.*

Proof of Lemma A1. Because the multi-agent estimates are formed iteratively, in order to show $\mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i-1},c_i} = \mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i-1},\hat{c}_i}$, we need to show her $(\hat{c}_i + 1)$-agent estimates are the same. Recall that when her cognitive ability is equal to $\hat{c}_i$, she uses the degenerate estimate $\mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i},\hat{c}_i} = \mathbf{b}_t^{ii_2\ldots i_{\hat{c}_i},\hat{c}_i}$ by equation (5). When her cognitive ability is above the cutoff ($c_i > \hat{c}_i$), then $\mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i},c_i}$ is formed by the Bayes-like formula (4). Despite this difference, we claim that if $c_i > \hat{c}_i$,

$$\mathbf{b}_t^{ii_1\ldots i_{\hat{c}_i},c_i} = \mathbf{b}_t^{ii_2\ldots i_{\hat{c}_i},c_i}. \tag{15}$$

This is true by the second part of Proposition A1 above. In other words, the $(\hat{c}_i + 1)$-agent estimate formed by formula (4) is the same as that set degenerately by equation (5). Thus, agent i's $\hat{c}_i$-agent estimate remains the same for all $c_i \geq \hat{c}_i$. □

Since the cognitively constrained learning rule is iterative, by Lemma A1, formula (1) and (3), agent i learns the same new information from her neighbors for all $c_i \geq \hat{c}_i$. Then by (2) and (4), her estimates are the same for all $c_i \geq \hat{c}_i$. This proves part 1 of Proposition 4.

Second, we show that if $c_i < \hat{c}_i$, agent i's c_i-agent estimate under c_i could differ from that under cognitive ability $\hat{c}_i$.

Lemma A2. *Suppose that $\mathbf{c}_{-i} \geq \hat{\mathbf{c}}_{-i}$. For any $c_i < \hat{c}_i$, there exist some t, X_T, and clique $\{i, i_1, \ldots, i_{c_i}\}$, such that agent i's $(c_i + 1)$-agent (possibly degenerate) estimate is affected by her cognitive ability: $\mathbf{b}_t^{ii_1\ldots i_{c_i},c_i} \neq \mathbf{b}_t^{ii_1\ldots i_{c_i},\hat{c}_i}$.*

Proof of Lemma A2. First, suppose $c_i \geq 2$ and $c_i < \hat{c}_i$. By definition, there exists some $(\hat{c}_i - 1)$-agent clique within g_i, say $\{i, i_1, \ldots, i_{\hat{c}_i-2}\}$, whose shared local network is not a clique itself. Then there exist at least two agents h_1 and h_2 who are both connected to all the agents in $\{i, i_1, \ldots, i_{\hat{c}_i-2}\}$, but h_1 and h_2 are not connected themselves. Let agent h_1 receive the only informative signal, which is his initial signal. Because agent i cannot form $(c_i + 1)$-agent estimates, she uses degenerate estimate $\mathbf{b}_t^{ih_2i_1\ldots i_{c_i-1},c_i} = \mathbf{b}_t^{ii_1\ldots i_{c_i-1},c_i}$ by equation (5). Notice that the latter estimate uses reports observed by all the agents in $\{i, i_1, \ldots, i_{c_i-1}\}$, including reports from both h_1 and h_2. Thus, agent i expects each of them to believe that other agents should believe there is only one informative signal, the one from agent h_1. Because of the degenerate estimate, she also expects h_2 to believe so: $\mathbf{b}_3^{ih_2i_1\ldots i_{c_i-1},c_i} = \mathbf{b}_1^{h_1}$. However, this expectation of agent i is incorrect because agent h_2 does not know agent h_1 and he only observes that each agent in $\{i, i_1, \dots, i_{\hat{c}_i-2}\}$ reports a signal, a total of $(\hat{c}_i - 1)$ signals. When agent i is sufficiently sophisticated with a cognitive ability of $\hat{c}_i$ or above, agent i would be able to set the correct expectation about agent h_2 according to Lemma A1. That is, if her cognitive ability is at the cutoff, agent i believes $\mathbf{b}_3^{ih_2i_1\dots i_{c_i-1},\hat{c}_i} = \eta\mathbf{b}_1^{h_1}$ for some integer $\eta \geq (\hat{c}_i - 1)$. Thus, $\mathbf{b}_3^{ih_2i_1\dots i_{c_i-1},c_i} \neq \mathbf{b}_3^{ih_2i_1\dots i_{c_i-1},\hat{c}_i}$. Second, if $c_i = 1$, then by definition, agent i's 2-agent estimate is always $\mathbf{0}$. It differs from her 2-agent estimate whenever $c_i \geq 2$ and there is an informative signal. □

Lemma A2 shows that if $c_i < \hat{c}_i$, then her $(c_i + 1)$-agent estimate differs from that she would have formed when her cognitive ability is above the cutoff level. If this difference occurs at period t, then her c_i-agent estimates in period $t+1$ differ under c_i and $\hat{c}_i$ by formula (3) and (4). This continues iteratively until in period $t + c_i$, her estimates differ under c_i and $\hat{c}_i$. This proves part 2 of Proposition 4. □

Derivation of estimates in Example 3. Recall from the text that each agent's estimate in each period depends only on how many copies of x^1 she has extracted so far. For instance, $b_t^i(s_1) = \eta_t^i\varphi$, $b_t^{ij}(s_1) = \eta_t^{ij}\varphi$ and similarly for their multi-agent estimates. Then, we claim that when $c_i = 2$ for all agents, their estimates evolve by the following expressions.

$$\eta_t^1 = 2\eta_{t-1}^2 - \eta_{t-1}^1, \text{ and } \eta_t^2 = 2\eta_{t-1}^1 - \eta_{t-2}^2. \tag{16}$$

To begin with, we can rewrite the formula (2) such that:

$$\eta_t^i = \eta_{t-1}^i + \sum_{j\in \mathrm{N}_i}(\eta_{t-1}^j - \eta_{t-1}^{ij}).$$

Similarly, rewrite expression (4) for 2-agent estimates, and we have

$$\eta_t^{ij} = \eta_{t-1}^{ij} + \sum_{h\in g_{ij}}(\eta_{t-1}^h - \eta_{t-1}^{ijh}).$$

Recall agent 1's cutoff is $\hat{c}_1 = 2$, by Lemma A1, $\eta_t^{121} = \eta_t^{12}$ and $\eta_t^{124} = \eta_t^{14}$. Since $\eta_t^{12} = \eta_t^{14}$ and $\eta_t^2 = \eta_t^4$, we have

$$\eta_t^1 = \eta_{t-1}^1 + 2\eta_{t-1}^2 - 2\eta_{t-1}^{12}.$$

Substitute into how agent 1 extracts signals, it is easy to see that $\eta_{t-1}^{12} = \eta_{t-1}^1$. Intuitively, agent 1's cutoff level of cognitive ability is 2 and he expects agent 2 and 4 to agree with him in every period because he believes (incorrectly) that they observe the same reports. The expression for agent 1's estimates in (16) follows immediately.

For agent 2, note that $\eta_{t-1}^{214} = \eta_{t-1}^{24}$ by her degenerate 3-agent estimate, and $\eta_{t-1}^{24} = \eta_{t-1}^4$ because 2 and 4 are symmetric and never learn from each other. Agent 2 also knows that agent 1 does not know 3. Therefore

$$\eta_t^{21} = \eta_{t-1}^1 + (\eta_{t-1}^2 - \eta_{t-1}^{21}).$$

Put it together with

$$\eta_t^2 = \eta_{t-1}^2 + 2(\eta_{t-1}^1 - \eta_{t-1}^{21})$$

and solve, we have:

$$2\eta_t^{21} = 2\eta_{t-1}^1 - \eta_{t-1}^2 + \eta_t^2.$$

Substitute into agent 2's estimate, we have the expression in (16) for agent 2. □

Proof of Proposition 5. Recall that ϕ_{mn}^i is the probability agent i observes signal $x^{i,m}$ when the true state is s_n, and $\phi_\emptyset^i$ is the probability she observes the uninformative signal. Then, decisive evidence requires that each $\phi_{mn}^i \in \{0, 1 - \phi_\emptyset^i\}$. That is, the signal $x^{i,m}$ can either completely rule out state s_n (when $\phi_{mn}^i = 0$) or it is the only informative signal when s_n is the true state (but it can be positive for multiple states). For simplicity, we consider initial signals only. Let s_1 be the true state, and let $\mathcal{P}^i(s_n) = \{s_{n'} : \phi_{mn}^i = \phi_{mn'}^i = 1 - \phi_\emptyset^i\}$, which is the subset of states whose signal is the same as state s_n's. Because agents believe some states are true with zero probability in this partitional model, let $\log 0 = -\infty$. Since agents can use any state they believe to be the true state with a positive probability to normalize, we let all agents normalize their estimates by reporting the log-likelihood ratios over s_1. For simplicity, we focus on the case that all the initial signals are informative. If agent i's initial signal is uninformative, we can replace $\mathcal{P}^i(s_1)$ with the trivial partition S, the entire state space, and the following proof goes through.

Note that at $t = 1$, for agent i, $b_1^i(s_n) = 0$ if $s_n \in \mathcal{P}^i(s_1)$, and $-\infty$ otherwise. Also, by assumption, $\mathbf{b}_1^{ij} = \mathbf{0}$ for all j. At $t = 2$, by formula (1) of the learning rule, $\boldsymbol{\alpha}_1^{ij} = \mathbf{b}_1^j - \mathbf{b}_1^{ij}$. Clearly, for all $s_n \in \mathcal{P}^j(s_1)$, $\alpha_1^{ij}(s_n) = 0$, and $-\infty$ otherwise. Thus, what agent i learns from agent j is the same as j's period-1 estimate: $\boldsymbol{\alpha}_1^{ij} = \boldsymbol{\beta}_1^j$. Let the intersection of the partitional elements containing the true state s_1 of all agents in g_i be $\mathcal{P}_1^{g_i}(s_1) \equiv \cap\{\mathcal{P}^h(s_1)\}_{h\in g_i}$. Then by formula (2), for all cognitive ability $c_i \geq 1$, agent i's estimates are

$$b_2^i(s_n) = b_1^i(s_n) + \sum_{h\in g_i} \alpha_1^{ih}(s_n) = \begin{cases} 0 & \text{if } s_n \in \mathcal{P}_1^{g_i}(s_1); \\ -\infty & \text{otherwise.} \end{cases}$$

The multi-agent estimates at $t \geq 2$ differ by agents' cognitive abilities, so we consider the two separate cases: $c_i \geq 2$ and $c_i = 1$. For all $c_i \geq 2$, let $\mathcal{P}_1^{g_{ij}}(s_1) \equiv \cap\{\mathcal{P}^h(s_1)\}_{h\in g_{ij}}$. By expression (4), agent i's estimate of agent j's estimate is $b_2^{ij}(s_n) = 0$ for $s_n \in \mathcal{P}_1^{g_{ij}}(s_1)$, and $-\infty$ otherwise. At $t = 3$, when $\boldsymbol{\beta}_2^j \neq \boldsymbol{\beta}_2^{ij}$, agent i extracts a new signal from agent j. We let $\alpha_2^{ij}(s_n) = 0$ if $b_2^{ij}(s_n) = b_2^j(s_n) = -\infty$ because agent i learns nothing about s_n from agent j. Note from the above definitions that $\mathcal{P}_1^{g_{ij}}(s_1) \supseteq \mathcal{P}_1^{g_j}(s_1)$. Thus if $s_n \notin \mathcal{P}_1^{g_{ij}}(s_1)$ then $s_n \notin \mathcal{P}_1^{g_j}(s_1)$. It then follows that if $b_2^{ij}(s_n) = -\infty$, then $b_2^j(s_n) = -\infty$. Then agent i only extracts new information if $b_2^j(s_n) = -\infty$ but $b_2^{ij}(s_n) = 0$ for some s_n. This implies that $\alpha_2^{ij}(s_n) = -\infty$ if $s_n \in \mathcal{P}_1^{g_{ij}}(s_1) \setminus \mathcal{P}_1^{g_j}(s_1)$, and $\alpha_2^{ij}(s_n) = 0$ otherwise. If $c_i = 1$ instead, $\mathbf{b}_2^{ij} = \mathbf{b}_1^j$, and then $\alpha_2^{ij}(s_n) = -\infty$ if $s_n \in \mathcal{P}^j(s_1) \setminus \mathcal{P}_1^{g_j}(s_1)$, and $\alpha_2^{ij}(s_n) = 0$ otherwise. Let $\mathcal{P}_2^{g_i}(s_1)$ be the set of states agent i thinks are still possible given her extracted signals at $t = 3$. Regardless of the cognitive abilities $c_i \geq 1$, the set $\mathcal{P}_2^{g_i}(s_1)$ remains the same and $\mathcal{P}_2^{g_i}(s_1) \subseteq \mathcal{P}_1^{g_i}(s_1)$. Formally,

$$\mathcal{P}_2^{g_i}(s_1) \equiv \cap\{\mathcal{P}_1^{g_h}(s_1)\}_{h\in g_i} = \cap\{\mathcal{P}^l(s_1)\}_{l\in g_i^2},$$

in which $g_i^d = \{l : d(il) \leq d\}$ where $d(il)$ is the *distance*—the length of the shortest path—between agent i and l. Also let D be the *diameter*, the longest distance between any two agents. Intuitively, $\mathcal{P}_2^{g_i}(s_1)$ is the intersection of the partitional elements containing s_1 of all i's neighbors and their neighbors. Therefore for all $c_i \geq 1$, $b_3^i(s_n) = 0$ if $s_n \in \mathcal{P}_2^{g_i}(s_1)$, and $-\infty$ otherwise.

Iteratively, we can show that for $t \geq 2$, $b_t^i(s_n) = 0$ if $s_n \in \mathcal{P}_{t-1}^{g_i}(s_1) \equiv \cap\{\mathcal{P}^h(s_1)\}_{h\in g_i^{t-1}}$, and $-\infty$ otherwise. By at most period $t = D + 1$, all the initial signals have reached agent i through her neighbors according to the travel path of the signals. Let $\mathcal{P}^g(s_1) \equiv \cap\{\mathcal{P}^l(s_1)\}_{l\in g}$ be the intersection of all agents' partitional elements containing state s_1. Then agent i's estimate is simply 0 if $s_n \in \mathcal{P}^g(s_1)$, and $-\infty$ otherwise. This holds for all networks and all cognitive levels, and thus the agents' learning outcomes are correct. □

References

Acemoglu, Daron, Dahleh, Munther A., Lobel, Ilan, Ozdaglar, Asuman, 2011. Bayesian learning in social networks. Rev. Econ. Stud. 78 (4), 1201–1236.
Alevy, Jonathan E., Haigh, Michael S., List, John A., 2007. Information cascades: evidence from a field experiment with financial market professionals. J. Finance 62, 151–180.
Anderson, Lisa R., Holt, Charles A., 1997. Information cascades in the laboratory. Am. Econ. Rev. 87, 847–862.
Aumann, Robert J., 1976. Agreeing to disagree. Ann. Stat. 4 (6), 1236–1239.
Bai, Jie, Golosov, Mikhail, Qian, Nancy, Kai, Yan, 2015. Understanding the Influence of Government Controlled Media: Evidence from Air Pollution in China. Working paper.
Bala, Venkatesh, Goyal, Sanjeev, 1998. Learning from neighbours. Rev. Econ. Stud. 65 (3), 595–621.
Banerjee, Abhijit V., 1992. A simple model of herd behavior. Q. J. Econ. 107, 797–819.
Bikhchandani, Sushil, Hirshleifer, David, Welch, Ivo, 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. J. Polit. Econ. 10 (1), 992–1026.
Cai, Hongbin, Chen, Yuyu, Fang, Hanming, 2009. Observational learning: evidence from a randomized natural field experiment. Am. Econ. Rev. 99 (3), 864–882.
Canen, Nathan, Schwartz, Jacob, Song, Kyungchul, 2020. Estimating local interactions among many agents who observe their neighbors. Quant. Econ. 11 (3), 917–956.
Celen, Bogachan, Kariv, Shachar, 2004. Distinguishing informational cascades from herd behavior in the laboratory. Am. Econ. Rev. 94, 484–498.
Chandrasekhar, Arun G., Larreguy, Horacio, Xandri, Juan Pablo, 2020. Testing models of social learning on networks: evidence from two experiments. Econometrica 88 (1), 1–32.
Charness, Gary, Levin, Dan, 2009. The origin of the winner's curse: a laboratory study. Am. Econ. J. Microecon. 1 (1), 207–236.
Coleman, James Samuel, 1961. The Adolescent Society: The Social Life of the Teenager and Its Impact on Education. Free Press of Glencoe.
Costa-Gomes, Miguel, Crawford, Vincent P., Broseta, Bruno, 2001. Cognition and behavior in normal form games: an experimental study. Econometrica 69 (5), 1193–1235.
Costa-Gomes, Miguel A., Crawford, Vincent P., 2006. Cognition and behavior in two-person guessing games: an experimental study. Am. Econ. Rev. 96 (5), 1737–1768.
DeGroot, Morris H., 1974. Reaching a consensus. J. Am. Stat. Assoc. 69 (345), 118–121.
DeMarzo, Peter M., Vayanos, Dimitri, Zwiebel, Jeffrey, 2003. Persuasion bias, social influence, and uni-dimensional opinions. Q. J. Econ. 118 (3), 909–968.
Ellison, Glenn, Fudenberg, Drew, 1993. Rules of thumb for social learning. J. Polit. Econ. 101 (4), 612–643.
Ellison, Glenn, Fudenberg, Drew, 1995. Word-of-mouth communication and social learning. Q. J. Econ. 110 (1), 93–125.
Eyster, Erik, Rabin, Matthew, 2010. Naive herding in rich-information settings. Am. Econ. J. Microecon. 2 (4), 221–243.
Feld, Scott L., 1991. Why your friends have more friends than you do. Am. J. Sociol. 96 (6), 1464–1477.
Gale, Douglas, Kariv, Shachar, 2003. Bayesian learning in social networks. Games Econ. Behav. 45 (2), 329–346.
Geanakoplos, John D., Polemarchakis, Heraklis M., 1982. We can't disagree forever. J. Econ. Theory 28 (1), 192–200.
Golub, Benjamin, Jackson, Matthew O., 2010. Naive learning in social networks and the wisdom of crowds. Am. Econ. J. Microecon. 2 (1), 112–149.
Golub, Benjamin, Sadler, Evan, 2017. Learning in social networks. Working paper, SSRN 2919146.
Grimm, Veronika, Mengel, Friederike, 2020. Experiments on belief formation in networks. J. Eur. Econ. Assoc. 18 (1), 49–82.
Hellman, E. Martin, Cover, Thomas M., 1970. Learning with finite memory. Ann. Math. Stat. 41 (3), 765–782.

Ho, Teck-Hua, Camerer, Colin F., Weigelt, Keith, 1998. Iterated dominance and iterated best response in experimental 'p-beauty contests'. Am. Econ. Rev. 88 (4), 947–969.
Jackson, Matthew O., Rogers, Brian W., 2005. The economics of small worlds. J. Eur. Econ. Assoc. 3 (2–3), 617–627.
Jadbabaie, Ali, Molavi, Pooya, Sandroni, Alvaro, Tahbaz-Salehi, Alireza, 2012. Non-Bayesian social learning. Games Econ. Behav. 76 (1), 210–225.
Kneeland, Terri, 2015. Identifying higher-order rationality. Econometrica 83 (2065–2079).
Kubler, Dorothea, Weizsacker, Georg, 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. Rev. Econ. Stud. 71, 425–441.
Levy, Gilat, Razin, Ronny, 2018. Information diffusion in networks with the Bayesian Peer Influence heuristic. Games Econ. Behav. 109, 262–270.
Li, Wei, Tan, Xu, 2020. Locally Bayesian learning in networks. Theor. Econ. 15, 239–278.
Mobius, Markus, Phan, Tuan, Szeidl, Adam, 2015. Treasure Hunt: Social Learning in the Field. NBER Working Paper No. 21014.
Molavi, Pooya, Tahbaz-Salehi, Alireza, Jadbabaie, Ali, 2018. A theory of non-Bayesian social learning. Econometrica 86 (2), 445–490.
Mossel, Elchanan, Sly, Allan, Tamuz, Omer, 2015. Strategic learning and the topology of social networks. Econometrica 83 (5), 1755–1794.
Mueller-Frank, Manuel, 2013. A general framework for rational learning in social networks. Theor. Econ. 8 (1), 1–40.
Mueller-Frank, Manuel, 2014. Does one Bayesian make a difference? J. Econ. Theory 154, 423–452.
Mueller-Frank, Manuel, Neri, Claudia, 2021. A general analysis of boundedly rational learning in social networks. Theor. Econ. 16 (1), 317–357.
Nagel, Rosemarie, 1995. Unraveling in guessing games: an experimental study. Am. Econ. Rev. 85 (5), 1313–1326.
Parikh, Rohit, Krasucki, Paul, 1990. Communication, consensus, and knowledge. J. Econ. Theory 52 (1), 178–189.
Penczynski, Stefan P., 2017. The nature of social learning: experimental evidence. Eur. Econ. Rev. 94, 148–165.
Piccione, Michele, Rubinstein, Ariel, 1997. On the interpretation of decision problems with imperfect recall. Games Econ. Behav. 20 (1), 3–24.
Roberts, Ben, Kroese, Dirk P., 2007. Estimating the number of s-t paths in a graph. J. Graph Algorithms Appl. 11 (1), 195–214.
Shields, Timothy, Xin, Baohua, 2012. Higher-order Beliefs in Simple Trading Models. ESI Working Paper 12-18.
Valiant, Leslie G., 1979. The complexity of enumeration and reliability problems. SIAM J. Comput. 8 (3), 410–421.
Wilson, Andrea, 2014. Bounded memory and biases in information processing. Econometrica 82 (2), 2257–2294.