

Locally Bayesian Learning in Networks*

Wei Li

University of British Columbia

Xu Tan

University of Washington

November 1, 2018

Abstract

Agents in a network want to learn the true state of the world from their own signals and their neighbors' reports. They only know their local networks, consisting of their neighbors and the links among them. Every agent is Bayesian with the (possibly misspecified) prior belief that her local network is the entire network. We present a tractable procedure to implement such an agent's learning: she extracts new information using the full history of observed reports in her local network. Despite their limited network knowledge, agents learn correctly when the network is a social quilt, a tree-like union of cliques. But they fail to learn when a network contains interlinked circles (echo chambers) despite an arbitrarily large number of correct signals.

JEL: D03, D83, D85

Keywords: locally Bayesian learning, rational learning with misspecified priors, iterated learning procedure

*We thank Andrea Blume, Jeff Ely, Ben Golub, Vitor Farinha Luz, Drew Fudenberg, Fahad Khalil, Jacques Lawarrée, Stephen Morris, Marit Rehavi, Tanya Rosenblat, Edward Schlee, Joel Sobel, Wing Suen, Joel Watson, Quan Wen, the editor and three referees for helpful comments. We are especially grateful for extensive discussions and numerous comments from Matt Jackson, Li, Hao and Mike Peters. We thank Anubhav Jha for excellent research assistance. We also thank the seminar participants at UBC, UW, 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 for financial support. Email: wei.li@ubc.ca; tanxu@uw.edu.

1 Introduction

People often learn from those they interact with, who in turn talk to and learn from their neighbors. In order to make correct decisions, they need to account for potential repetitions in the information learned from their social networks. Failure to do so can lead to learning errors with real world consequences such as political polarization, entrenched poverty, and disease outbreaks. For instance, in Minnesota’s close-knit Somali community, MMR vaccination rates among children dropped from 92% in 2004 to 43% in 2013. If a new mother in this community hears from her neighbors that MMR causes autism, she may decide not to vaccinate her baby. Her neighbors may have heard this news from their neighbors. Thus one piece of fake news such as a fraudulent research paper linking MMR to autism, fully retracted in 2010, may reach her through multiple neighbors creating the illusion of multiple sources. As a consequence, she believes erroneously—and increasingly if the same information travels back to her again—that MMR is dangerous even though there is a large amount of correct information to the contrary.¹

Motivated by this phenomenon, we propose a novel model of locally Bayesian learning. It is *Bayesian* in that each agent updates her beliefs rationally using all the observed reports from her neighbors. In particular, she tracks the changes in each neighbor’s reports over time, and attributes any *unexpected* change to new information. It is *local* in that each agent only knows and discerns new information within her local network, consisting of her neighbors and the links among them.² We show that, surprisingly, despite limited knowledge of the network, agents are capable of partialing out repeated information and forming correct beliefs in certain networks. To characterize these networks, we identify an intuitive and clean relationship between two specific features of a finite network and the agents’ learning outcomes. Agents learn correctly if a network has these two features; otherwise, they cannot avoid learning errors even if they receive an arbitrarily large number of correct signals. These results complement the large literature on network learning which focuses on when the Law

¹The Minnesota Department of Public Health has had very limited success in changing these beliefs, even as they encountered the largest and growing measles outbreak in two decades. For more information, see Howard, Jacqueline. 2017. “Anti-vaccine groups blamed in Minnesota measles outbreak.” CNN, May 8. In the result sections, we will show why the retraction of fake news and the announcements from the public health officials may not overturn such erroneous beliefs.

²The limited and local network knowledge is consistent with evidence from surveys. For instance, [Krackhardt \(1990\)](#) finds that the accuracy of knowing other people’s connections is 15% - 48% in a small startup of 36 people, and [Casciaro \(1998\)](#) finds the accuracy is around 45% in a research center of 25 people. Moreover, [Breza, Chandrasekhar, and Tahbaz-Salehi \(2016\)](#) find that each agent’s knowledge about the network is highly localized, declining steeply with the pair’s network distance from the agent.

of Large Numbers holds and agents can learn asymptotically in large networks.³ Moreover, because the (finite) network features we identify are observable, our model is potentially testable in the lab or in the field even with small datasets.

We focus on an important feature of Bayesian learning that has been under-studied in the literature. That is, agents have perfect memory—the full history of observed reports—and use it to update their beliefs by Bayes’ rule. To make the model tractable, we make the following behavioral assumption: each agent believes her local network is the entire network (and such belief is common knowledge). Formally, our model studies the learning outcomes of Bayesian agents who focus entirely on their local networks due to these (possibly misspecified) priors of the network. This assumption reflects the heavy cognitive and computational burden agents face if they were to properly update their beliefs about the entire network. In our view, modeling perfect memory is a necessary step toward modeling how agents try to avoid using repeated and distorted information from the network, a topic under increasing scrutiny in recent years.

Our first contribution is to identify an iterative procedure that implements locally Bayesian learning. Specifically, suppose there are finitely many states, and agents want to learn the true state, such as whether MMR causes autism in our opening example. Each agent learns by forming and updating her beliefs about the state distribution, such as the probability that MMR truly leads to autism. In period 1, each agent first forms her beliefs using her initial signal. She then simultaneously reports her beliefs to each neighbor and hears their reports. Next, she gets another signal from nature and period 1 ends. In each ensuing period, each agent first infers any new information contained in her neighbors’ most recent reports. The main innovation of our procedure is that each agent remembers existing information by forming her (interim) second-order beliefs—her beliefs about each neighbor’s beliefs conditional on the reports that they both have seen. She then compares these (interim) second-order beliefs with a neighbor’s actual report and attributes any difference to a “new” signal. Under the behavioral assumption, she thinks this new signal must be an independent signal from nature. She then incorporates all the newly inferred signals and updates her beliefs using Bayes’ rule. Next, she exchanges reports with her neighbors and receives a signal from nature. The learning procedure continues to the next period.

Our second contribution is to show if a network has the following two features, the agents learn correctly even though they only know their local networks and have possibly misspecified beliefs about the entire network. The first feature is that the network contains *no simple*

³It often requires that each agent has a negligible influence on the limit beliefs of the network. See [Golub and Jackson \(2010\)](#) and [Mossel, Sly, and Tamuz \(2015\)](#) among others.

circles. A simple circle is a path-connected subnetwork that contains at least four agents, each of whom has exactly two links to other agents in the subnetwork. The second feature is that the network satisfies *local connection symmetry*. That is, if any pair of agents have common neighbors, then these neighbors must be connected themselves. We show that the only type of networks with both features is a social quilt, a tree-like union of cliques. First, because every pair of agents in a clique is connected, it satisfies local connection symmetry. When information arrives at one member of the clique, all the other members identify it as new. More importantly, everyone believes that all the other members in the clique have learned this information from the same agent. Thus they avoid overcounting the information. Second, the overall tree structure of a social quilt ensures that no simple circle exists. Each signal travels through the network once and only once, and thus the agents learn all signals without repetition. In this way, all the information is aggregated efficiently.

If a network contains simple circles or if it fails local connection symmetry, we show that agents must make learning errors. First, in networks with simple circles, agents overcount signals due to these circles, despite their ability to remember old information locally. For example, consider a simple circle with four agents 1, 2, 3 and 4. Agent 1 receives the only signal, which reaches agent 2 and 4 next. Agent 3, not knowing the existence of agent 1, thinks that there are two copies of the signal. The signal then travels along the circle in both directions, clockwise and counterclockwise. In this way, in every four periods, each agent believes there are two new copies of this signal. The problem is exacerbated in networks with multiple simple circles, or *echo chambers*. Because each signal travels both within each simple circle and between each pair of simple circles, the signal repetition grows at an exponential rate.⁴ As a result, the Law of Large Numbers may fail: everyone believes in a wrong initial signal despite an arbitrarily large number of correct signals.

Second, we identify the error of opinion swings if the network fails local connection symmetry. This more novel type of learning error could lead to the instability in poll results and opinion surveys without any exogenous changes in information. To see the intuition, add a link between agent 2 and 4 to the simple circle (1234) above, in which agent 1 and 3 remains unconnected. Recall that agent 1 receives the only informative signal. Then, agent 3 starts off believing in two copies of this signal as in the simple circle. But agent 2 and 4 only believe in one copy because they are connected and know they have inferred the same signal

⁴For instance, add another simple circle (1567) to the previous one so that agent 1 connects both simple circles. Then, after each round trip within the simple circle (1234), the two new copies of the signal reach agent 1 and then travel to the simple circle (1567). After another four periods, they get doubled again and come back to (1234) in the form of four copies. Because the number of copies is doubled every four periods (and more because of repetition within simple circles), the repetition increases exponentially.

from agent 1. Agent 3 has to rationalize the fact that 2 and 4 do not learn a second copy from each other. He does so by inferring that agent 2 and 4 each must have received a signal negatively correlated with their original signal. These two negatively correlated signals offset agent 3's earlier two signals. In this way, agent 3 believes in two copies of the initial signal and no informative signal in alternating periods. More generally, any subnetwork that fails local connection symmetry generates both positively and negatively correlated copies of the informative signals endogenously. As a consequence, the agents may exhibit large and divergent opinion swings.

Literature review

It is well-documented that we learn from our social networks.⁵ One strand of the theoretical literature on network learning shows that agents can form the correct Bayesian beliefs (asymptotically) if the network is common knowledge (see [Gale and Kariv \(2003\)](#), [Mueller-Frank \(2013\)](#), [Mossel, Sly, and Tamuz \(2015\)](#), [Golub and Sadler \(2017\)](#), among others), or if the agents can communicate in complex ways. The other strand of the theoretical literature eschews the complexity of Bayesian learning by assuming that agents learn by following reasonable rules of thumb.⁶ In the classic model of [DeGroot \(1974\)](#), agents treat their neighbors' reports in each period as new information and update their opinions by taking a weighted average of these reports. In our model, agents do not employ any mechanical learning rule, instead they are Bayesian when they learn from their neighbors' reports (subject to the behavioral assumption).

Several more closely-related papers consider quasi-Bayesian learning in networks. In [Bala and Goyal \(1998\)](#), each agent updates beliefs about her optimal action rationally based on the outcomes seen in her local network, but she does not infer information from the actions chosen by her neighbors. They focus on the long-run convergence of actions in any network, while we study how particular network structures affect agents' learning outcomes. More recently, several papers feature imperfect memory in the context that is otherwise the same as our model—agents apply Bayes' rule to all the information they believe are independent ([Molavi, Tahbaz-Salehi, and Jadbabaie \(2016\)](#); [Mueller-Frank and Neri \(2017\)](#); [Levy and](#)

⁵For instance, [Conley and Udry \(2001\)](#) show that pineapple farmers in Ghana learn to use fertilizer from neighbors. [Duflo and Saez \(2002\)](#) find employee participation in retirement savings plans is strongly influenced by their peers. [Mobius and Rosenblat \(2001\)](#) study the opposite side—the effect of isolation and the resulting reduced opportunities to learn from social networks—on inner-city neighborhoods in Chicago.

⁶See [DeGroot \(1974\)](#), [Ellison and Fudenberg \(1993, 1995\)](#), [DeMarzo, Vayanos, and Zwiebel \(2003\)](#), [Golub and Jackson \(2010\)](#), [Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi \(2012\)](#), among many others.

Razin (2016); Dasaratha and He (2017)). The underlying assumption of this class of models, as shown by Molavi, Tahbaz-Salehi, and Jadbabaie (2016), is that in each period, each agent treats a neighbor’s most recent report as sufficient statistics of all the information available to that neighbor.⁷ We differ from these models in a new and significant way: our agents have perfect memory and learn from the full history of reports. They extract new information from unexpected changes in their neighbors’ reports, and thus account for correlations of information locally.

Our paper is also related to the literature on behavioral (mis)learning where rational agents may have some misspecified beliefs. In a non-network context, Eyster and Rabin (2005) assume that each agent correctly predicts the distribution of other agents’ actions, but underestimates the degree to which these actions are correlated with other players’ information. In a sequential social learning context,⁸ Eyster and Rabin (2014) illustrate that rational agents would anti-imitate some predecessors to remove repeated information. They then show that, akin to our model, if the directed network satisfies local connection symmetry, then the agents’ learning outcomes are correct. In addition, Bohren (2016) and Bohren and Hauser (2018) allow agents to have incorrect beliefs about primitives such as the signal distribution or others’ preferences. Our model differs from these papers in that first, we study undirected networks with repeated exchanges of information. Therefore, our agents’ beliefs evolve in a more complex manner due to the large set of reports they receive over time. Second, our misspecified beliefs are about the network structure, which implies that locally, each agent is Bayesian in how she processes information from her neighbors.⁹

Many experiments have studied observational learning in networks.¹⁰ Recently, Enke and Zimmermann (2015) and Golosov, Qian, and Kai (2015) show that people often struggle with distinguishing new information from the old ones in their network. Chandrasekhar,

⁷In addition, Mueller-Frank and Neri (2017) motivate their model as each agent treats each period as period 2, that is, she treats each neighbor’ action as if it depends only on that neighbor’ private signal. In Levy and Razin (2016), agents use a Bayesian Peer Influence heuristic, namely, they believe each neighbor’s belief only contains independent information. In Alatas, Banerjee, Chandrasekhar, Hanna, and Olken (2016), agents know more about the network and treat all signals received as independent.

⁸Our model is also related to the classic papers on herding and social learning in which agents take one and only one action sequentially after observing their predecessors’ actions (see Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), Lee (1993), Smith and Sorensen (2000), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Harel, Mossel, Strack, and Tamuz (2014), among others). If the agents report their posterior beliefs instead of actions, all agents’ learning outcomes are correct because a linear chain is a social quilt.

⁹In this sense, our paper is related to Lipnowski and Sadler (2018) who assume that agents only form correct conjectures about their neighbors strategies. Consequently, if the network is a clique, their agents use Nash equilibrium strategies, just as our agents learn correctly.

¹⁰See Anderson and Holt (1997), Celen and Kariv (2004), Alevy, Haigh, and List (2007), Cai, Chen, and Fang (2009) and Mobius, Phan, and Szeidl (2015) among others.

Larreguy, and Xandri (2012) compare Bayesian learning with naive learning when the network is common knowledge. They find that the naive learning model explains 86-88% of the actions taken versus 82% for the Bayesian model. Grimm and Mengel (2014) show that while some subjects seem to be naive learners, others tried to account for old information, for instance by reducing the weight they attach to their neighbors' later reports.

Section 2 sets up our model and Section 3 introduces our learning procedure. Section 4 shows when agents can learn correctly, and Section 5 characterizes and quantifies their learning errors when they cannot. All proofs are in the Appendix.

2 The model

Consider a network (g, G) : $g = \{1, 2, \dots, I\}$ 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.¹¹ The network is undirected, so information flows both ways: $ij \in G$ if and only if $ji \in G$. It is also path-connected so information diffuses to everyone. That is, for any $i, h \in g$, there is a path $(i_0 i_1 \dots i_l)$ such that $i_0 = i$, $i_l = h$ and $i_k i_{k+1} \in G$ for all $k < l$. A subset of agents in g is a *clique* if any pair of agents in this subset is connected.

Let the set of agent i 's neighbors be $N_i = \{j : ij \in G\}$. Agent i 's *local network* consists of herself, all her neighbors, and all the links among them in the original network. We denote her local network as (g_i, G_i) , where $g_i = N_i \cup \{i\}$ and $G_i = \{hj : h, j \in g_i \text{ and } hj \in G\}$. Agent i and her neighbor j 's *shared local network* is the intersection of their local networks, consisting of themselves, their common neighbors, and all the links among them. We denote their shared local network as (g_{ij}, G_{ij}) , where $g_{ij} = g_i \cap g_j$ and $G_{ij} = G_i \cap G_j$. Similarly, the shared local network of any clique $\{i, j, \dots, k\} \subseteq g_i$ consists of themselves, common neighbors to all of them, and all the links among them.¹² We denote this shared local network as $(g_{ij\dots k}, G_{ij\dots k})$, where $g_{ij\dots k} = g_i \cap g_j \cap \dots \cap g_k$, and $G_{ij\dots k} = G_i \cap G_j \cap \dots \cap G_k$. For instance, consider a triangle network: $g = \{1, 2, 3\}$ and $G = \{12, 13, 23\}$. The shared local network of any pair of agents, or that of all three agents, is the triangle: $g_1 = g_{12} = g_{123} = g$ and $G_1 = G_{12} = G_{123} = G$.

Agents in the network want to learn an unknown state. The set of possible states S is finite: $S = \{s_1, \dots, s_N\}$. All the states are *a priori* equally likely: $\Pr(s_n) = 1/N$ for all $s_n \in S$. The true state is realized before learning begins.

¹¹Throughout this paper, the generic agent is agent i ("she"), and her generic neighbor is agent j ("he").

¹²From now on, we use $(ij\dots k)$ to denote a sequence of agents in which the order matters such as those in a path, and $\{i, j, \dots, k\}$ to denote a set of agents whose order does not matter.

Agents receive signals from nature about the state. The support of agent i 's signals is finite: $X^i = \{x^\emptyset, x^{i,1}, \dots, x^{i,M_i}\}$. Agent i receives the uninformative signal x^\emptyset with probability $\phi^i \in (0, 1)$. For informative signals, let $\phi_{mn}^i = \Pr(x^{i,m} | s_n) \in (0, 1)$ be agent i 's conditional probability of receiving signal $x^{i,m}$ if the state is s_n . Time is discrete: $t = 0, 1, 2, \dots$. In each period up to period T_i , agent i observes a realized signal x_t^i according to the information structure above. She does not receive informative signals at or after $T_i \in [1, \infty]$.¹³ It is common knowledge that the signals are independent across agents and time conditional on the state. Agent i only knows her own information structure $(M_i, \phi^i, \{\phi_{mn}^i\}_{m \leq M_i, n \leq N}, T_i)$, that is, the number of her informative signals, the probability of x^\emptyset and that of each informative signal given each state, and the stopping time of informative signals. She does not know any neighbor j 's information structure except that it has full support. That is, $M_j \in [1, \infty)$, $\phi^j \in (0, 1)$, $\phi_{mn}^j \in (0, 1)$ for all $m \leq M_j$ and $n \leq N$, and $T_j \in [1, \infty]$.

Agent i learns about the underlying state based on her own signals and the reports from her neighbors. In each period t , agent i first forms her beliefs about the state distribution. We denote agent i 's beliefs as $\mathbf{p}_t^i = (p_t^i(s_1), \dots, p_t^i(s_N))$, where $p_t^i(s_n)$ is the probability agent i believes the true state is s_n in period t .¹⁴ To simplify notations, we use the log-likelihood ratios of these beliefs and call them agent i 's *estimates* at period t . That is,

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

Agent i reports her estimates to her neighbors, and simultaneously receives their reports.¹⁵ She then receives a signal (x_t^i) from nature, and period t ends. Given this time line, agent i 's estimates \mathbf{b}_t^i is based on the reports and signals she observed prior to period t . We formally describe how these estimates are formed in Section 3.

¹³Notice that $T_i = \infty$ corresponds to the case that agent i can receive an infinite number of signals, and $T_i = 1$ corresponds to the case of an initial signal only. The latter is the focus of many existing models, while we consider a more general setup allowing for the possibility that signals arrive over time.

¹⁴We use boldface letters to denote vectors throughout this paper.

¹⁵We do not model a utility function formally, but each agent's report is consistent with her maximizing the following quadratic utility function. Namely, agent i myopically chooses a report \mathbf{r}_t^i at period t to maximize the following expected utility using her beliefs at period t , $\mathbb{E}_{\mathbf{p}_t^i} \left[-\sum_{s_n} (r_t^i(s_n) - \mathbb{1}_{s_n=s^*})^2 \right]$, where s^* denotes the true state. It is easy to verify that the optimal report must be her beliefs about the state distribution at period t , that is, $\mathbf{r}_t^i = \mathbf{p}_t^i$.

2.1 The behavioral assumption

Throughout the paper, we assume each agent i can only observe her local network (g_i, G_i) . We now introduce our behavioral assumption on agents' beliefs about the network. Intuitively, each agent treats her local network as the entire network, ignoring the outside network that she cannot observe.¹⁶

ASSUMPTION 1. *Every agent believes her local network is the entire network and this is common knowledge.*

Under this prior on the network structure, agent i does not update her belief about the network from any information she receives. We call an agent with the above belief (or who acts as if she has the above belief) *locally Bayesian*. Each locally Bayesian agent learns as a Bayesian agent within her local network.

This assumption allows us to focus on an important feature of Bayesian learning, namely, agents have *perfect memory*: a Bayesian agent learns from the entire history of her neighbors' reports and her signals. Hitherto understudied in the literature, this feature sets our model apart from the myopic learning models (DeGroot (1974), DeMarzo, Vayanos, and Zwiebel (2003), Golub and Jackson (2010), among others), and the more recent learning models with imperfect memory (Molavi, Tahbaz-Salehi, and Jadbabaie (2016), Mueller-Frank and Neri (2017), Levy and Razin (2016), Dasaratha and He (2017)). The behavioral assumptions of these papers can be decomposed into two parts: agents only focus on their local networks (the same as Assumption 1), and agents can only remember their neighbors' most recent reports and treat them as new information in every period. Our model does not assume the second part since we explicitly study perfect memory. It is a necessary step toward modeling how agents try to avoid forming wrong beliefs due to repeated and distorted information from the network. Perfect memory, however, adds significantly to the complexity of characterizing the agents' short-run learning dynamics and long-run learning outcomes.¹⁷ This motivates

¹⁶The idea of a local network is flexible. Our model can be generalized to smaller local networks when each agent only observes her own links, or larger local networks including an agent's indirect neighbors. Moreover, we can model agents' learning according to her *perceived* local network, which may differ from the real one.

¹⁷In a model where agents only recall the most recent reports (often beliefs) from their neighbors, an agent's belief in period t depends only on the period- $(t - 1)$ beliefs of her neighbors. Therefore it satisfies the memoryless property of Markov chains. Classic results such as the Perron-Frobenius theorem apply, which lead to nice characterizations on learning dynamics, convergence and steady-state beliefs. In contrast, with perfect memory, an agent's belief in period t depends directly on the new information—the difference between period- $(t - 1)$ beliefs and beliefs based on earlier information shared in her local network, and thus it depends indirectly on her earlier beliefs in an iterative fashion. While there are some explorations in the theory of Markov chain with finite memory, the convergence and dynamics are cumbersome and there are no simple sufficient conditions.

us to develop a tractable learning procedure to implement locally Bayesian learning and to derive its useful properties in the next section.

Locally Bayesian learning removes an important component—learning about the network structure—from Bayesian learning. In the classic Bayesian framework, agents with non-degenerate priors should learn about the network as well as the true state from their neighbors’ reports. It is well-known, however, that the cognitive and computational cost of Bayesian learning about an unknown (and possibly large) network is very high.¹⁸ Our way of modeling reflects the high cost agents face if they were to properly account for correlations in their information by updating their beliefs about the outside network. Instead, our agents behave *as if* all the information from outside her local network are new, exogenous signals.

Before proceeding, we make two further remarks on the model.

Communication protocols. Our agents report their most up-to-date estimates of the state distribution, similar to [Lee \(1993\)](#) and [Eyster and Rabin \(2014\)](#). On one hand, these reports contain more information than actions used in many models of observational learning such as [Bala and Goyal \(1998\)](#) and [Mossel, Sly, and Tamuz \(2015\)](#). Observing these estimates is equivalent to observing actions whenever the action space is “rich” enough, so that an agent’s action fully reveals her estimates. On the other hand, these reports contain far less information than that used in [Acemoglu, Bimpikis, and Ozdaglar \(2014\)](#) where agents can tag each signal with its full travel history. When there are a large number of signals, these messages may become too complex for agents to use in reality. We adopt a commonly-used message space to aid comparison with the existing models. But the concept of locally Bayesian learning is applicable if agents communicate more or less with their neighbors.

Distribution of the state and signals. Our model is applicable to other information structures. First, it can be extended to the information partition model: agent i ’s information structure can be represented by a mapping from each state s_n to a non-empty subset of states $\mathcal{P}^i(s_n)$. When the true state is s_n , agent i considers $\mathcal{P}^i(s_n)$ to be the set of possible states. In our context, each initial signal x_0^i informs agent i of the subset $\mathcal{P}^i(s^*)$ that contains the true state s^* . Over time, agents learn to remove states that deemed impossible from their neighbors’ and indirect neighbors’ signals. Second, our model can be easily adapted to study some other standard information structures, such as the the one with uniformly distributed states and normally distributed signals.

¹⁸An agent must first form beliefs about the total number of agents in the network. For each fixed number, say I , the number of total possible networks is $2^{I(I-1)/2}$. For each of the path-connected networks among them, she assigns probabilities to all the possible signals and travel paths through which a signal may reach her. She also needs to update all these beliefs every period.

3 The locally Bayesian learning procedure

Locally Bayesian learning is easy and intuitive to define, but it is not obvious how agents with perfect memory actually learn. Therefore, we first present an iterative learning procedure that implements locally Bayesian learning. In section 3.1, we describe the key and new component of our learning procedure, that is, how agents extract new signals from their neighbors' reports. Then, section 3.2 describes the full procedure.

3.1 Extracting one new signal from each neighbor

A locally Bayesian agent remembers all the reports she and each of her neighbors shared. We use a convenient measure to store their shared reports. Specifically, agent i uses her (*interim*) *second-order* estimates about agent j 's estimates to store all the information she believes neighbor j has learned from his neighbors' reports.

Let agent i 's (*interim*) second-order estimates be \mathbf{b}_t^{ij} . It stores the information contained in the set of reports that they both observe: $\{\mathbf{b}_\tau^h : h \in g_{ij}, \tau < t\}$ (recall that $g_{ij} = g_i \cap g_j$). Formally, let $p_t^{ij}(s_n) = \Pr(s_n | \mathbf{b}_\tau^h, h \in g_{ij}, \tau < t)$ be agent i 's belief about agent j 's belief that the true state is s_n given the reports they both observed.¹⁹ As before, we use the log-likelihood ratios of these probabilities:

$$\mathbf{b}_t^{ij} = \{b_t^{ij}(s_1), \dots, b_t^{ij}(s_N)\}, \text{ where } b_t^{ij}(s_n) = \log p_t^{ij}(s_n) - \log p_t^{ij}(s_N). \quad (1)$$

By definition, \mathbf{b}_t^{ij} is an interim estimate that only stores information from their shared reports,²⁰ making it more suitable to model the agents' memory. Similarly, all the higher-order estimates in our learning procedure are interim in that they are based on past reports only. We directly refer to them as estimates instead of interim estimates from now on.

The key to locally Bayesian learning is how agent i uses her memory to extract new information from neighbor j 's report. Under Assumption 1, in agent i 's mind, agent j believes that the network is simply (g_{ij}, G_{ij}) . Therefore agent i believes that \mathbf{b}_t^{ij} stores all the information agent j has learned, except for his private signal x_{t-1}^j . This implies that agent i attributes any difference between agent j 's actual report and her estimates about j 's estimates to his private signal x_{t-1}^j . To differentiate the actual signal x_{t-1}^j from what agent i believes to

¹⁹Note that this probability is computed based on agent i 's belief of the network. That is, given Assumption 1, she believes $g = g_i, G = G_i$. It is in general different from agent j 's belief because they have different knowledge and thus beliefs about the network.

²⁰It differs from agent i 's complete estimates about agent j 's estimates, which also include her estimates of agent j 's most recent private signal (x_{t-1}^j) using her most recent signal x_{t-1}^i .

be this signal, we denote the latter as y_{t-1}^{ij} . Formally, agent i recovers y_{t-1}^{ij} (which she thinks is x_{t-1}^j) by using \mathbf{p}_t^{ij} as her prior and agent j 's actual belief \mathbf{p}_t^j (recovered from his report \mathbf{b}_t^j) as her posterior. By Bayes' rule, for any $s_n \in S$,

$$p_t^j(s_n) = \frac{p_t^{ij}(s_n)\Pr(y_{t-1}^{ij} | s_n)}{\sum_{n'=1}^N p_t^{ij}(s_{n'})\Pr(y_{t-1}^{ij} | s_{n'})}. \quad (2)$$

Taking the log-likelihood ratios of state s_n over state s_N , we have

$$\log \frac{p_t^j(s_n)}{p_t^j(s_N)} = \log \frac{p_t^{ij}(s_n)}{p_t^{ij}(s_N)} \cdot \log \frac{\Pr(y_{t-1}^{ij} | s_n)}{\Pr(y_{t-1}^{ij} | s_N)}.$$

Recall that \mathbf{b}_t^j and \mathbf{b}_t^{ij} are defined as the log-likelihood ratios according to beliefs \mathbf{p}_t^j and \mathbf{p}_t^{ij} . Then the above equation becomes:

$$b_t^j(s_n) = b_t^{ij}(s_n) + \log \frac{\Pr(y_{t-1}^{ij} | s_n)}{\Pr(y_{t-1}^{ij} | s_N)}.$$

Let α_t^{ij} be the log-likelihood ratios of the state distribution according to signal y_{t-1}^{ij} , we have

$$\alpha_t^{ij}(s_n) \equiv \log \frac{\Pr(y_{t-1}^{ij} | s_n)}{\Pr(y_{t-1}^{ij} | s_N)} = b_t^j(s_n) - b_t^{ij}(s_n). \quad (3)$$

Intuitively, agent i extracts the new signal by removing old information from agent j 's report, which is the right hand side of (3). From now on, we abuse notations slightly and refer to the log-likelihood ratios α_t^{ij} as the new signal agent i infers from neighbor j in period $t + 1$ (even though this may not be agent j 's signal at all).

3.2 The learning procedure

3.2.1 General learning procedure for any \mathbf{b}_t^{ij}

Recall that at $t = 0$, agent i receives signal x_0^i . At $t = 1$, agent i updates her beliefs about the state distribution by Bayes' rule. Her period-1 estimates are \mathbf{b}_1^i , where for each s_n ,

$$b_1^i(s_n) = \log \Pr(s_n | x_0^i) - \log \Pr(s_N | x_0^i).$$

Agent i 's second-order estimates are based on reports she and each neighbor observed. At $t = 1$, there is no such report. Since agents have symmetric priors, for each s_n ,

$$b_1^{ij}(s_n) = \log \Pr(s_n) - \log \Pr(s_N) = 0.$$

She reports b_1^i to each of her neighbors and simultaneously receives each neighbor's report b_1^j . Next, she observes her signal from nature x_1^i . Period 1 ends.

At the beginning of each period $t \geq 2$, for any given b_{t-1}^{ij} , agent i forms her period- t estimates b_t^i in two steps.

Step 1: Extracting new information. Agent i extracts a new signal α_{t-1}^{ij} from each neighbor j . From expression (3), we have,

$$\alpha_{t-1}^{ij} = b_{t-1}^j - b_{t-1}^{ij}. \quad (4)$$

Furthermore, let $\alpha_{t-1}^{ii} = \{\alpha_{t-1}^{ii}(s_1), \dots, \alpha_{t-1}^{ii}(s_N)\}$ be her signal from nature in period $t - 1$. That is, $\alpha_{t-1}^{ii}(s_n) \equiv \log \Pr(x_{t-1}^i | s_n) - \log \Pr(x_{t-1}^i | s_N)$ for each s_n .

Step 2: Updating. Agent i then updates her period- t estimates using the signals extracted from each neighbor and from nature. By Bayes' rule:

$$b_t^i = b_{t-1}^i + \sum_{h \in g_i} \alpha_{t-1}^{ih}. \quad (5)$$

She reports b_t^i to each of her neighbors and simultaneously receives each neighbor's report b_t^j . Then, she observes her signal from nature x_t^i . Period t ends. \parallel

It is worth emphasizing that an agent can use this procedure regardless of *how* she forms her second-order estimates. In particular, it easily accommodates the familiar DeGroot learning model, as well as models in which agents have imperfect memory. To see this, let agent i always set her estimates of each neighbor j 's estimates to be the uninformative prior: $\tilde{b}_{t-1}^{ij} = \mathbf{0}$ for any $t \geq 2$. Intuitively, at period t , she does not recall the reports in period $1, \dots, t - 2$. Then by expression (4), $\tilde{\alpha}_{t-1}^{ij} = \tilde{b}_{t-1}^j$. That is, she treats each neighbor's entire report at period $t - 1$ as a new signal and computes her estimates.

3.2.2 Perfect memory: using higher-order estimates to form b_t^{ij}

While the procedure above works for any second-order estimates, a locally Bayesian agent forms hers using the full history of reports she believes her neighbor j has observed. To do

so, she accounts for what she believes agent j has learned from each of his neighbors and stores it in her *third-order* estimates. She uses these third-order estimates to form her second-order estimates. Similarly, agent i forms her *higher-order* estimates about her neighbors in each clique she belongs to, and uses them to form the next lower-order estimates iteratively.

Formally, consider such a clique $\{i, j, \dots, k\}$. Recall that $g_{ij\dots k} = g_i \cap g_j \cap \dots \cap g_k$. Under Assumption 1, in agent i 's mind, agent j believes that \dots agent k believes that the network is simply $(g_{ij\dots k}, G_{ij\dots k})$. Therefore, agent i 's estimates about j 's estimates \dots about k 's estimates store all the information agent i believes agent j believes \dots agent k learned from his neighbors prior to period t . Let her beliefs based on such information be $\mathbf{p}_t^{ij\dots k}$, such that $p_t^{ij\dots k}(s_n) = \Pr(s_n \mid \mathbf{b}_\tau^h, h \in g_{ij\dots k}, \tau < t)$ for each s_n . Then, her corresponding estimates are $\mathbf{b}_t^{ij\dots k}$, where for each s_n ,

$$b_t^{ij\dots k}(s_n) = \log p_t^{ij\dots k}(s_n) - \log p_t^{ij\dots k}(s_N).$$

In this iterative procedure, agents use η th-order estimates in time t to form $(\eta - 1)$ th-order estimates in $t + 1$. To complete this procedure, without loss of generality, we let agents set *degenerate* estimates: when the last agent is repeated, agent i sets the estimates to be equal to her estimates without the last agent.²¹ Formally, for each $h \in \{i, j, \dots, k\}$, let $\mathbf{b}_t^{ij\dots kh} = \mathbf{b}_t^{ij\dots k}$. We use these degenerate estimates to reduce the agents' computations, because their learning outcomes do not change even if they form more higher-order estimates.

At $t = 1$, since there is no previous report, the initial values for all the higher-order estimates are $\mathbf{b}_1^{ij\dots k} = \mathbf{b}_1^{ij\dots kh} = \mathbf{0}$. At $t \geq 2$, agent i forms her higher-order estimates in two steps just like how she forms her estimates in Section 3.2.1.

Step 3a: Extracting higher-order new information. For any $h \in g_{ij}$, \mathbf{b}_{t-1}^{ijh} stores, from agent i 's perspective, all the old information agent j believes agent h has. In each period, agent i believes agent j extracts a new signal from herself by comparing \mathbf{b}_{t-1}^{ijj} with her report \mathbf{b}_{t-1}^j . Similarly, agent i believes agent j extracts a new signal from each common neighbor $k \in N_i \cap N_j$ by comparing \mathbf{b}_{t-1}^{ijk} with his report \mathbf{b}_{t-1}^k . Denote these new signals as α_{t-1}^{ijh} .²²

²¹We show in appendix A.1 that using these degenerate estimates is without loss of generality. Intuitively, only distinct agents matter in agent i 's higher-order estimates, because they depend on the shared reports among these distinct agents. This implies that agent i does not need to form any other estimates, for instance those with more than one repeated agent.

²²Observe that α_{t-1}^{ijj} , which is the signal agent i believes agent j gets from nature, is always equal to α_{t-1}^{ij} . This is because we set degenerate estimates $\mathbf{b}_{t-1}^{ijj} = \mathbf{b}_{t-1}^{ij}$, and then expression (6) is the same as expression (4) when agent i extracts a new signal from j . We use α_{t-1}^{ijj} to simplify formulas in step 3b.

Similar to how agent i extracts a new signal α_{t-1}^{ij} in expression (3), we have

$$\alpha_{t-1}^{ijh} = \mathbf{b}_{t-1}^h - \mathbf{b}_{t-1}^{ijh}. \quad (6)$$

Next, agent i extracts new information involving more agents in a similar way. For every clique $\{i, j, \dots, k\}$, agent i believes that agent j believes \dots that agent k extracts $\alpha_{t-1}^{ij\dots kh}$ from agent $h \in g_{ij\dots k}$, where

$$\alpha_{t-1}^{ij\dots kh} = \mathbf{b}_{t-1}^h - \mathbf{b}_{t-1}^{ij\dots kh}. \quad (7)$$

Step 3b: Updating higher-order estimates. Agent i updates her estimates of each neighbor j 's estimates by applying Bayes' rule to (what she believes to be) the new signals agent j has extracted, namely α_{t-1}^{ijh} from (6). Then we have

$$\mathbf{b}_t^{ij} = \mathbf{b}_{t-1}^{ij} + \sum_{h \in g_{ij}} \alpha_{t-1}^{ijh}. \quad (8)$$

Similarly, for every clique $\{i, j, \dots, k\}$, agent i updates her estimates of agent j 's estimates of \dots agent k 's estimates by applying Bayes' rule to (what she believes agent j believes \dots are) the new signals agent k has extracted. Namely $\alpha_{t-1}^{ij\dots kh}$ for each $h \in g_{ij\dots k}$ according to (7) above. Then we have

$$\mathbf{b}_t^{ij\dots k} = \mathbf{b}_{t-1}^{ij\dots k} + \sum_{h \in g_{ij\dots k}} \alpha_{t-1}^{ij\dots kh}. \quad (9)$$

Agent i forms all her higher-order estimates \mathbf{b}_t^{ij} , \mathbf{b}_t^{ijk} , up to $\mathbf{b}_t^{ij\dots k}$ simultaneously as she forms estimates \mathbf{b}_t^i . All these estimates are formed *before* agent i exchanges period- t reports with her neighbors. ||

To summarize, agent i learns from her neighbors in period t as in Figure 1.

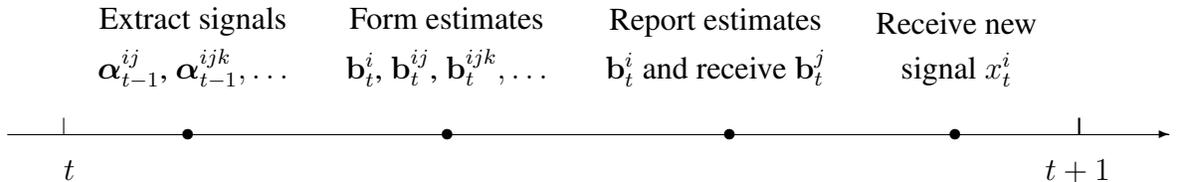


Figure 1: Time line

3.2.3 Implementing locally Bayesian learning

Agents who use our procedure have the same learning outcomes as locally Bayesian agents. In this sense, our learning procedure above can be viewed as an algorithm to implement their learning.

OBSERVATION 1. *Each agent's estimates formed in section 3.2.1 and 3.2.2 are exactly (the log-likelihood ratios of) her locally Bayesian posterior in the corresponding period.*

Intuitively, under Assumption 1, agent i believes that she knows all the links among her neighbors, and thus she can form estimates just like them. Specifically, for each neighbor j , she believes her estimates of agent j 's estimates include all the information j has learned from the reports so far, except for his most recent private signal x_{t-1}^j . She also believes that she can correctly extract x_{t-1}^j after hearing agent j 's report containing that signal. This implies that each agent i believes that she can account for any mistake a neighbor makes, and thus her own estimates b_t^i correctly includes all the signals that reach her, including $\{x_0^j, \dots, x_{t-2}^j\}$ for each neighbor j as well as her own signals.²³

While the learning procedure in section 3.2.1 and 3.2.2 yields the same outcomes as the learning of locally Bayesian agents, it is not the only way a locally Bayesian agent can learn. We choose this simple learning procedure, because it uses interim estimates which does not require the agents to know the details of their neighbors' information structure.²⁴

3.2.4 A two-agent example

The following example illustrates how our learning procedure works. Note that we use the special case of binary states and binary symmetric informative signals for all our examples.

²³It will become clear using the results in the next section that Observation 1 can be generalized by generalizing Assumption 1. Observation 1 holds if every agent believes that the network outside her local network is either empty or it consists of one or multiple unconnected components, each of which is a tree-like union of cliques with the root being one of her neighbors, and this belief is common knowledge. These types of beliefs are consistent with Fainmesser and Goldberg (2016) who show that when the population is large and the number of each agent's neighbors is bounded, each agent believes asymptotically that the network is a random tree where she is the root agent.

²⁴One alternative procedure is that in each period, agent i can update her complete estimates about agent j 's estimates including her estimates about j 's most recent signal. For example, given her initial signal x_0^i , agent i can update her complete estimates of agent j 's estimates, which is her estimates of agent j 's initial signal, if she knows agent j 's information structure (the mapping from states to j 's signals). In this case, her complete estimates of agent j 's estimates at the beginning of $t = 1$ is no longer the symmetric prior. But she does not need to do so in our learning procedure, because after she hears agent j 's report at $t = 1$, she can directly extract this signal and incorporate it into her estimates.

EXAMPLE 1. *The network has two agents and one link: $g = \{1, 2\}$ and $G = \{12\}$. The states are binary: $S = \{s_1, s_2\}$. The set of signals is $X = \{x^0, x^1, x^2\}$ (recall that x^0 is uninformative). Let agent 1 receive the only informative signal: $x_0^1 = x^1$. The Bayesian posterior given the signals is $\Pr(s_1 | x_0^1) = \phi$, $\Pr(s_2 | x_0^1) = 1 - \phi$. Let $\varphi = \log(\phi/(1 - \phi))$. The corresponding log-likelihood ratios relative to s_2 are $\{\varphi, 0\}$.*

At $t = 0$, agent 1 observes x_0^1 .

At $t = 1$, agent 1 reports her estimates based on x_0^1 : $\mathbf{b}_1^1 = \{\varphi, 0\}$. Since the states are binary and the estimates are in log-likelihood ratios, we only keep track of $b_t^i(s_1)$ in our examples. Agent 2 has no signal and reports $b_1^2(s_1) = 0$. The initial second-order estimates are $b_1^{12}(s_1) = 0$ and $b_1^{21}(s_1) = 0$. This is summarized in the first row of Table 1.

	Agent 1			Agent 2		
	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

Table 1: A two-agent example: locally Bayesian learning.

Observation: *The locally Bayesian agents' learning outcomes are correct.*

At the beginning of $t = 2$, agent 1 notices agent 2's report in period 1 agrees with what agent 1 expects him to report: $b_1^2(s_1) = b_1^{12}(s_1) = 0$. Thus agent 1 learns nothing from 2, $\alpha_1^{12}(s_1) = 0$ by expression (4). Her report does not change by expression (5): $b_2^1(s_1) = \varphi$. But agent 2 notices that $b_1^1(s_1) \neq b_1^{21}(s_1)$ and infers agent 1's signal $\alpha_1^{21}(s_1) = \varphi$. Agent 2 then adds this signal to his estimates: $b_2^2(s_1) = b_1^2(s_1) + \alpha_1^{21}(s_1) = \varphi$. Next, following step 3, agent 1 updates $b_2^{12}(s_1) = \varphi$, since she expects 2 to infer the signal from her. Similarly, because agent 2 learns that agent 1 has the signal and expects agent 1 to learn nothing from him, he updates $b_2^{21}(s_1) = \varphi$. For all $t > 2$, the agents' estimates and their estimates of each other's estimates agree. There is no new information: $\alpha_t^{12}(s_1) = \alpha_t^{21}(s_1) = 0$, and their estimates remain unchanged. Both agents believe the true state is s_1 with probability ϕ , which is the Bayesian posterior. \diamond

We now illustrate how agents learn with imperfect memory in this example. Recall from Section 3.2.1 that in DeGroot (1974) and related models, agent i treats each neighbor j 's latest report as new: $\tilde{b}_t^{ij}(s_1) = 0$ and $\tilde{\alpha}_t^{ij}(s_1) = \tilde{b}_t^j(s_1)$ for all $t \geq 2$.

Observation: *Exponential repetition of signals with imperfect memory.*

At $t = 2$, just like before, agent 1 and agent 2 report $\tilde{b}_2^1(s_1) = \varphi$ and $\tilde{b}_2^2(s_1) = \varphi$. At $t = 3$, however, agent 1 treats 2's report as new information, $\tilde{\alpha}_2^{12}(s_1) = \varphi$, and updates her estimates to $\tilde{b}_3^1(s_1) = 2\varphi$ as if she has received two copies of x_0^1 . Similarly, agent 2 reports $\tilde{b}_3^2(s_1) = 2\varphi$. At each period $t \geq 4$, each agent treats the neighbor's entire report as new information in the same way as in $t = 3$. Their estimates grow exponentially: $\tilde{b}_t^1(s_1) = \tilde{b}_t^2(s_1) = 2^{t-2}\varphi$.

	Agent 1			Agent 2		
	\tilde{b}_t^1	\tilde{b}_t^{12}	$\tilde{\alpha}_{t-1}^{12}$	\tilde{b}_t^2	\tilde{b}_t^{21}	$\tilde{\alpha}_{t-1}^{21}$
$t = 1$	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: learning with imperfect memory.

It is worth examining why imperfect memory leads to severe overcounting in such a simple network. At $t = 3$, because agent 1 does not recall agent 2 has learned the signal from herself at $t = 2$, she believes agent 2 receives a new signal. Thus a signal goes back and forth between every two neighbors, and is treated as new each time. To remove such repetition, each agent in a pair must recall who has received the signal first. Both should expect the other agent to learn it in the next period. This is exactly what our learning procedure does.

3.3 Main properties of locally Bayesian learning

Two nice properties of locally Bayesian learning make our procedure tractable and simplify the ensuing analysis. First, the Bayesian part implies that a signal travels through the network independent of other signals. Specifically, the learning outcomes of an agent given multiple signals can be decomposed as follows: divide the full sequence of realized signals by the end of period $t - 1$, X_{t-1} , into any two disjoint sets of signals, X_{t-1}^μ and X_{t-1}^ν . Recall that \mathbf{b}_t^i is agent i 's estimates when X_{t-1} is the set of signals from nature. Let $\mathbf{b}_t^{\mu,i}$ and $\mathbf{b}_t^{\nu,i}$ be her estimates when the set of signals from nature is X_{t-1}^μ and X_{t-1}^ν , respectively.

LEMMA 1. *For any $t \geq 1$,*

$$\mathbf{b}_t^i = \mathbf{b}_t^{\mu,i} + \mathbf{b}_t^{\nu,i}; \quad (10)$$

$$\mathbf{b}_t^{ij} = \mathbf{b}_t^{\mu,ij} + \mathbf{b}_t^{\nu,ij}. \quad (11)$$

Lemma 1 shows that the agent’s estimates under X_{t-1} are equal to the sum of her estimates under X_{t-1}^μ and X_{t-1}^ν . Intuitively, because agent i treats all her inferred signals as independent, one signal travels, with all its possible repetitions and distortions, independently from another. This lemma allows us to study one signal at a time. It implies that if the agents’ learning outcomes are correct under every signal, their learning outcomes are also correct under a sequence of these signals.

The second property characterizes the travel of each signal over time through the network. Recall that a locally Bayesian agent uses Bayes’ rule in each period to extract information (expression (4)) and to incorporate it into her own estimates (expression (5)). So do her neighbors. Combining these two steps, the new signal agent i extracts from agent j is the unexpected change in j ’s report, which is due to what j has learned in the previous period. Thus, we have an iterative rule characterizing how each signal travels.

LEMMA 2. For any $t \geq 2$,

$$\alpha_t^{ij} = \sum_{l \in (g_j \setminus g_i) \cup \{j\}} \alpha_{t-1}^{jl} + \sum_{h \in g_{ij} \setminus \{j\}} (\alpha_{t-1}^{jh} - \alpha_{t-1}^{ijh}). \quad (12)$$

On the left hand side, α_t^{ij} is the signal agent i extracts from neighbor j at the beginning of period $t + 1$, which can be decomposed into two parts according to equation (12). The first part consists of what agent j has just learned from nature (α_{t-1}^{jj}) and from his neighbors who are not connected to agent i (α_{t-1}^{jl} for $l \in g_j \setminus g_i$). That is, the first part is information from sources outside agent i ’s local network. The second part consists of a potential error term whenever agent i and j share at least one common neighbor, say agent h . Each of the differences ($\alpha_{t-1}^{jh} - \alpha_{t-1}^{ijh}$) is the difference between what agent j learned from h and what agent i believes agent j learned from agent h . The second term is zero if either agent i and j have no common neighbors, or when their local networks satisfy a property—local connection symmetry—defined in the next section.

Lemma 2 suggests that in locally Bayesian learning, a signal travels to agent i through her neighbor j from non-neighbors of agent i . That is, the first part of expression (12) does not include what agent j has learned from i (no α_{t-1}^{ji}). So there is no *pairwise feedback*. It also does not include what agent j has learned from a common neighbor k (no α_{t-1}^{jk} , $k \in N_i \cap N_j$). So there is no *triangular feedback*. In this sense, locally Bayesian learning involves far fewer repetitions than models with imperfect memory such as DeGroot (1974), which include both pairwise and triangular feedbacks.²⁵

²⁵In models with imperfect memory, we count the number of paths with a length of at most τ between agent i

4 When are learning outcomes efficient?

Can agents learn correctly given the signals a network receives? How do their learning outcomes depend on the network structure? Before answering this central question of any network learning model, we lay out our notions of correct learning. Our benchmark is for each agent to learn correctly in every period *given the travel paths of signals*.

We begin with the set of signals that can reach agent i in period t . Recall that X_t contains all the realized signals the network receives by the end of period t . It is the union of X_t^i , the set of signals agent i receives up to and including period t from nature. Also, recall that agent i receives no informative signals at or after period $T_i \in [1, \infty]$. Let $T = \max_i T_i$. Then X_T is the sequence of all the realized signals the network receives. Let $d(il)$ be the *distance*, or the length of the shortest path, between agent i and agent l , $l = 1, \dots, I$, with $d(ii) = 0$. The *diameter* of the network is D , which is the longest distance between any two agents. It takes one period for agent l to incorporate a private signal into his report, and then $d(il)$ periods for the signal to travel from l to i . Therefore at the beginning of period t , the set of agent l 's signals that can reach agent i is $X_{t-d(il)-1}^l$ (we let $X_{t-d(il)-1}^l = \emptyset$ if $t < d(il) + 1$). Suppose that agent i correctly learns every signal that has reached her at the beginning of t , then for every $s_n \in S$, her Bayesian posterior is:²⁶

$$q_t^i(s_n) = \Pr(s_n \mid X_{t-d(i1)-1}^1, \dots, X_{t-d(iI)-1}^I). \quad (13)$$

DEFINITION 1. For all sequences of realized signals X_T ,

- Agent i 's learning is **strongly efficient** if her report in period t is the log-likelihood ratio of her Bayesian posterior: $b_t^i(s_n) = \log q_t^i(s_n) - \log q_t^i(s_N)$.
- Agent i 's learning is **efficient** if her report converges to the log-likelihood ratio of the Bayesian posterior: $\lim_{t \rightarrow \infty} b_t^i(s_n) = \log \Pr(s_n | X_T) - \log \Pr(s_N | X_T)$.
- Agent i 's learning is **asymptotically efficient** if she learns the true state almost surely as $t \rightarrow \infty$ when everyone receives an arbitrarily large number of signals ($T_l = \infty$ for every $l \in g$).

and h , to find the total number of copies of signal x_0^i that will travel to agent h in period $\tau + 1$. Locally Bayesian learning reduces this count, and thus the repetition, dramatically by removing paths involving pairwise feedback or triangular feedback.

²⁶More precisely, $q_t^i(s_n)$ is the probability that agent i believes the state is s_n when she knows the entire network, all agents report all the signals they know, and each signal is tagged with its time of arrival and the agent who initially receives it.

Strong efficiency is the strongest notion of correct learning in this context.²⁷ Therefore we use strong efficiency to prove our positive result, showing that the agents learn correctly in every period, not just eventually. Efficiency and asymptotic efficiency are weaker notions we use to prove our negative results about the agents’ learning errors. When everyone receives an arbitrarily large number of signals, we adopt asymptotic efficiency—the most commonly used measure of learning outcomes in the literature. However, it is not appropriate when the agents only receive a finite number of signals (T is finite), as the correct Bayesian posterior is bounded away from 0 and 1. In this case, we use the notion of efficient learning, which requires the agent’s estimates in the long run to match the Bayesian posterior.

4.1 Strongly efficient learning in social quilts

To learn correctly, an agent must treat a signal as new information once and only once. In particular, the agent must not count it as a new signal at any point after her first encounter with the signal. Given that each agent only exchanges reports with her neighbors, her local network as well as the entire network (even though she does not know it) need to meet certain conditions for strongly efficient learning. We now show that a particular type of networks, *a social quilt*, and only this type of networks, meets these conditions. We say a path $(i_1 \dots i_l)$ is a *circle* if $i_1 i_l \in G$.

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

Definition 2 requires that in a social quilt, any circle must be embedded in a clique. In a *tree*, any two nodes are connected by a unique path. Intuitively, a social quilt can be thought of as a tree of cliques. Figure 2 shows a social quilt, which in general could include subnetworks such as the well-known trees, cliques, stars, lines, and some of the core-periphery networks.²⁸ Our main result is an intuitive and clean relationship between social quilts and strongly efficient learning outcomes.

²⁷It implies that when T is finite, all agents form the correct posterior at or before period $T + D$. This is a strong notion because even when the network is common knowledge and all the agents are Bayesian, it often takes much longer than the diameter of the network for agents to learn (see Mossel, Olsman, and Tamuz (2016)). Strong efficiency implies efficiency, and both are stronger than asymptotic efficiency if the information structure is such that the agents can differentiate every state from another at the network level. We assume such an information structure in section 5 when we study asymptotic efficiency because otherwise asymptotic efficiency can never be achieved regardless of the number of signals a network receives.

²⁸The overall tree structure is important theoretically. For example, the limit of a large Erdős-Rényi network with bounded degree is a random tree, and the binary tree has high *expansiveness* as defined by Ambrus, Möbius, and Szeidl (2014) which they show are important for risk sharing networks. In addition, some networks with the core-periphery structure are social quilts, which are important for financial markets (Babus and Kondor

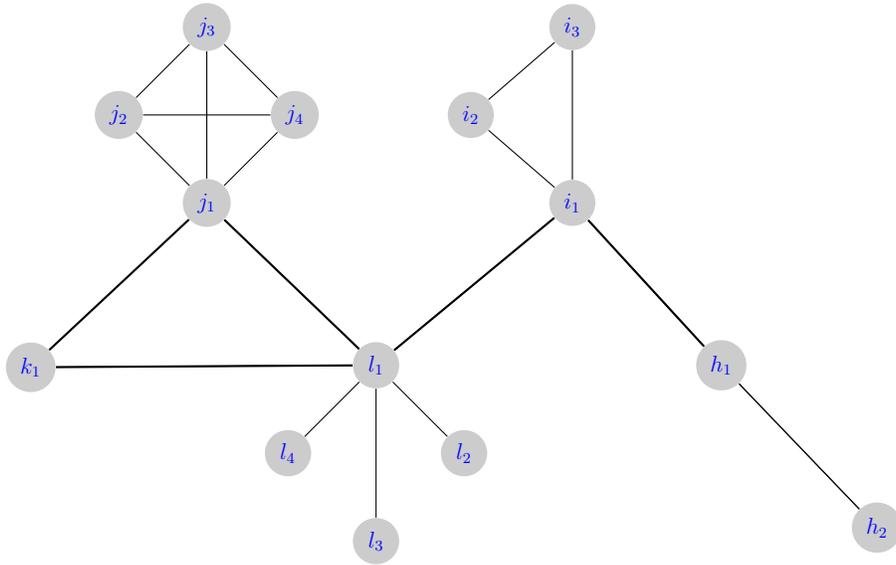


Figure 2: A social quilt

PROPOSITION 1. *All agents' learning outcomes are strongly efficient if the network is a social quilt. Otherwise, there exists some sequence of realized signals such that at least one agent's learning outcomes are not strongly efficient.*

Intuitively, two important features of social quilts—a global tree and local cliques—enable locally Bayesian agents to learn correctly. First, a global tree makes sure each signal will not travel back to an agent for a second time, which only happens via a circle (that is not embedded in a clique). Second, each local clique makes sure information is not distorted locally. That is, whenever one agent gets a new piece of information, all her neighbors in the clique learn it from her and expect all others to do the same. These two features ensure that each of the signals an agent extracts from her neighbors is truly independent, and thus the locally Bayesian agents' learning outcomes are strongly efficient.

We now examine these two features and their respective roles in Proposition 1 in more depth. A global tree means that there is no *simple circle* in a social quilt.

DEFINITION 3. *A **simple circle** is a circle that contains at least four agents and each agent has exactly two links to other agents in the circle.*

(2017)). This occurs when a few core members are connected in a clique and peripheries are connected to one core member. Jackson, Rodriguez-Barraquer, and Tan (2012) and Ali and Miller (2013) show social quilts and cliques are important for favor exchanges and cooperation in the network.

Whenever a network has simple circles, there are multiple paths between one agent and another.²⁹ As a result, each signal could travel along these simple circles and reach an agent repeatedly. For example, if $(ijkh)$ is a simple circle, then agent i will double count if, say two of her neighbors j and k inferred the same signal from agent h . This learning error is not present in a social quilt because there is no simple circle. Therefore if there is one informative initial x_0^l , agent i learns this signal for the first time at period $d(il) + 1$ through the unique shortest path from l to i . This path partially coincides with the unique path from agent l 's clique to agent i 's clique. When each signal reaches the “terminal” clique(s) of the tree, the signal stops traveling.

Next, local cliques ensure every agent's local network satisfies the following property.

DEFINITION 4. *Agent i 's local network satisfies **local connection symmetry** if g_{ij} is a clique for every $j \in N_i$.*

First, local connection symmetry for agent i (LCS_i from now on) holds if for any neighbor $j \in N_i$, $N_i \cap N_j = \emptyset$, which is the case when they are part of a simple circle or a line. For example, consider agent h_1 in Figure 2. She does not have any common neighbor with any of her neighbors i_1 and h_2 , so LCS_{h_1} holds. It also holds if each pair of agent i and j 's common neighbors k and l are connected. For instance, consider subnetwork $\{j_1, j_2, j_3, j_4\}$ in Figure 2. Clearly, each agent's shared local network with another agent is a clique. In contrast, suppose that we take away the link j_2j_4 , then agent j_1 's local network fails this property, because j_1 and j_3 have two common neighbors j_2 and j_4 who are not connected.

We say that a network satisfies *local connection symmetry* (LCS for now on) if property LCS_i holds for all $i \in g$. This property makes agents' learning particularly simple.

LEMMA 3. *For any agent $j \in N_i$, $\mathbf{b}_t^{ij} = \mathbf{b}_t^{ji}$. Moreover, if (g_i, G_i) satisfies LCS_i , then for any clique $\{i, j, l, \dots, k\}$,*

$$\mathbf{b}_t^{ij} = \mathbf{b}_t^{il} = \dots = \mathbf{b}_t^{ik}, \text{ and } \mathbf{b}_t^{ij} = \mathbf{b}_t^{ijl} = \dots = \mathbf{b}_t^{ijl\dots k}. \quad (14)$$

The first part of the lemma says, in any network, agent i and j have the same second-order estimates about each other, even if their estimates may differ. This is very intuitive: agent i forms her second-order estimates using reports from their shared local network, namely, $\{\mathbf{b}_\tau^h : h \in g_{ij}, \tau < t\}$, and so does agent j . Then, the lemma goes on to show if agent i 's local network satisfies LCS_i , her estimates about her two connected neighbors j and l must

²⁹More precisely, $C = (i_1 \dots i_k)$ is a simple circle if $k \geq 4$ and $N_{i_l} \cap C = \{i_{l-1}, i_{l+1}\}$ for any $2 \leq l \leq k-1$. Moreover $N_{i_1} \cap C = \{i_2, i_k\}$ and $N_{i_k} \cap C = \{i_1, i_{k-1}\}$.

be the same. The reason is that given LCS_i , $g_{ij} = g_{il}$ if her neighbors j and l are connected, which is the clique all three of them belong to. From agent i 's perspective, agents in g_{ij} share the same reports as those in g_{il} , implying that $\mathbf{b}_t^{ij} = \mathbf{b}_t^{il}$. Take the clique $\{j_1, j_2, j_3, j_4\}$ in Figure 2. Agent j_1 thinks j_2 can see the reports from all four of them, and thinks j_3 sees the same set of reports, so $\mathbf{b}_t^{j_1 j_2} = \mathbf{b}_t^{j_1 j_3}$. The second part of (14) says that agent i 's higher-order estimates are the same as her second-order estimates for neighbors in the same clique. Therefore she does not need to keep track of estimates higher than second-order ones.

To see this lemma at work, recall that the second part (the error term) of the iterative rule characterizing a signal's travel in expression (12) from Lemma 2 is $\sum_{h \in g_{ij} \setminus \{j\}} (\alpha_{t-1}^{jh} - \alpha_{t-1}^{ijh})$. If the network satisfies LCS, then information is not distorted locally because all these differences are zero. More precisely, if i, j, h belong to the same clique, the signal agent i thinks j has inferred from h is exactly what agent j inferred from h : $\alpha_{t-1}^{ijh} = \alpha_{t-1}^{jh}$. This follows from Lemma 3 because $\mathbf{b}_{t-1}^{ijh} = \mathbf{b}_{t-1}^{ij} = \mathbf{b}_{t-1}^{ji} = \mathbf{b}_{t-1}^{jh}$.

In this way, every agent learns each signal correctly the first time it reaches her clique by local connection symmetry, and it never travels back to her again because there are no simple circles. Then by Lemma 1, the agents learn all the signals correctly. Specifically, agent i 's estimates at period t include signals observed by each agent l from period 0 to period $t - d(il) - 1$, and thus her learning outcomes are strongly efficient. Clearly, Proposition 1 implies that the agents' learning outcomes are asymptotically efficient under mild conditions (defined formally above Proposition 2 in Section 5.1).

Proposition 1 shows that social quilts are also necessary for strongly efficient learning to hold for any realized sequence of signals. When a network is not a social quilt, it must either contain simple circles or fail LCS. Each of them leads to a specific type of learning error, which we turn to in the next section.

5 When efficient learning is impossible

5.1 Repetition in echo chambers

To isolate the learning error caused by simple circles, we consider a network that satisfies LCS, but it is not a social quilt. Thus it contains at least one simple circle. In such a network, all agents make the error of repetition, believing they receive many independent signals which are in fact all perfectly correlated copies. Intuitively, because each agent only knows her local network, she may keep inferring “new” signals from her neighbors when it is the

same signal reaching her again and again. Consider the following example.

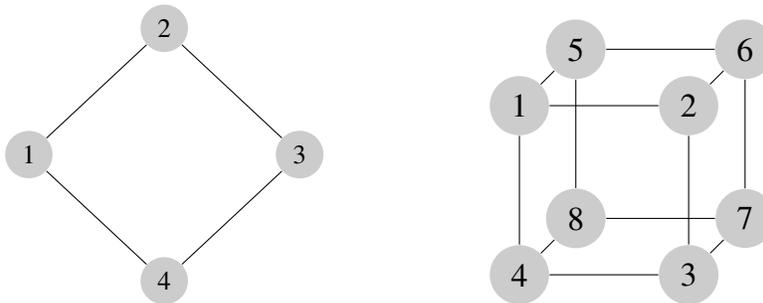


Figure 3: (a) a four-agent simple circle

(b) a cube

EXAMPLE 2. *Four agents are connected in a simple circle as in Figure 3(a). Recall from Example 1 that $S = \{s_1, s_2\}$, $X = \{x^0, x^1, x^2\}$, and $\Pr(x^1 | s_1) = \Pr(x^2 | s_2) = \phi$. Agent 1 receives the only informative signal $x_0^1 = x^1$. The corresponding log-likelihood ratio is $\log(\Pr(s_1 | x^1)/\Pr(s_2 | x^1)) = \varphi$.*

	b_t^1	$b_t^2 = b_t^4$	b_t^3
$t = 1$	φ	0	0
$t = 2$	φ	φ	0
$t = 3$	φ	φ	2φ
$t = 4$	φ	2φ	2φ
$t = 5$	3φ	2φ	2φ

Table 3: Learning in a simple circle

The signal x_0^1 travels from agent 1 in both directions around the simple circle. Agent 1 incorporates x_0^1 into her estimates at $t = 1$. At $t = 2$, agent 2 and 4 learn it and incorporate it into their reports. At $t = 3$, agent 3 learns two copies of the signal, one from 2 and the other from 4. At $t = 4$, expression (4) yields $\alpha_3^{23}(s_1) = \alpha_3^{43}(s_1) = 1$. That is, agent 2 (and 4) infers a second copy of the signal from agent 3. This is because he expects agent 3 to learn one copy from himself but agent 3 reports 2φ . At $t = 5$, agent 1 learns two copies of the signal, one from agent 2 and the other from agent 4, and thus she believes there are three copies of the signal. In every four-period interval, the agents learn two additional copies of the signal in the same way as in the four periods from $t = 2$ to $t = 5$. In particular, in each period $t = 4\tau + 1$, where $\tau = 0, 1, \dots$, agent 1 believes in $2\tau + 1$ copies of the signal and all other three agents believe in 2τ copies. \diamond

The error of repetition occurs in networks with simple circles more generally, and can persist even when the network receives a large number of informative signals. To show this, we impose two restrictions on the agents' information structure in this section. First, we assume for any pair of states s and s' , there exists some agent i 's possible signal $x^{i,m}$ such that $\Pr(x^{i,m} | s) \neq \Pr(x^{i,m} | s')$. This ensures that it is possible to differentiate every pair of states at a network level. Second, for each signal $x^{i,m}$, there exists a unique state s such that $s = \arg \max_{s_n} \Pr(s_n | x^{i,m})$. That is, each signal favors a unique state. The first assumption is necessary for the agents to learn the true state. After all, if the agents cannot differentiate one state from another collectively, they can not learn which one is the true state regardless of the number of signals. The second assumption ensures that each signal indicates a unique mostly likely state, which always holds for informative signals when the states are binary.³⁰

PROPOSITION 2. *Suppose that a network satisfies LCS, but it contains $\kappa_{sc} \geq 1$ simple circles.*

1. *With a finite number of informative signals, no agent's learning outcome is efficient.*
2. *When each agent receives an infinite number of informative signals, if $\kappa_{sc} = 1$, the agents' learning outcomes are asymptotically efficient; if $\kappa_{sc} > 1$, the agents' learning outcomes are not asymptotically efficient with a positive probability.*

The first part of the result generalizes the error of repetition from Example 2. Consider the case of only one informative signal (x_0^i); it is repeatedly learned by agents in the network because of the simple circle(s). As time goes on ($t \rightarrow \infty$), every agent is wrong, because they believe in the state that is most likely given x_0^i . But the correct Bayesian posterior is bounded away from 0 and 1.

To see whether the agents' learning is asymptotically efficient, we need to study the rate of repetition. We begin with one simple circle of k agents and agent i learns a signal at time t . The signal travels in both directions, reaching all other $k - 1$ agents in the simple circle. At time $t + 1 + k$, agent i infers two new copies of this signal. Similarly, each agent in the simple circle learns two new copies every k periods after the signal reaches him initially, just like in Example 2. The key is that all these repeatedly inferred signals grow at the same rate—two additional copies per k periods—for each signal that reaches the simple circle. Therefore with multiple signals, only the relative precision of these signals, not their arrival

³⁰The first assumption is commonly used to study asymptotically efficient learning in the literature. It allows for the possibility that one agent's signals cannot tell certain states apart, and thus she needs to learn from others in the network. The second assumption makes it simpler to characterize the learning outcomes since each signal indicates a unique most likely state. While learning errors persist more generally even when the second assumption fails, the exposition is more cumbersome.

times, matters. When each agent receives an infinite number of informative signals, the Law of Large Numbers still holds and everyone learns asymptotically.

With multiple simple circles, however, the agents’ learning outcomes become qualitatively worse because of the feedbacks among simple circles. Specifically, each signal travels both within each simple circle, and travels back and forth from one simple circle to another. Agents in one simple circle keep inferring more and more new signals from all the other $\kappa_{sc} - 1$ simple circles, and passing their own repeatedly inferred signals to them. This leads to an exponential growth of the number of copies of each signal. The second part of the result shows the severity of the error of repetition. In any network with two or more simple circles, there exists a period such that after that period, agents can receive an arbitrarily large number of correct signals—signals that are the most informative of the true state—but they still believe in the wrong state. This is because each of the correct new signals arrives too late, and is dominated by the rapidly growing existing signals. This error is more likely to occur if the agents receive uninformative signals with a high probability in each period. In this case, any early informative signal has a lot of time to grow and to dominate later signals.

This suggests that fake news—propaganda and disinformation pretending to be real news—may thrive in networks containing multiple simple circles (“echo chambers”).³¹ Moreover, “facts might not beat falsehoods”: an objective source of information has limited ability to reduce the influence of fake news in the presence of echo chambers. To be more concrete, consider the network depicted in the right panel of Figure 3.

EXAMPLE 3. *Eight agents are connected in a cube as in Figure 3(b). The information structure is the same as before. The true state is s_1 . Suppose that each agent observes $x_0^i = x^2$ at $t = 0$, and $x_t^i = x^1$ for all $t \geq 1$. As $t \rightarrow \infty$, everyone believes the true state is s_2 with a probability arbitrarily close to 1.*

Why do they believe in state s_2 even when they receive so many opposing (and correct) signals from $t = 1$ onward? Observe that at $t = 1$, each agent reports $b_1^i(s_1) = -\varphi$ which is based on the initial signal x^2 . At $t = 2$, each agent infers three signals of x^2 from their neighbors net her own signal of x^1 received by the end of $t = 1$, so her count of copies of x^2 increases by two and she reports her estimates $b_2^i(s_1) = -3\varphi$. Her estimates of each neighbor j ’s estimates are $b_2^{ij}(s_1) = -2\varphi$, because she thinks that j learns a signal of x^2

³¹This is a common theme of discussions following the Brexit campaign. For instance, see Bell, Emily. “The truth about Brexit didn’t stand a chance in the online bubble.” Guardian, July 3, 2016. Moreover, if we extend the model such that agents shares fake news more often than the truth as suggested by Vosoughi, Roy, and Aral (2018), then with echo chambers, a slight increase in the sharing of fake news can lead to their total dominance.

from herself plus his own signal of x^2 . Therefore at $t = 3$, each agent infers another x^2 from each neighbor, $\alpha_2^{ij}(s_1) = b_2^j(s_1) - b_2^{ij}(s_1) = -\varphi$. That is, agent i infers three new signals of x^2 from neighbors, net of one copy of x^1 from nature, which is exactly the same as period 2. In fact, the agents' learning in each later period is identical to that in period 2. In the limit, they believe the true state is s_2 with probability 1. \diamond

5.2 Opinion swings due to local asymmetric information

Proposition 1 shows that for strongly efficient learning, the network must contain no simple circles and satisfy LCS. We now isolate the role of the second feature by considering a network that fails LCS even though it has no simple circles. At the end of this subsection, we discuss the agents' learning outcomes when both features fail.

If a network fails LCS, a novel type of learning error arises, namely, belief oscillation and non-convergence. To see this, consider the following example. In the network in Figure 4(a), agent 2's local network fails LCS_2 (similarly for agent 4). This means that agent 2 and 4 know when they learn from a common neighbor 1, but agent 3 does not.

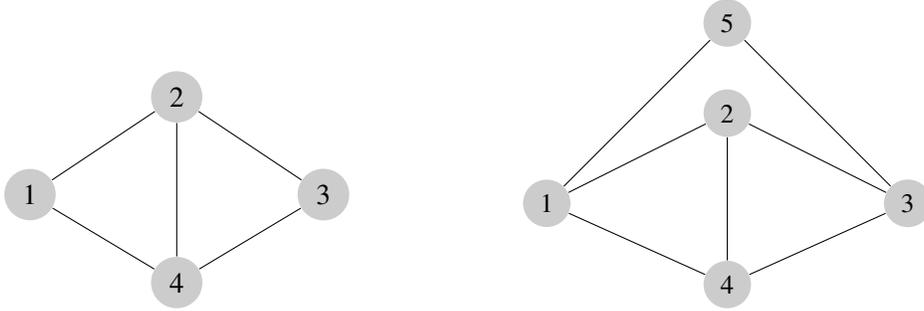


Figure 4: (a) A diamond with a link

(b) Expanded diamond with a link

EXAMPLE 4. Consider Figure 4(a), which is a diamond with a link. The information structure is the same as before. Let $x_0^1 = x^1$ be the only informative signal. The corresponding log-likelihood ratio remains $\log(\Pr(s_1 | x^1) / \Pr(s_2 | x^1)) = \varphi$.

At $t = 1$, agent 1 incorporates the signal x_0^1 and reports $b_t^1(s_1) = \varphi$. At $t = 2$, agent 2 and 4 learn the signal from agent 1, and thus $b_2^2(s_1) = b_2^4(s_1) = \varphi$. Since agent 2 and 4 know the entire network, they form the correct Bayesian posterior. So can agent 1 since he will not learn new information from 2 and 4, that is, $b_t^1(s_1) = b_t^2(s_1) = b_t^4(s_1) = \varphi$ for $t \geq 2$.

At $t = 3$, agent 3 learns two signals, one from agent 2 and one from agent 4, so $b_3^3(s_1) = 2\varphi$. Also, agent 3 believes agent 2 and 4 should learn from each other because he believes

	b_t^1	$b_t^2 = b_t^4$	b_t^3	$\alpha_{t-1}^{32} = \alpha_{t-1}^{34}$
$t = 1$	φ	0	0	n/a
$t = 2$	φ	φ	0	0
$t = 2\tau + 1, \tau \in \mathbb{N}$	φ	φ	2φ	φ
$t = 2\tau + 2, \tau \in \mathbb{N}$	φ	φ	0	$-\varphi$

Table 4: Learning in a diamond with a link.

these two signals are independent. That is, $b_3^{32}(s_1) = b_3^{34}(s_1) = 2\varphi$. More interestingly, at $t = 4$, agent 3 compares $b_3^2(s_1) = \varphi$ with $b_3^{32}(s_1) = 2\varphi$, and infers $\alpha_3^{32}(s_1) = -\varphi$, a copy that is negatively correlated with the initial signal. He infers another negatively correlated copy from agent 4, and thus $b_4^3(s_1) = 0$. Intuitively, a locally Bayesian agent 3 can only justify the fact that agent 2 and 4 do not learn from each other by believing that they have each learned an offsetting signal. Agent 3's learning in the later periods oscillate in the same way. That is, in each odd period, he reports 2φ and in each even period, he reports 0.

In contrast with the simple circle in Example 2, both agent 2 and 4 expect that agent 3 believes in two copies in odd periods and 0 in even periods because they know agent 3 does not know the existence of agent 1. They expect agent 3 to learn this way, and thus their own estimates are not affected by agent 3's opinion swings. \diamond

The above example illustrates that in a small network, the failure of LCS affects agents differently: those who know more about their local networks may learn correctly, but those who know less have long-lasting opinion swings. This oscillation and non-convergence could persist even if the network receives a large number of signals.

PROPOSITION 3. *Consider a network with no simple circles, but fails LCS. Then there exists a sequence of signals X_∞ such that at least one agent's learning outcomes are not efficient (and not converging), even though $\Pr(s' | X_\infty) = 1$ for some state s' ,*

A key feature of networks without simple circles is that a signal travels sequentially, away from the agent who receives it. If agent i receives a signal, we can classify the agents by their distance to agent i , $N_i^d = \{h \in g : d(ih) = d\}$. Then no agent in N_i^d infers any new signals from her successors in N_i^{d+1} . Otherwise the same information must reach an agent through two different paths, implying that a simple circle exists. This one-way travel of signals allows us to keep track of agents' learning. Suppose that agent i has a neighbor j whose local network fails LCS_j (for instance, agent i is agent 1 in Example 4). If she gets one correct signal, the estimates of at least one agent (say agent l , who is agent 3 in

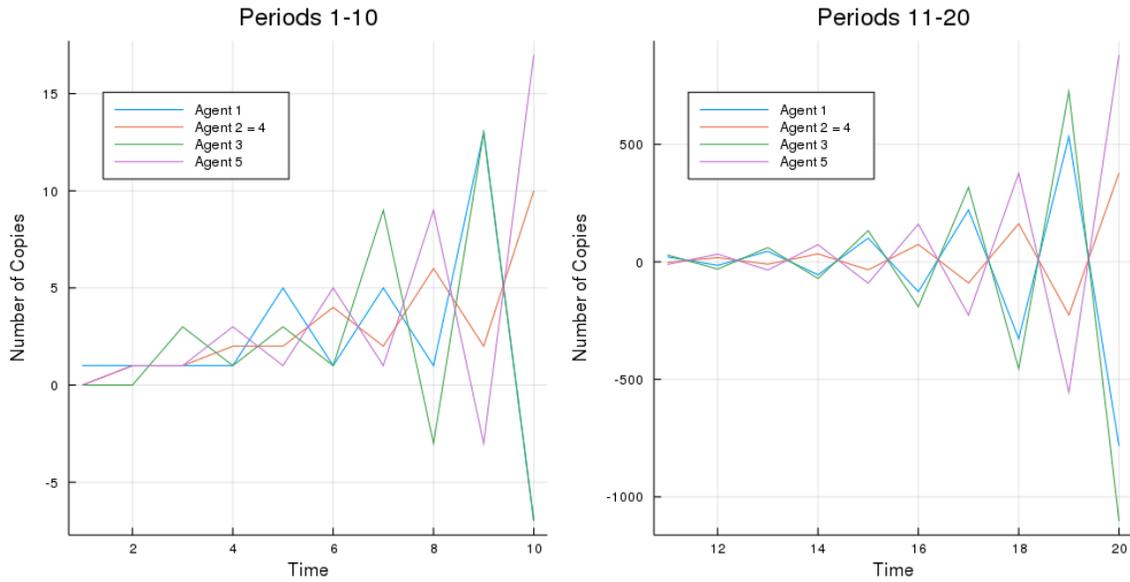


Figure 5: Simulation results for the expanded diamond with a link

the example above) cannot stop oscillating once he starts. If agent i receives more correct signals, it exacerbates agent l 's oscillation. Moreover, all the successors of agent l would have opinion swings—possibly divergent opinion swings if any of their local networks fails LCS. This type of learning error may lead to unreliable poll results, experimental outcomes and other empirical observations taken at a particular time.

If a network has simple circles and fails LCS, both repetition and belief oscillations occur locally. But we cannot provide a full characterization of the agents' learning outcomes because this problem does not have enough structure in general. This is because whenever a signal reaches a subnetwork that fails LCS, some agent in the subnetwork will infer signals negatively correlated with the original signal. Unlike in Proposition 3, the presence of simple circles means that both the positively correlated copies of this signal (due to repetition) and the negatively correlated copies (due to belief oscillation) are propagated throughout the network. There is no simple rule to calculate the net number of signals for each network.³²

We conjecture that non-convergence is robust in networks that have simple circles and also fail LCS. The intuition is that the (endogenously) generated negatively correlated signals are just as strong as the positively correlated signals. For example, consider the network

³²While one can treat each agent's estimates and all her higher-order estimates as one set of estimates to form a memoryless Markov process, each of these estimates are updated via a matrix with both positive and negative entries (negative signs from removing old information). There is no sufficient condition for convergence, without which it is difficult to characterize the long-run outcomes.

in Figure 4(b), an expanded diamond with a link which contains two simple circles (1235) and (1435). What happens if agent 1 receives an initial signal of x^1 ? It travels through both the simple circles and the diamond with a link. The agents initially believe the true state is more likely to be s_1 due to the simple circles. But each time these positively correlated signals reach agent 3 through the diamond with a link, she will infer as many negatively correlated copies. In short, to every positive correlated signal there is always an equal negatively correlated signal. Figure 5 shows that the agents begin to oscillate quickly. As time goes on, every agent alternates between believing in s_1 and s_2 . Other simulation results suggest similar patterns of diverging opinion swings in these networks.

6 Conclusion

Our modeling approach is primarily positive: we want to study the agents' learning outcomes even if they only know their local networks. The agents try to discern new information from old information in a locally Bayesian way. This approach brings the predictions of our model closer to the actual learning outcomes of agents with limited network knowledge. It also adds more sophisticated Bayesian reasoning to existing models with imperfect memory. Moreover, locally Bayesian learning is far more tractable than Bayesian learning. As such, it is potentially useful for other network learning models.

Our results also have normative implications. First, if we relax the behavioral assumption which makes agents believe information from outside their local networks is independent, agents may account for repeated information from outside their local networks by some simple rule. For example, they dismiss any signal they have already inferred as old information. We can show that with this simple rule, their learning outcomes are strongly efficient in any network if all signals are generic—each signal has its own idiosyncratic noise—and they reach the *same* agent initially. Therefore it may be desirable for a policymaker to disseminate information through one central agent over time. Second, our results show that while it is easy for wrong beliefs to propagate in networks with multiple simple circles, it is much harder to correct these beliefs. Giving everyone the correct information alone can be ineffective, but exaggerating the truth may lead to extreme views in the opposite direction. It may be more fruitful to work on the network structure, say by building more links among agents to create a (local) clique; or by removing links to avoid simple circles. In addition, one can encourage the agents to share the sources of their information by tagging their reports locally similar to [Mobius, Phan, and Szeidl \(2015\)](#). These are possible topics for future study.

References

- ACEMOGLU, D., K. BIMPIKIS, AND A. OZDAGLAR (2014): “Dynamics of information exchange in endogenous social networks,” *Theoretical Economics*, 9(1), 41–97.
- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): “Bayesian Learning in Social Networks,” *Review of Economic Studies*, 78(4), 1201–1236.
- ALATAS, V., A. BANERJEE, A. G. CHANDRASEKHAR, R. HANNA, AND B. A. OLKEN (2016): “Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia,” *American Economic Review*, 106(7), 1663–1704.
- ALEVY, J. E., M. S. HAIGH, AND J. A. LIST (2007): “Information Cascades: Evidence from a Field Experiment with Financial Market Professionals,” *Journal of Finance*, 62, 51–180.
- ALI, S. N., AND D. A. MILLER (2013): “Enforcing Cooperation in Networked Societies,” *working paper*.
- AMBRUS, A., M. MÖBIUS, AND A. SZEIDL (2014): “Consumption risk-sharing in social networks,” *American Economic Review*, 104(1), 149–182.
- ANDERSON, L. R., AND C. A. HOLT (1997): “Information Cascades in the Laboratory,” *American Economic Review*, 87, 847–862.
- BABUS, A., AND P. KONDOR (2017): “Trading and information diffusion in over-the-counter markets,” *Working paper*.
- BALA, V., AND S. GOYAL (1998): “Learning from Neighbours,” *The Review of Economic Studies*, 65(3), 595–621.
- BANERJEE, A. V. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–819.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of Political Economy*, 10(1), 992–1026.
- BOHREN, AISLINN, J., AND D. HAUSER (2018): “Social Learning with Model Misspecification: A Framework and a Robustness Result,” *mimeo*.
- BOHREN, J. A. (2016): “Informational herding with model misspecification,” *Journal of Economic Theory*, 163, 222–247.
- BREZA, E., A. G. CHANDRASEKHAR, AND A. TAHBAZ-SALEHI (2016): “Seeing the forest for the trees? an investigation of network knowledge,” *working paper*.
- CAI, H., Y. CHEN, AND H. FANG (2009): “Observational Learning: Evidence from a Randomized Natural Field Experiment,” *American Economic Review*, 99(3), 864–882.
- CASCIARO, T. (1998): “Seeing things clearly: social structure, personality, and accuracy in social network perception,” *Social Networks*, 20(4), 331–351.
- CELEN, B., AND S. KARIV (2004): “Distinguishing Informational Cascades from Herd Behavior in the Laborator,” *American Economic Review*, 94, 484–498.

- CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2012): “Testing Models of Social Learning on Networks: Evidence from a Framed Field Experiment,” mimeo.
- CONLEY, T., AND C. UDRY (2001): “Social Learning through Networks: The Adoption of New Agricultural Technologies in Ghana,” *American Journal of Agricultural Economics*, 83(3), 668–673.
- DASARATHA, K., AND K. HE (2017): “Network Structure and Naive Sequential Learning,” *working paper*.
- DEGROOT, M. H. (1974): “Reaching a Consensus,” *Journal of the American Statistical Association*, 69(345), 118–121.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Uni-Dimensional Opinions,” *Quarterly Journal of Economics*, 118(3), 909–968.
- DUFLO, E., AND E. SAEZ (2002): “Participation and investment decisions in a retirement plan: The influence of colleagues choices,” *Journal of Public Economics*, 85, 121–148.
- ELLISON, G., AND D. FUDENBERG (1993): “Rules of Thumb for Social Learning,” *Journal of Political Economy*, 101(4), 612–643.
- (1995): “Word-of-Mouth Communication and Social learning,” *Quarterly Journal of Economics*, 110(1), 93–125.
- ENKE, B., AND F. ZIMMERMANN (2015): “Correlation Neglect in Belief Formation,” mimeo.
- EYSTER, E., AND M. RABIN (2005): “Cursed equilibrium,” *Econometrica*, 73(5), 1623–1672.
- (2014): “Extensive imitation is irrational and harmful,” *The Quarterly Journal of Economics*, 129(4), 1861–1898.
- FAINMESSER, I. P., AND D. A. GOLDBERG (2016): “Cooperation in partly observable networked markets,” *working paper*.
- GALE, D., AND S. KARIV (2003): “Bayesian learning in social networks,” *Games and Economic Behavior*, 45(2), 329–346.
- GOLOSOV, M., N. QIAN, AND Y. KAI (2015): “Understanding the Influence of Government Controlled Media: Evidence from Air Pollution in China,” *working paper*.
- GOLUB, B., AND M. O. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2(1), 112–49.
- GOLUB, B., AND E. SADLER (2017): “Learning in social networks,” .
- GRIMM, V., AND F. MENGEL (2014): “Experiments on Belief Formation in Networks,” mimeo.
- HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2014): “When more information reduces the speed of learning,” *working paper*.

- JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (2012): “Social Capital and Social Quilts: Network Patterns of Favor Exchange,” *American Economic Review*, 102(5), 1857–1897.
- JADBABAIE, A., P. MOLAVI, A. SANDRONI, AND A. TAHBAZ-SALEHI (2012): “Non-Bayesian social learning,” *Games and Economic Behavior*, 76(1), 210–225.
- KRACKHARDT, D. (1990): “Assessing the Political Landscape: Structure, Cognition, and Power in Organizations,” *Administrative Science Quarterly*, 35(2), 342–369.
- LEE, I.-H. (1993): “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 61, 395–411.
- LEVY, G., AND R. RAZIN (2016): “Information diffusion in networks with the Bayesian Peer Influence heuristic,” *working paper*.
- LIPNOWSKI, E., AND E. SADLER (2018): “Peer-confirming equilibrium,” *Econometrica*.
- MOBIUS, M., T. PHAN, AND A. SZEIDL (2015): “Treasure Hunt: Social Learning in the Field,” *NBER Working Paper No. 21014*.
- MOBIUS, M., AND T. S. ROSENBLAT (2001): “The Process of Ghetto Formation: Evidence from Chicago,” mimeo.
- MOLAVI, P., A. TAHBAZ-SALEHI, AND A. JADBABAIE (2016): “Foundations of non-bayesian social learning,” *working paper*.
- MOSSEL, E., N. OLSMAN, AND O. TAMUZ (2016): “Efficient Bayesian Learning in Social Networks with Gaussian Estimators,” *working paper*.
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): “Strategic Learning and the Topology of Social Networks,” *Econometrica*, 83(5), 1755–1794.
- MUELLER-FRANK, M. (2013): “A general framework for rational learning in social networks,” *Theoretical Economics*, 8(1), 1–40.
- MUELLER-FRANK, M., AND C. NERI (2017): “Quasi-Bayesian Updating in Social Networks,” *working paper*.
- SMITH, L., AND P. SORENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68(2), 371–398.
- VOSOUGHI, S., D. ROY, AND S. ARAL (2018): “The spread of true and false news online,” *Science*, 359(6380), 1146–1151.

A Appendix: An extension and proofs

A.1 An equivalent learning procedure without degenerate estimates

In our learning procedure described in Section 3, agent i forms higher-order estimates for each clique $\{i, j, \dots, k\}$ within her local network. Moreover, she sets degenerate estimates

when the last agent is a repeated agent, that is, for $h \in \{i, j, \dots, k\}$, she sets $\mathbf{b}_t^{ij\dots kh} = \mathbf{b}_t^{ij\dots k}$. Agent i does not set or form any estimates involving two or more repeated agents. One may wonder whether our learning procedure is with loss because we truncate the agent's higher-order estimates. We say a sequence of agents is *fully-connected* if there are at least two distinct agents and every distinct agent is connected to every other distinct agent in the sequence. Will an agent learn any differently if she forms all higher-order estimates involving fully-connected neighbors in her local network?

In this section, we show that the answer is no. We describe an equivalent learning procedure without degenerate estimates. In other words, in this learning procedure, agent i forms all higher-order estimates, including those with repeated agents. We then show the learning outcomes of the two procedures are the same, and clearly the procedure in the text economizes on computation.

At $t = 0$, agent i receives signal x_0^i .

At $t = 1$, agent i updates her estimates \mathbf{b}_1^i about the state distribution using x_0^i as before. Since there is no previous report, the initial values for all the higher-order estimates are $\mathbf{b}_1^{ij\dots l} = \mathbf{0}$ for each sequence of fully-connected (possibly repeated) agents $(ij\dots l)$. She exchanges reports with each neighbor in her local network. Then she observes her signal from nature x_1^i . Period 1 ends.

At $t \geq 2$, agent i forms her estimates and higher-order estimates in three steps just like the learning procedure in Section 3.

Step 1: Extracting new information. Agent i extracts a new signal α_{t-1}^{ij} from each neighbor j .

$$\alpha_{t-1}^{ij} = \mathbf{b}_{t-1}^j - \mathbf{b}_{t-1}^i.$$

Furthermore, α_{t-1}^{ii} is based on her signal x_{t-1}^i like before.

Step 2: Updating. Agent i then updates her period- t estimates using the signals extracted from each neighbor and from nature. By Bayes' rule:

$$\mathbf{b}_t^i = \mathbf{b}_{t-1}^i + \sum_{h \in g_i} \alpha_{t-1}^{ih}.$$

Step 3: Updating higher-order estimates. For each sequence of fully-connected agents $(ij\dots lh)$, agent i believes that agent j believes \dots that agent l extracts $\alpha_{t-1}^{ij\dots lh}$ from agent h ,

where

$$\alpha_{t-1}^{ij\dots lh} = \mathbf{b}_{t-1}^h - \mathbf{b}_{t-1}^{ij\dots lh}.$$

Then, agent i updates her estimates of agent j 's estimates of . . . of agent l 's estimates by applying Bayes' rule to (what she believes that j believes . . . are) the new signals agent l has extracted. Namely, $\alpha_{t-1}^{ij\dots lh}$ for each $h \in g_{ij\dots l}$ above. Then we have

$$\mathbf{b}_t^{ij\dots l} = \mathbf{b}_{t-1}^{ij\dots l} + \sum_{h \in g_{ij\dots l}} \alpha_{t-1}^{ij\dots lh}.$$

Next, she reports \mathbf{b}_t^i to each of her neighbors and simultaneously receives each neighbor's report \mathbf{b}_t^j . Then she observes her signal from nature x_t^i . Period t ends. ||

In the procedure above, agent i can form infinitely higher-order estimates involving her neighbors, which is a conceptual device involving a large amount of computation. We now show that under this learning procedure, agent i 's estimates for two sequences of agents are the same if these sequences include the same set of distinct agents. Therefore there is no loss in setting the degenerate estimates as we do in the main text to avoid forming estimates that involving repeated agents.

LEMMA 4. *Fix a sequence of realized signals X_T . For any sequence of fully-connected agents $(l_1 \dots l_z)$, if the set of distinct agents in the sequence is $\{i, j, \dots, k\}$, then $\mathbf{b}_t^{l_1 \dots l_z} = \mathbf{b}_t^{ij\dots k}$ for all $t \geq 1$.*

Proof of Lemma 4: We prove this lemma by induction on time t . Clearly, the induction hypothesis is true at $t = 1$. Because there is no shared report when agents form their estimates at $t = 1$, $\mathbf{b}_t^{l_1 \dots l_z} = \mathbf{0}$ for all sequences of fully-connected agents.

Next, suppose the induction hypothesis holds at time t . Then, we will show it also holds at time $t + 1$. To begin with, because $\{i, j, \dots, k\}$ is the set of distinct agents in the sequence $(l_1 \dots l_z)$, the shared local network is the same one: $g_{ij\dots k} = g_{l_1 \dots l_z}$. For any $h \in g_{ij\dots k}$, using the first part in step 3, we have

$$\alpha_t^{ij\dots kh} = \mathbf{b}_t^h(s_n) - \mathbf{b}_t^{ij\dots kh} = \mathbf{b}_t^h - \mathbf{b}_t^{l_1 \dots l_z h} = \alpha_t^{l_1 \dots l_z h}. \quad (15)$$

The second equality uses the induction hypothesis at time t . Then, using the second part in

Step 3, we have,

$$\mathbf{b}_{t+1}^{ij\dots k} = \mathbf{b}_t^{ij\dots k} + \sum_{h \in g_{ij\dots k}} \boldsymbol{\alpha}_t^{ij\dots kh} = \mathbf{b}_t^{l_1\dots l_z} + \sum_{h \in g_{l_1\dots l_z}} \boldsymbol{\alpha}_t^{l_1\dots l_z h} = \mathbf{b}_{t+1}^{l_1\dots l_z}.$$

The second equality uses both the induction hypothesis and equation (15). ■

A.2 Proofs

Proof of Observation 1: Let $B^i(z)$ be what agent i believes z is under Assumption 1. We expand the result and show the following three claims are true. For every agent $i, j \in N_i$, any $t \geq 2$ and any clique $\{i, j, \dots, k\}$:

- (i) $b_t^i(s_n) = B^i \left(\log \frac{\Pr(s_n | \{x_\tau^h\}_{\tau \leq t-2, h \in g, x_{t-1}^i})}{\Pr(s_N | \{x_\tau^h\}_{\tau \leq t-2, h \in g, x_{t-1}^i})} \right)$;
- (ii) $\mathbf{b}_t^{ij} = B^i(\mathbf{b}_t^j - \boldsymbol{\alpha}_{t-1}^{jj})$;
- (iii) $\mathbf{b}_t^{ij\dots k} = B^i(\mathbf{b}_t^{j\dots k})$.

That is, (i) agent i 's estimates at period t are what she believes to be the log-likelihood ratios of the state distribution conditional on all the signals the network receives up to period $t - 2$ plus her own signal at $t - 1$; and (ii) her estimates of neighbor j 's estimates is equal to what agent i believes to be agent j 's estimates except for his most recent signal x_{t-1}^j ; and (iii) agent i believes all her higher-order estimates are correct.

We prove this result by induction. First, agent i gets the initial signal x_0^i and the log-likelihood ratio of her Bayesian posterior is $b_1^i(s_n) = \log \Pr(s_n | x_0^i) - \log \Pr(s_N | x_0^i)$ by definition. All the higher-order estimates $\mathbf{b}_1^{ij\dots k} = \mathbf{0}$ reflect the symmetric prior because the agents have learned nothing from their neighbors.

At $t = 2$, agent i infers the signal $\boldsymbol{\alpha}_1^{ij}$ for all $j \in N_i$, which are the log-likelihood ratios of agent j 's initial signal x_0^j . By Assumption 1, agent i believes she can see the entire network, that is, she believes $g_i = B^i(g)$. Using expression (5), we have

$$b_2^i(s_n) = \log \frac{\Pr(s_n | \{x_0^h\}_{h \in g_i, x_1^i})}{\Pr(s_N | \{x_0^h\}_{h \in g_i, x_1^i})} = B^i \left(\log \frac{\Pr(s_n | \{x_0^h\}_{h \in g, x_1^i})}{\Pr(s_N | \{x_0^h\}_{h \in g, x_1^i})} \right).$$

Assumption 1 also implies that agent i believes she can see neighbor j 's entire local network,

$g_{ij} = B^i(g_j)$. Using expression (8), we have:

$$b_2^{ij}(s_n) = \log \frac{\Pr(s_n | \{x_0^h\}_{h \in g_{ij}})}{\Pr(s_N | \{x_0^h\}_{h \in g_{ij}})} = B^i \left(\log \frac{\Pr(s_n | \{x_0^h\}_{h \in g_j})}{\Pr(s_N | \{x_0^h\}_{h \in g_j})} \right).$$

This means that agent i believes any difference between $b_t^j(s_n)$ and $b_2^{ij}(s_n)$ can only be attributed to x_1^j . Moreover, for any clique $\{i, j, \dots, k\}$, agent i believes that $g_{ij\dots k} = B^i(g_{j\dots k})$. We have,

$$b_2^{ij\dots k}(s_n) = B^i \left(\log \frac{\Pr(s_n | \{x_0^h\}_{h \in g_{j\dots k}})}{\Pr(s_N | \{x_0^h\}_{h \in g_{j\dots k}})} \right) = B^i \left(b_2^{j\dots k}(s_n) \right).$$

Next, suppose the induction hypothesis is true at t . Then, we will show it also holds at time $t + 1$. At $t + 1$, according to (ii) above, agent i believes that the difference between \mathbf{b}_t^j and \mathbf{b}_t^{ij} is caused by j 's private signal x_{t-1}^j . That is, she believes that α_t^{ij} reflects x_{t-1}^j and is new information. Agent i then follows expression (5) to update her estimates. She believes that she is using Bayes' rule to incorporate all the x_{t-1}^j and her own signal x_t^i . This is exactly (i) for period $t + 1$:

$$b_{t+1}^i(s_n) = B^i \left(\log \Pr(s_n | \{x_\tau^h\}_{\tau \leq t-1, h \in g}, x_t^i) - \log \Pr(s_N | \{x_\tau^h\}_{\tau \leq t-1, h \in g}, x_t^i) \right).$$

Moreover, since agent i believes she observes all of neighbor j 's neighbors, she must update her estimates \mathbf{b}_{t+1}^{ij} according to expression (8) of our procedure. By the induction hypothesis, $\alpha_t^{ijh} = B^i(\alpha_t^{jh})$, and thus agent i believes that she can infer all the signals agent j inferred from his neighbors. The only thing missing is agent j 's own signal, which is exactly point (ii). To see this, note that

$$\mathbf{b}_{t+1}^{ij} = \mathbf{b}_t^{ij} + \sum_{h \in g_{ij}} \alpha_t^{ijh} = \mathbf{b}_t^j + B^i \left(\sum_{h \in g_j \setminus \{j\}} \alpha_t^{jh} \right)$$

while

$$B^i(\mathbf{b}_{t+1}^j) = \mathbf{b}_t^j + B^i \left(\sum_{h \in g_j} \alpha_t^{jh} \right).$$

Thus, $\mathbf{b}_{t+1}^{ij} = B^i(\mathbf{b}_{t+1}^j - \alpha_t^{jj})$, and (ii) holds. Similarly, because $\alpha_t^{ij\dots kh} = B^i(\alpha_t^{j\dots kh})$, we have $\mathbf{b}_{t+1}^{ij\dots k} = B^i(\mathbf{b}_{t+1}^{j\dots k})$ for any clique $\{i, j, \dots, k\}$, and (iii) holds. ■

Proof of Lemma 1: Recall the definition of the disjoint sets (X_t^μ, X_t^ν) . For each agent i , let $\{x_t^{\mu,i}, x_t^{\nu,i}\} = \{x_t^i, x_t^\emptyset\}$. That is, agent i is uninformed in one and learns x_t^i in the other. In addition to equations (10) and (11) in the lemma, we claim that for any clique, $\{i, j, \dots, k\}$ and $t \geq 1$,

$$\mathbf{b}_t^{ij\dots k} = \mathbf{b}_t^{\mu,ij\dots k} + \mathbf{b}_t^{\nu,ij\dots k}. \quad (16)$$

We now prove all three equations by induction on time t .

By the definition of $\{x_0^{\mu,i}, x_0^{\nu,i}\}$, we have $\{\mathbf{b}_1^{\mu,i}, \mathbf{b}_1^{\nu,i}\} = \{\mathbf{b}_1^i, \mathbf{0}\}$. Also, all the higher-order estimates are $\mathbf{0}$ by definition since there has been no previous report. Thus equations (10), (11) and (16) hold at $t = 1$.

Next, suppose equations (10), (11) and (16) hold at time t . We now show they also hold at time $t + 1$. Recall that agent i 's inferred signals under X_t^μ and X_t^ν are respectively

$$\alpha_t^{\mu,ij} = \mathbf{b}_t^{\mu,j} - \mathbf{b}_t^{\mu,ij}, \text{ and } \alpha_t^{\nu,ij} = \mathbf{b}_t^{\nu,j} - \mathbf{b}_t^{\nu,ij}.$$

Further, by the induction hypothesis, from (10) and (11), we have:

$$\alpha_t^{ij} = \mathbf{b}_t^j - \mathbf{b}_t^{ij} = (\mathbf{b}_t^{\mu,j} + \mathbf{b}_t^{\nu,j}) - (\mathbf{b}_t^{\mu,ij} + \mathbf{b}_t^{\nu,ij}) = \alpha_t^{\mu,ij} + \alpha_t^{\nu,ij}. \quad (17)$$

Since $\{x_t^{\mu,i}, x_t^{\nu,i}\} = \{x_t^i, x_t^\emptyset\}$, $\{\alpha_t^{\mu,ii}, \alpha_t^{\nu,ii}\} = \{\alpha_t^{ii}, \mathbf{0}\}$ which implies $\alpha_t^{ii} = \alpha_t^{\mu,ii} + \alpha_t^{\nu,ii}$.

$$\mathbf{b}_{t+1}^i = \mathbf{b}_t^i + \sum_{h \in g_i} \alpha_t^{ih} = \mathbf{b}_t^{\mu,i} + \mathbf{b}_t^{\nu,i} + \sum_{h \in g_i} (\alpha_t^{\mu,ih} + \alpha_t^{\nu,ih}) = \mathbf{b}_{t+1}^{\mu,i} + \mathbf{b}_{t+1}^{\nu,i}.$$

The second equality holds by (10) and (17), and the last equality holds because it is expression (5) of the learning procedure under X_t^μ and X_t^ν respectively. Thus, (10) holds at time $t + 1$. Moreover, all the new information agent i believes one neighbor has learned from another under X_t can be expressed as the sum of the corresponding new information under X_t^μ and X_t^ν similar to equation (17). Specifically,

$$\alpha_t^{ijh} = \alpha_t^{\mu,ijh} + \alpha_t^{\nu,ijh} \text{ and } \alpha_t^{ij\dots kh} = \alpha_t^{\mu,ij\dots kh} + \alpha_t^{\nu,ij\dots kh}.$$

Then we can show that:

$$\mathbf{b}_{t+1}^{ij} = \mathbf{b}_t^{ij} + \sum_{h \in g_{ij}} \alpha_t^{ijh} = \mathbf{b}_t^{\mu,ij} + \mathbf{b}_t^{\nu,ij} + \sum_{h \in g_{ij}} (\alpha_t^{\mu,ijh} + \alpha_t^{\nu,ijh}) = \mathbf{b}_{t+1}^{\mu,ij} + \mathbf{b}_{t+1}^{\nu,ij}.$$

In a similar way, we can show for all cliques $\{i, j, \dots, k\}$, $\mathbf{b}_{t+1}^{ij\dots k} = \mathbf{b}_{t+1}^{\mu, ij\dots k} + \mathbf{b}_{t+1}^{\nu, ij\dots k}$. Thus (11) and (16) also hold at time $t + 1$. ■

Proof of Lemma 2: By definition, for any $t \geq 2$,

$$\begin{aligned} \alpha_t^{ij} = \mathbf{b}_t^j - \mathbf{b}_t^{ij} &= \left(\mathbf{b}_{t-1}^j + \sum_{k \in g_j} \alpha_{t-1}^{jk} \right) - \left(\mathbf{b}_{t-1}^{ij} + \sum_{h \in g_{ij}} \alpha_{t-1}^{ijh} \right) \\ &= \left(\mathbf{b}_{t-1}^j + \sum_{k \in g_j} \alpha_{t-1}^{jk} \right) - \left(\mathbf{b}_{t-1}^j + \sum_{h \in g_{ij} \setminus \{j\}} \alpha_{t-1}^{ijh} \right) \\ &= \sum_{l \in (g_j \setminus g_i) \cup \{j\}} \alpha_{t-1}^{jl} + \sum_{h \in g_{ij} \setminus \{j\}} \left(\alpha_{t-1}^{jh} - \alpha_{t-1}^{ijh} \right). \end{aligned}$$

The first term concerns what agent j learns from his neighbors (and nature) who are not connected to agent i . The second term concerns their common neighbors. ■

Proof of Proposition 1: We begin with two properties of social quilts. First, if $d(ih) = d$, then there must be a unique path of length d from i to h . Suppose instead, there are two such distinct paths between them. Let these two paths be $(ii_1i_2 \dots i_{d-1}j)$ and $(ij_1j_2 \dots j_{d-1}j)$, with $i = i_0 = j_0$ and $j = i_d = j_d$. Then there must exist parts of the two paths that differ, that is there must exist two numbers k and h , $0 \leq k < h \leq d$ and $h - k \geq 2$ such that $i_k = j_k$, and $i_h = j_h$, but $i_l \neq j_l$ if $k < l < h$. Clearly, $(i_k i_{k+1} \dots i_h j_{h-1} \dots j_{k+1})$ must be a circle, going from i_k to herself through distinct agents. The agents are distinct because by assumption $i_l \neq j_l$ for any $l \in (k, h)$, and since $d(ii_l) = l$ and $d(ij_{l'}) = l'$, $i_l \neq j_{l'}$ whenever $l \neq l'$. In a social quilt, any two agents in a circle are linked. Thus agent i_k and i_h must be linked, but this contradicts $(ii_1i_2 \dots i_{d-1}j)$ being a shortest path.

The second property of social quilts is that if agent i 's signal travels from agent l to k , and then inferred by k 's neighbor h who is not linked to l , then h must be further away from i . Specifically, if l is the agent before k on the shortest path from i to k , such that $d(ik) = d(il) + 1$ and $kl \in G$, then for any h with $hk \in G$ and $hl \notin G$, the shortest path from i to h must go through l and k : $d(ih) = d(ik) + 1$. To see this, note that since $hk \in G$, the maximum possible distance between i and h is $d(ih) \leq d(ik) + 1$. Next, if $d(ih) \leq d(ik) - 1$, then the path through l cannot be the unique shortest path between i and k . If $d(ih) = d(ik)$, then the shortest path between i and h must not involve k , or agent l since $hl \notin G$. Thus we have a circle involving $\{h, k, l\}$ and i 's shortest path to agent h and l , which is a contradiction to the definition of social quilts. Therefore $d(ih) = d(ik) + 1$.

We now proceed to prove the proposition. By Lemma 1, if we can show that agents' learning outcomes are strongly efficient for each signal, then it is also true for multiple signals. Without loss of generality, let agent i receive an initial signal $x_0^i = x^{i,m}$. Let $\varphi_{nm}^i = \log \Pr(s_n | x^{i,m}) - \log \Pr(s_N | x^{i,m})$ for each $s_n \in S$, and $\boldsymbol{\varphi}^{i,m} = (\varphi_{1m}^i, \dots, \varphi_{Nm}^i)$. By the first property, there is a unique shortest path from i to each agent h . That is, there is a unique neighbor k of h who is on h 's shortest path to i . We want to show that agent h infers the signal at $t = d(ih) + 1$ from this neighbor k (who can be agent i), and this is the only signal agent h infers from his neighbors at any time. Specifically, for any $k' \in N_h$ and any time t , $\boldsymbol{\alpha}_t^{hk'} = \boldsymbol{\varphi}^{i,m}$ if and only if $t = d(ik) + 1 = d(ih)$. Otherwise, $\boldsymbol{\alpha}_t^{hk'} = \mathbf{0}$. Notice that this implies agent h learns the signal and changes his estimates once at $t = d(ih) + 1$.

We prove this claim by induction on time t . First, this holds at $t = 2$. If $d(ih) = 1$, or $h \in N_i$, then agent h infers the signal from agent i 's report \mathbf{b}_1^i such that $\boldsymbol{\alpha}_1^{hi} = \boldsymbol{\varphi}^{i,m}$. No other agents (including agent i) infer any new signal from their neighbors. Next, if $\boldsymbol{\alpha}_1^{hk} = \boldsymbol{\varphi}^{i,m}$, then clearly $k = i$ and $d(ik) = 0$, $d(ih) = 1$.

Next, suppose this holds at time t , we show it also holds at time $t+1$. First, if $\boldsymbol{\alpha}_t^{hk} = \boldsymbol{\varphi}^{i,m}$ at time $t+1$, then using the iterative relationship between inferred signals in equation (12) and the fact that the second term is zero by Lemma 3 (which is proved next), we have

$$\boldsymbol{\alpha}_t^{hk} = \sum_{l \in (g_k \setminus g_h) \cup \{k\}} \boldsymbol{\alpha}_{t-1}^{kl}.$$

That is, agent k must infer the signal from someone (say l) outside g_h in the previous period, so $hl \notin G$. By the induction hypothesis, since $\boldsymbol{\alpha}_{t-1}^{kl} = \boldsymbol{\varphi}^{i,m}$, we have $d(ik) = t - 1$ and $d(il) = t - 2$. By the second property above, it must be true that $d(ih) = t$. Second, if $d(ih) = t$ and $d(ik) = t - 1$, by the induction hypothesis $\boldsymbol{\alpha}_{t-1}^{kl} = \boldsymbol{\varphi}^{i,m}$ for some neighbor l . Because $d(il) = t - 2$ and $d(ih) = t$, l is not connected to h , $l \in g_k \setminus g_h$. Since agent h has not learned any new information so far, $\boldsymbol{\alpha}_t^{hk} = \boldsymbol{\varphi}^{i,m}$. Thus $\boldsymbol{\alpha}_t^{hk} = \boldsymbol{\varphi}^{i,m}$ if and only if $d(ih) = t$ and $d(ik) = t - 1$. Since agent h incorporates signal x_0^i exactly once at period $d(ih) + 1$, $\mathbf{b}_t^h = \boldsymbol{\varphi}^{i,m}$ if $t > d(ih)$ and $\mathbf{b}_t^h = \mathbf{0}$ otherwise. Thus the learning outcomes are strongly efficient with signal x_0^i .

Lastly, if the network is not a social quilt, there exists some sequence of realized signals such that at least one agent's learning outcomes are not strongly efficient. We first characterize a social quilt as follows.

LEMMA 5. *Network (g, G) is a social quilt if and only if it contains no simple circle and satisfies LCS.*

Proof: For necessity, if a network is a social quilt, it does not contain a simple circle by definition. Moreover, (g_i, G_i) satisfies LCS_i because for any $j \in N_i$, if there exist agents k and k' such that $k, k' \in N_i \cap N_j$, then $(kik'j)$ must be a circle. In a social quilt, $kk' \in G$, and thus every agent i 's local network satisfies LCS_i .

For sufficiency, we show by induction that if the network satisfies LCS and contains no simple circle, then any circle of more than three agents must be a clique. First, any four agent circle must be part of a clique. No simple circle means that there must be at least one link between two nonadjacent agents. Since the network satisfies LCS, all four agents must be a clique. Next, suppose any circle of $l \geq 4$ agents is part of a clique. Consider a circle of $l + 1$ agents. Because it is not a simple circle, there exists at least one link between two nonadjacent agents ij . The original circle is now divided into two smaller circles of no more than l agents, and thus each must be a clique by the induction hypothesis. In addition, any pair of agents, one from each smaller circle, are common neighbors of i and j . Because agent i 's local network satisfies LCS_i , they are connected. Therefore this circle of $l + 1$ agent must be a clique, which is the definition of a social quilt. Next, if the network satisfies LCS and there is no circle, then the network is a tree and thus also a social quilt. Finally, in a circle of three agents, a triangle, clearly all three agents are connected. ■

Lemma 5 shows that when a network is not a social quilt, it must either contain simple circles or violate LCS. We show in Proposition 2 and 3 that both lead to learning errors. ■

Proof of Lemma 3: First, $\mathbf{b}_t^{ij} = \mathbf{b}_t^{ij}$ is immediate from Lemma 4 in Appendix A.1 because they are estimates involving the same distinct agents.

We now prove the second part of the lemma by induction on time t . At $t = 1$, all the higher-order estimates reflect the symmetric prior because there has been no previous reports. So for any clique $\{i, j, l, \dots, k\}$, $\mathbf{b}_1^{ij} = \mathbf{b}_1^{il} = \mathbf{b}_1^{ijl} = \dots = \mathbf{b}_1^{ijl\dots k} = \mathbf{0}$.

Next, suppose this is true at time t , we want to show it also holds at time $t + 1$. Notice that by LCS_i , g_{ij} is a clique, implying $g_{ij} = g_{il}$ for all l such that agent $\{i, j, l\}$ form a triangle. By the induction hypothesis, for any $h \in g_{ij} = g_{il}$,

$$\alpha_t^{ijh} = \mathbf{b}_t^h - \mathbf{b}_t^{ijh} = \mathbf{b}_t^h - \mathbf{b}_t^{ilh} = \alpha_t^{ilh}.$$

Then, using expression (8), we have:

$$\mathbf{b}_{t+1}^{ij} = \mathbf{b}_t^{ij} + \sum_{h \in g_{ij}} \alpha_t^{ijh} = \mathbf{b}_t^{il} + \sum_{h \in g_{il}} \alpha_t^{ilh} = \mathbf{b}_{t+1}^{il}.$$

Similarly, since g_{ij} is a clique, $g_{ij} = g_{ijl\dots k}$ for all cliques $\{i, j, l, \dots, k\}$ containing i and j . By the induction hypothesis, for any $h \in g_{ij} = g_{ijl\dots k}$,

$$\alpha_t^{ijh} = \mathbf{b}_t^h - \mathbf{b}_t^{ijh} = \mathbf{b}_t^h - \mathbf{b}_t^{ijl\dots kh} = \alpha_t^{ijl\dots kh}.$$

Then, using expression (8) and (9),

$$\mathbf{b}_{t+1}^{ij} = \mathbf{b}_t^{ij} + \sum_{h \in g_{ij}} \alpha_t^{ijh} = \mathbf{b}_t^{ijl\dots k} + \sum_{h \in g_{ijl\dots k}} \alpha_t^{ijl\dots kh} = \mathbf{b}_{t+1}^{ijl\dots k}.$$

Thus, $\mathbf{b}_{t+1}^{ij} = \mathbf{b}_{t+1}^{il} = \mathbf{b}_{t+1}^{ijl} = \dots = \mathbf{b}_{t+1}^{ijl\dots k}$. ■

Proof of Proposition 2: For Part 1, by our definition of efficient learning, it suffices to show that the agents' learning outcomes are not efficient for some sequence of realized signals X_T . We now show this is the case if the network receives only one initial informative signal. We begin with the repetition of one signal, $x_0^i = x^{i,m}$, within a simple circle. For any k -agent simple circle $C = (i_1 \dots i_k)$, there are two cases: agent $i \in C$ or $i \notin C$. First, suppose that $i \in C$ and without loss, let $i = i_k$. Then at $t = 2$, agent i_1 and i_{k-1} 's inferred signals are $\alpha_1^{i_1 i} = \alpha_1^{i_{k-1} i} = \alpha_0^{ii}$. Recall that LCS holds, and thus the second term of the iterative relationship between inferred signals in equation (12) is zero. Also, by assumption, $\alpha_t^{ll} = 0$ for any $t > 0, l \in g$. Then equation (12) can be rewritten as

$$\alpha_{t+1}^{jh} = \sum_{l \in g_h \setminus g_j} \alpha_t^{hl}. \quad (18)$$

At period $t = k + 1$, the signal finishes traveling around the simple circle in both directions, and thus $\alpha_k^{i_{k-1} i} = \alpha_0^{ii}$ and $\alpha_k^{i_1 i} = \alpha_0^{ii}$. At this point, agent i learns a total of three copies of her original signal and everyone else in the simple circle learns two copies. From now on agent i and all other agents in the simple circle infer two copies of x_0^i in every k periods.

Next, if $i \notin C$, then the first time this signal arrives at the circle, it must reach either only one agent (say i_k), or two linked agents (say i_k and i_1 learn from their common neighbor). To see this, suppose to the contrary, i_k and i_l learn the signal at the same time, but either $l \neq 1, k - 1$; or i_l learns from a different source. Then there is another simple circle inside the path from i to i_k , i_k to i_l through C , and i_l to i . It contradicts the assumption that C is the only simple circle. Moreover, once the signal reaches the circle, agents in C do not infer any other new signal from outside C , because there is no other simple circle through which information can travel back. Without loss of generality, assume i_k (and i_1) learns the signal

from some agent j (who could be i) outside the simple circle, such that $\alpha_t^{i_k j} = \alpha_0^{ii}$ for some $j \in N_{i_k}$. Because i_1 and i_{k-1} are not linked by definition of a simple circle and (g_{i_k}, G_{i_k}) is assumed to satisfy LCS_{i_k} , j cannot be linked with i_{k-1} . Then $\alpha_{t+1}^{i_{k-1} i_k} = \alpha_0^{ii}$, and it is passed on to i_{k-2} and so on. Also, the signal travels through i_1 to i_2 , because i_1 learns from either j or i_k . Similar to the first case, we can show agent i_k and all other agents in the simple circle infer two more copies of x_0^i every k periods. Recall D is the *diameter* of network. These newly inferred signals will travel to all the other agents outside the simple circle in at most D periods. Clearly all agents believe in the state most likely given signal x_0^i as $t \rightarrow \infty$. Therefore, the agents' learning outcomes are not efficient.

Similarly, in a network with multiple simple circles, we can show that the agents' estimates are wrong when there is one initial informative signal. Let k be the number of agents in the largest simple circle. For any $z \in \mathbf{R}$, $\lceil z \rceil$ is the smallest integer that is greater or equal to z . Then simple algebra can show that at any $t \in [\tau(D + \lceil \kappa/2 \rceil) + 1, (\tau + 1)(D + \lceil \kappa/2 \rceil)]$, any agent l in a simple circle believes there are at least two copies of x_0^i if $\tau = 1$; and at least

$$2\tau + 2 \sum_{\tau'=1}^{\tau-1} (2(\kappa_{sc} - 1))^{\tau'} \quad (19)$$

copies of signal x_0^i if τ is an integer larger than 1. The first part captures the signal repetition in one simple circle, and the second part shows that agents in one simple circle keep inferring more and more new signals from all the other $\kappa_{sc} - 1$ simple circles, and passing their own repeatedly inferred signals to them. As $t \rightarrow \infty$, each agent believes in the state most likely given x_0^i while the Bayesian posterior is bounded away from 0 and 1.

For Part 2 of the result, we begin with a network with one simple circle. Specifically, to study asymptotic efficiency, we consider the case with a finite number of informative signals ($T_i < \infty$ for all $i \in g$), and then let each T_i go to infinity. Recall that $T = \max_{i \in g} T_i$. When T is finite, at time $t = T + D$, all signals must have reached the simple circle. Let $\eta_{T+D}^{i_k}(x_t^l)$ be the number of copies of signal x_t^l agent i_k believes in at time $T + D$, then:

$$\mathbf{b}_{T+D}^{i_k} = \sum_{l \in g, t \leq T} (\eta_{T+D}^{i_k}(x_t^l) \cdot \alpha_t^{ll}). \quad (20)$$

As before, in every k periods, agent i_k must receive two more copies of each signal due to

the repetition in the simple circle, such that for any integer o ,

$$\mathbf{b}_{T+D+o\kappa}^{i\kappa} = \sum_{l \in g, t \leq T} ((\eta_{T+D}^{i\kappa}(x_t^l) + 2o) \cdot \boldsymbol{\alpha}_t^l). \quad (21)$$

Suppose that there are a subset of states S^* such that $S^* = \arg \max_{s_n \in S} \Pr(s_n | X_T)$. For any given T , as $o \rightarrow \infty$, the agents believe that only the states in S^* can be the true state, and $b_{T+D+o\kappa}^{i\kappa}(s') \rightarrow 0$ if $s' \notin S^*$. The case is similar for any other t between $T + D + o\kappa$ and $T + D + (o + 1)\kappa$ and any other agent in the network. Thus, all agents believe the true state is some states in S^* with probability arbitrarily close to 1 as $t \rightarrow \infty$. When each agent in the network receives an infinite number of signals, by the Law of Large Numbers, $s^* = \arg \max_{s_n \in S} \Pr(s_n | X_T)$ is the true state as $T_i = \infty$ for all $i \in g$. This is due to our assumption that collectively, the signals can tell every pair of states apart. Therefore the agents' estimates are correct in the limit.

When the network has multiple simple circles, 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 $s = s^*$. By assumption, there exists a possible signal $x^{i,m}$ belonging to some agent i such that for some other state $s' \neq s^*$, $s' = \arg \max_{s_n} \Pr(s_n | x^{i,m})$. Denote this signal as x' . Clearly, $\Pr(s' | x') > \Pr(s^* | x')$. Consider the following sequence of finite signals. First, let nature inject signal x' to agent i in every period from $t = 0$ to $t = k$. Recall that the largest simple circle has k agents. This insures that starting from some finite time, each simple circle receives new copies of x' from every other simple circle in every period. In particular, k consecutive signals is sufficient for non-stop signal transmission among simple circles because we have shown in part 1 of the proof that each signal comes back to each agent in the simple circle every k periods. Second, from $t = k + 1$ to period $t = t^*$ (which to be determined next), there are no informative signals. This interval of periods allows each signal x' to reach every other simple circle and travels back to the initial simple circle. Two steps to determine t^* . 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

$$\frac{\Pr(s' | k' \text{ copies of } x')}{\Pr(s^* | k' \text{ copies of } x')} \geq \frac{\Pr(s^* | x^*)}{\Pr(s' | x^*)}, \text{ where } x^* \equiv \arg \max_{x \in X} \frac{\Pr(s^* | x)}{\Pr(s' | x)}. \quad (22)$$

To avoid carrying this likelihood ratio for the rest of the proof, for any signal x (or set of

signals), we introduce

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

In the second step, we require that in each period from period $t^* - k$, the repetition must be strong enough such that every signal one simple circle infers from any other simple circle includes at least $(2k + D + 1)Ik'$ copies of x' (excluding other later exogenous signals). Here $I = |g|$ is the number of agents in the network. We let this start from period $t^* - k$ so that by period t^* , everyone in each simple circle has inferred such strong signals.

Next, we claim that regardless of the signals agents receive from nature after period t^* , all agents believe s' is increasingly more than s^* over time. That is, $\lim_{t \rightarrow \infty} b_t^h(s') - b_t^h(s^*) = \infty$ for all $h \in g$. We consider the signal one simple circle (for instance the largest one, $C = (i_1 i_2 \dots i_k)$) infers from another simple circle. Without loss, suppose the signal is learned by agent i_1 from her neighbor j who has only one link to C (more links only make it easier to dominate the later signals). By design, for $t \geq t^*$, from agent i_1 's perspective, use the second result from Observation 1, we have:

$$\begin{aligned} \alpha_t^{i_1 j}(s') - \alpha_t^{i_1 j}(s^*) &= B^{i_1} (\alpha_t^{j j}(s') - \alpha_t^{j j}(s^*)) \\ &= B^{i_1} (\log \Pr(s' | x_t^j) - \log \Pr(s^* | x_t^j)) \\ &\geq b(s', s^* | (2k + D + 1)Ik' \text{ copies of } x'). \end{aligned} \quad (23)$$

That is, the signal i_i infers from j should favor s' over s^* at least as much as $(2k + D + 1)Ik'$ copies of x' since period t^* (excluding other later exogenous signals). This reflects the fact that agent j has inferred at least $(2k + D + 1)Ik'$ copies of x' from a neighbor not connected to i_1 .

Next, $\alpha_t^{i_1 j}$ travels around the simple circle C clockwise and counterclockwise, and each time it overwhelms the exogenous signal(s) from the agent it reaches along the simple circle. Formally, in the next period, using equation (12), agent i_2 infers $\alpha_{t+1}^{i_2 i_1}$ from agent i_1 , such that

$$\alpha_{t+1}^{i_2 i_1}(s') - \alpha_{t+1}^{i_2 i_1}(s^*) \geq b(s', s^* | ((2k + D + 1)Ik' - Ik') \text{ copies of } x').$$

This is because agent i_1 gets fewer than I exogenous signals from the nature and from her neighbors outside the simple circle in each period. Moreover, each of these new exogenous signals can offset a maximum of k' copies of signal x' by the definition of k' in equation (22). The same is true for agents i_3, i_4, \dots, i_k at period $t+3, \dots, t+k$. By period $t+k+1$, agent i_k

and i_2 each must pass on a signal to i_1 . Note that $(2k + D + 1)Ik' - kIk' = (k + D + 1)Ik'$, and thus

$$\alpha_{t+k}^{i_1 i_k}(s') - \alpha_{t+k}^{i_1 i_k}(s^*) \geq b(s', s^* \mid (k + D + 1)Ik' \text{ copies of } x').$$

And the same is true for $\alpha_{t+k}^{i_1 i_2}$. Use equation (12) again for the next period, we have

$$\alpha_{t+k+1}^{j i_1}(s') - \alpha_{t+k+1}^{j i_1}(s^*) \geq b(s', s^* \mid (2k + 2D + 1)Ik' \text{ copies of } x').$$

That is, the signal agent j infers from agent i_1 includes $\alpha_{t+k}^{i_1 i_k}$ and $\alpha_{t+k}^{i_1 i_2}$ (net of the exogenous signals reaching agent i_1 in time $t + k$). Then this signal $\alpha_{t+k+1}^{j i_1}$ travels to all the other agents in the network. For example, it reaches agent l_1 at simple circle $C' = (l_1 \dots l_z)$ from agent h at time τ . Since the travel takes at most D periods, the strength of the signal favoring s' over s^* is reduced by at most DIk' copies of x' , so

$$\alpha_{\tau}^{l_1 h}(s') - \alpha_{\tau}^{l_1 h}(s^*) \geq b(s', s^* \mid (2k + D + 1)Ik' \text{ copies of } x').$$

This shows that the initial condition about the signal one simple circle infers from outside that simple circle (expression (23)) persists regardless of the exogenous signals reaching the network after period t^* . Therefore the process we described above will last forever. Because in each period each inferred signal increases the likelihood of state s' over that of s^* , all agents believe s^* is not the true state with probability arbitrarily close to 1 as $t \rightarrow \infty$.

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 all agents believe in s' with probability arbitrarily close to 1 as $t \rightarrow \infty$. 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 3: Since there exists some agent, whose local network does not satisfy LCS. We consider a neighbor of this agent, and denote this neighbor as agent l . Suppose agent l receives $x_0^l = x^{l,m}$, which is the only informative signal. We can classify all agents based on their distance to l , that is, $N_l^d = \{h \in g : d(lh) = d\}$, and $N_l^1 = N_l$. To begin with, we claim that if agent a and $b \in N_l^d$ are both linked to some agent h in N_l^{d+1} , then $ab \in G$. To see why, find a 's connection to some agent f in N_l^{d-1} , then agent f and h are not linked, because their distance must be 2. Similarly the agent who is linked to b in N_l^{d-1} , say f' , cannot be linked to h . If agent a and b are not linked, then there exists a simple circle consisting of agent f, a, h and b (with possibly other agents like f' and l), which is a

contradiction.

We first show a general feature of learning in networks without simple circles: agents in N_l^d never infer new signals from their neighbors in N_l^{d+1} . Suppose to the contrary, the first time some agent infers from her successor is agent a in N_l^d infers a new signal from h in N_l^{d+1} . Notice that in the previous period, h does not infer new signal from her successors, so the new signal a infers must come from h 's neighbors in either N_l^d or N_l^{d+1} . Suppose that the new information a infers comes from some b in N_l^d to h then to a , then by the first claim, a is linked to all h 's neighbors in N_l^d . Thus a knows all the information h learns from agents in N_l^d , contradicting the fact that a infers new information from h . The other possibility is that the new information a infers comes from agent h' in N_l^{d+1} , which reaches h and then to a . Then ah' must not be linked, because otherwise a can learn directly from h' , contradicting the assumption that a infers from h is the first time any agent learns from a successor. There are again several cases. The first one is agent h' has learned the new information from b in N_l^d . To make sure no simple circle exists, bh must be linked, so h would have learned it at the same time as h' from b . So we are back to the first possibility where the new information goes from b to h then to a , which is impossible. The other case is that h' has learned the new information from another peer h'' in N_l^{d+1} , which can be ruled out using a very similar argument. Since N_l^{d+1} contains finitely many agents, we can show a cannot learn from anyone in N_l^{d+1} .

The argument above shows that agent l never learns any new information and thus her estimates remain at $b_t^l = \alpha_0^l$ (which reflects her initial signal x_0^l). Moreover, the estimates of agents in N_l must remain at α_0^l . This is because, first, they cannot infer new information from their successors. Second, for any linked agents in N_l , they learn from agent l simultaneously and expect each other to learn it. Thus they cannot infer new information from each other.

Lastly, we claim that there must exist some agent $l' \in N_l^2$, who is linked to at least two agents in N_l but does not infer new signals from his peers (those with the same distance to l as him). Therefore the estimates of agent l' oscillate and his learning outcomes do not converge. Recall that by definition, there exist $i, j \in N_l$ and $k \in N_l^2$ such that $k \in g_{ij}$. Start with this agent k who is linked to i and j , and possibly more agents in N_l . If k does not infer new signals from his peers in N_l^2 , then he must keep oscillating. Because by the claim above, agents in N_l who are linked to k must be linked with each other. So k keeps inferring multiple copies of x_0^l in odd periods, and multiple copies of the signal that offsets x_0^l in even periods for $t \geq 3$.

Suppose instead agent k infers new information from one of his peers. The first case is

that he learns from agent $h \in N_l^2$, whose new signal comes from some agent $j' \in N_l$ different from i and j . Then $j'h$ are linked, while $j'k$ are not linked. Consider the circle $(ljkhj')$, in which lk , lh and $j'k$ cannot be linked. Because there can be no simple circles, jj' and jh must be linked. Similarly, ij' and ih must be linked, otherwise there will be a simple circle $(lj'hki)$. This implies that h never infers new signals from k because h is linked to all k 's neighbors in N_l . If h does not learn new information from his peers in N_l^2 , then his estimates must oscillate.

In the second case, agent k learns new information indirectly from some peer $h' \in N_l^2$. That is, he learns new information from h' through agent h . Suppose agent h learns information from h' , who learns the information from some agent $j' \in N_l$. The arguments are similar to the case above. We can show that i , j , and j' are all linked to agent h' while kj' and hj' cannot be linked. Moreover, ih must also be linked here to avoid a simple circle, so in this case $\{i, j, h, k\}$ is a clique. In fact, $\{i, j, h, h'\}$ is also a clique. Therefore h' is linked to more agents in N_l than agent k and h . Agent h' does not learn anything from agent h , and her estimates keep oscillating if she does not learn anything from her peers. If instead, k learns new information from h'' through h and h' . and agent h'' learns the new information from some agent in N_l , then we can show he does not learn anything from agent h' and his estimates must oscillate. This is because like before, we can show agents $\{i, j, k, h\}$ is a clique, then $\{i, j, h, h'\}$ has to be a clique, $\{i, j, h', h''\}$ has to be a clique, and so on. Since there are a finite number of agents, there must be one last agent who learns new information from some agent in N_l , but who has no peer to learn from. And this agent's estimates must oscillate because he is linked to multiple agents (more than i, j) in N_l . We denote this agent in N_l^2 who does not learn from peers as agent k^* .

Next, we construct a sequence of signals X_∞ , under which the Bayesian posterior is believing in a unique state. By assumption, the signal $x^{l,m}$ uniquely favors one state, and the Bayesian posterior under an arbitrarily large number of $x^{l,m}$ is to believe this unique state is the true state with probability arbitrarily close to 1. Let nature give this signal to agent l initially and also in every even period until $T_l = \infty$. That is, $x_\tau^l = x_0^l = x^{l,m}$ for all even τ . Recall from above that each such signal x_τ^l makes some agent k^* infers multiple copies of x_0^l in odd periods, and multiple copies of the signal that offsets x_0^l in even periods for all $t \geq \tau + 3$. In total, the estimates of agent k^* never converge. In fact, the swing of his estimates increases and goes to infinity as $t \rightarrow \infty$. ■