Locally Bayesian Learning in Networks*

Wei Li
University of British Columbia

Xu Tan
University of Washington

March 30, 2019

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 learning rule to implement such *locally Bayesian learning*: each agent 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, efficient learning in finite networks

*This paper has previously been circulated under the title "Learning in Local Networks." 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, Kyungchul (Kevin) Song, Wing Suen, Joel Watson, Quan Wen, the editor and three anonymous 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, UCR, 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 information overlaps and distortions when they learn from their social networks. Failure to do so can lead to learning errors with serious 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 influence the opinions of many of her neighbors. As a consequence, she believes erroneously—and increasingly if the same information travels back to her again in the guise of stronger opinions against MMR—that MMR is dangerous. Eventually she may believe MMR causes autism despite overwhelming evidence 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 and her own signals. In particular, she tracks the changes in each neighbor's reports over time, and attributes any *unexpected* change to new, independent information. It is *local* in that each agent only knows and extracts new information within her local network, consisting of her neighbors and the links among them.[2] We show that, despite the limited network knowledge, locally Bayesian agents are capable of partialing out repeated information and forming correct beliefs in *social quilts*, networks in which any two agents in the same circle must be connected.[3] Moreover, social quilts are also necessary for the agents to learn correctly. Because our correct learning result holds for finite networks, it complements the existing literature focusing on when the Law of Large Numbers holds and

[1]The 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 the fraudulent paper and announcements from public health officials may not overturn such erroneous beliefs.

[2]Agents having limited knowledge of their network 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.

[3]A *path* is an ordered sequence of agents, and each pair of adjacent agents in the sequence are connected. A *circle* is a path going from one agent back to the same agent.

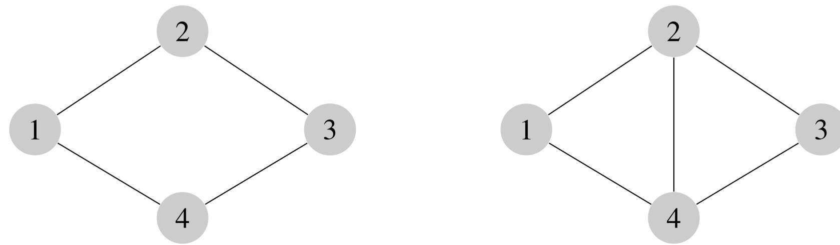

Figure 1: (a) A four-agent simple circle (b) Failure of local connection symmetry

agents can learn asymptotically in large networks.[4] In addition, the features of a social quilt are observable, and thus our result is potentially testable in the lab or in the field even with small datasets.

Because a locally Bayesian agent treats any "unexpected" change in her neighbors' reports as new, independent information, two network features we identify are crucial for their learning. First, if the network contains *simple circles*, as illustrated in Figure 1(a), then locally Bayesian agents treat correlated information as independent signals. For instance, in Figure 1(a), agent 1's information travels through a four-agent simple circle in both directions to reach agent 3. Agent 3 (not knowing agent 1) believes these two copies of the signal are independent and thus double count it. The problem is exacerbated in networks with multiple simple circles, or *echo chambers*, in which duplicate copies of each signal also travel among the simple circles repeatedly and grow exponentially. 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, similar to the MMR example above. Second, if a network fails *local connection symmetry*, then some locally Bayesian agents learn a wrong signal negatively correlated with the exogenous signal. This leads to the error of opinion swings and belief non-convergence. Local connection symmetry requires that if a pair of connected agents has two common neighbors, then these two neighbors must be connected. Figure 1(b) shows a failure of this property. In this *diamond-with-a-link* network, agent 2 and 4 know they both learn from agent 1, so they would not double count agent 1's information. But agent 3 expects them to double count because agent 3 believes their reports are independent. When they don't, agent 3 believes that they receive private signals negatively correlated with the original signal, and thus overreacts in the opposite direction.

A social quilt is characterized by two features: it contains no simple circles and it satisfies local connection symmetry. Therefore neither of the two types of learning errors mentioned

[4]It 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.

above is present. Each piece of information reaches an agent once and only once because there are no simple circles. Moreover, agents do not make local learning errors due to local connection symmetry. In short, any unexpected change in a neighbor's report is truly due to new signals in the network, and thus locally Bayesian agents learn correctly.

Our main theoretical contribution is that we retain an important feature of Bayesian learning: our agents have *perfect memory* and use it to update their beliefs by Bayes' rule. Specifically, in each period, each agent uses all her neighbors' reports from the first period to the previous period to form her belief. In contrast, several existing quasi-Bayesian learning models such as Molavi, Tahbaz-Salehi, and Jadbabaie (2018) assume each agent forms her belief in each period using only her neighbors' reports in the previous period. Intuitively, the more memory an agent uses, the fewer learning errors she makes. Modeling perfect memory makes it possible for agents to learn correctly given any finite number of signals, and thus it allows us to identify the effect of network structure on learning outcomes. Perfect memory has been understudied in the literature, possibly due to a lack of tractability. To make the model tractable, we make a crucial behavioral assumption: each agent believes her local network is the entire network (and each agent holding such belief is common knowledge). Formally, our model studies the learning outcomes of Bayesian agents who focus entirely on their local networks due to their (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. It also allows us to study boundedly rational updating when agents use all the local network information efficiently. In our view, modeling perfect memory is a necessary step toward modeling how people can avoid being misled by repeated and distorted information from their social networks, a topic under increasing scrutiny in recent years.

While locally Bayesian learning is easy to define and conceptualize, it may not be easy to analyze. Methodologically, we identify an iterative learning rule 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 belief about the state distribution, such as the probability that MMR truly leads to autism. Time is finite, and each agent receives one signal at the end of each period. From the second period onward, each agent first extracts any new information contained in her neighbors' most recent reports, which is the unexpected change mentioned above. The main innovation of our learning rule is that we identify a set of statistics—closely related to the agent's higher-order beliefs—that each agent can use to identify and to remember existing information. For example, they include her *second-order* beliefs*—her belief about each neighbor's belief in the event that the neighbor's most recent private signal is uninformative. She then compares these second-order* beliefs with a neighbor's actual report and attributes any difference to the neighbor's "new" signal. Under the behavioral assumption, she believes this new signal must be an independent signal from nature. She then incorporates all the newly extracted signals and updates her belief using Bayes' rule. This iterative learning rule is tractable and allows us to study when the agents' learning outcomes are correct and when they make learning errors.

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 Bayesian agents can learn (asymptotically) if the network is common knowledge (see Gale and Kariv (2003), Mueller-Frank (2013), Mossel, Sly, and Tamuz (2015), among others). 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] For instance, 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. A related literature in computer science studies consensus when agents use certain mechanical rules to compute the changes in opinions, say as a function of the differences between an agent and her neighbors' opinions (see Xie and Wang (2012), Yang, Meng, Dimarogonas, and Johansson (2014), among others). 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).

More closely related to our paper is the growing literature on quasi-Bayesian learning in networks. In Bala and Goyal (1998), each agent updates her belief about the optimal action rationally based on the outcomes observed 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, whereas we study how network structures affect the agents' learning outcomes. Several more recent papers feature imperfect memory in the context that is other-

[5] For 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 reduced opportunities to learn from social networks—on inner-city neighborhoods in Chicago. See Golub and Sadler (2017) for a detailed survey on the progress and challenges of learning in social networks.

[6] See 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.

wise the same as our model—agents apply Bayes' rule to all the information they believe are independent (Molavi, Tahbaz-Salehi, and Jadbabaie (2018); Mueller-Frank and Neri (2017); Levy and Razin (2018)). The underlying assumption of these models as shown by Molavi, Tahbaz-Salehi, and Jadbabaie (2018) is that in each period, each agent treats a neighbor's most recent report as a sufficient statistic for 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 can account for correlations of information locally. Thus, the learning errors of locally Bayesian agents (if any) are driven by their failures to learn about the entire network. As a result, these agents' learning outcomes, including their learning errors, have a clean relationship with the network structure. In contrast, the learning errors in quasi-Bayesian learning are primarily driven by the agents' imperfect memory.

Our paper is also related to the social learning literature in which each agent takes one and only one action sequentially.[8] In the context of misspecified beliefs, Eyster and Rabin (2010) assume that each agent believes each of her predecessors chooses an action by following his own private signal, even though her predecessors learn from their own predecessors in reality. Eyster and Rabin (2014) point out that rational agents should anti-imitate some predecessors to remove repeated information. But if agents fail to account for the redundancy in their predecessors' actions, they imitate too much. 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 learning and information aggregation in the lab and in the

[7] More specifically, Mueller-Frank and Neri (2017) assume the agent treats each neighbor' action as if it depends only on that neighbor' private signal. In Levy and Razin (2018), 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.

[8] Examples include 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), Dasaratha and He (2017), among many others). Note that if the agents report their posterior beliefs instead of actions, all our agents' learning outcomes are correct because a linear chain is a social quilt.

[9] Our paper is also related to Lipnowski and Sadler (2018) who assume that agents only form correct conjectures about their neighbors' strategies. Thus in a complete network in which every pair of agents is connected, their agents use Nash equilibrium strategies, just as our agents learn correctly.

field.[10] Recently, Golosov, Qian, and Kai (2015) and Enke and Zimmermann (2017) show people often fail to remove repeated and correlated information when they learn. In particular, they found people are prone to double counting and opinion swings in simple (directed) networks. It stands to reason that such learning errors may persist if the environment is more complex as in a typical network we study. Chandrasekhar, Larreguy, and Xandri (2018) and Grimm and Mengel (2018) compare the two benchmark models of Bayesian learning and naive learning. Grimm and Mengel (2018) find 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 the locally Bayesian learning rule. 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

2.1 Network and beliefs about the network

Let a network be (g, G): $g = \{1, 2, \ldots, I\}$ represents a finite set of agents, and G represents the set of the links among them; $ij \in G$ if i and j are linked.[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: for any $i, h \in g$, there is a path $(i_0 i_1 \ldots i_l)$ 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 $\mathrm{N}_i = \{j : ij \in G\}$. 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 $g_i = \mathrm{N}_i \cup \{i\}$ 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 $\{i, j, \ldots, l\} \subseteq g_i$ consists of themselves, common neighbors to all of them, and all the links among them.[12] We denote this shared local network as $(g_{ij\ldots l}, G_{ij\ldots l})$, where $g_{ij\ldots l} = g_i \cap g_j \cap \ldots \cap g_l$, and $G_{ij\ldots l} = G_i \cap G_j \cap \ldots \cap G_l$. For

[10] See 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.

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

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

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 $G_1 = G_{12} = G_{123} = G$.

Each agent i is assumed to observe only her local network (g_i, G_i). What does an agent believe about the entire network? Intuitively, we assume each agent treats her local network as the entire network, ignoring what she cannot observe.

ASSUMPTION **1.** *Every agent believes that her local network is the entire network: $g_i = g, G_i = G$. Moreover, each agent holding this belief is common knowledge.*

Under this (possibly misspecified) prior, agent i does not update her belief about the network when she communicates with her neighbors about the true state (defined in the next subsection). We call an agent with the above belief, or who acts *as if* she has the above belief, *locally Bayesian*. Each locally Bayesian agent processes information as a Bayesian agent within her local network. It has two implications. First, Assumption 1 uniquely pins down each agent's higher-order beliefs about the network. Since agent i believes (g_i, G_i) is the entire network, she believes that her neighbor j's local network is their shared local network (g_{ij}, G_{ij}). Consequently, agent i believes that j believes (g_{ij}, G_{ij}) is the entire network. Similarly, for any clique $\{i, j, \ldots, l\}$, agent i believes that j believes ... that agent l believes the shared local network $(g_{ij\ldots l}, G_{ij\ldots l})$ is the entire network. Second, because it is common knowledge that each agent believes no other agents exist outside her local network, an agent only forms higher-order beliefs for cliques of agents within her local network. In Figure 1(b), for example, the set $\{1, 2, 3, 4\}$ is not a clique because agent 1 and 3 are not connected. Thus agent 2 does not form belief about 1's belief about 3's belief, because she knows that 1 believes 3 does not exist. We remark on this behavioral assumption further after setting up the model in section 2.3.

2.2 Information structure

Agents in the network want to learn an unknown state, which takes values from a finite state space $S = \{s_1, \ldots, s_N\}$. All the states are *a priori* equally likely: $\Pr(s_n) = 1/N$ for all $s_n \in S$. Agents receive signals from nature about the state.

The support of each agent i's signals is finite: $X^i = \{x^{i,1}, \ldots, x^{i,M_i}\}$. That is, she can observe at least two possible signals, $M_i \geq 2$. For each signal $x^{i,m}$, let $\phi^i_{mn} = \Pr(x^{i,m} \mid s_n)$ be agent i's conditional probability of receiving signal $x^{i,m}$ if the state is s_n. Each agent's information structure $(M_i, \{\phi^i_{mn}\}_{m \leq M_i, n \leq N})$ is identically and independently drawn. For simplicity, we assume $M_i \in \mathbb{N}\setminus\{1\}$ is drawn randomly according to a geometric distribution with rate p_M.[13] For each state s_n, the signals' probability distribution $(\phi^i_{1n}, \ldots, \phi^i_{M_i n})$ is independently uniformly drawn from the *interior* of the set of all probability distributions with M_i outcomes, $\{(\rho_1, \ldots, \rho_{M_i}) : \sum_{m=1}^{M_i} \rho_m = 1 \text{ and } \rho_m > 0\ \forall m\}$. Notice that we assume no signal can completely rule out a state, because every $\phi^i_{mn} > 0$.[14]

Time is discrete: $t = 0, 1, \ldots$. In each period up to period T, agent i observes a realized signal x^i_t according to the information structure above. No informative signal arrives at or after period $T \in \mathbb{N} \cup \{\infty\}$, which is randomly drawn from an (improper) uniform distribution over $\mathbb{N} \cup \{\infty\}$.[15]

Agents' common knowledge includes the (prior) distribution over S, the distribution of each agent i's information structure $(M_i, \{\phi^i_{mn}\}_{m \le M_i, n \le N})$, and the distribution of T. Moreover, it is common knowledge that the signals are independent across agents and time conditional on the state. The true state, each agent's information structure, and T are realized before learning begins. Each agent privately observes her own information structure, and she does not observe the true state, T, or the other agents' information structure.

The above assumptions on the agents' information structure and knowledge are stronger than necessary: they are made merely to ease exposition. It will become clear after Section 3 that all we need is for every agent's potential signals to have full support (and this is common knowledge among agents). That is, agent i can rationalize any posterior belief of neighbor j, because she believes that there is a potential signal in X^j (with the appropriate conditional probabilities) that can generate that particular posterior of agent j given j's prior belief.

2.3 Communication and learning

Agent i learns about the true 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 period-t belief as $\mathbf{b}^i_t = (b^i_t(s_1), \ldots, b^i_t(s_N))$, where $\mathbf{b}^i_t \in \Delta(S)$. Throughout this paper, we use boldface letters to denote vectors. To ease exposition, we use the log-likelihood

[13] That is, $\Pr(M_i = z) = (1 - p_M)^{z-2} p_M$ for each $z \in \mathbb{N} \setminus \{1\}$. We use the geometric distribution merely to fixed ideas. Any discrete probability distribution such that $\Pr(M_i = z) > 0$ for each $z \in \mathbb{N} \setminus \{1\}$ suffices.

[14] Our model can easily accommodate the case when signals can rule out some state s_n, that is, $\phi^i_{mn} = 0$ for some signal $x^{i,m}$. This assumption merely eases the notation since we use log-likelihood ratios of the agents' beliefs throughout this paper.

[15] If $T = \infty$, the agents can receive an infinite number of signals, and if $T = 1$, the agents can receive their 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.

ratios of these beliefs and call them agent i's *estimates* at period t, namely,

$$\boldsymbol{\beta}_t^i = (\beta_t^i(s_1), \ldots, \beta_t^i(s_N)), \text{ where } \beta_t^i(s_n) = \log b_t^i(s_n) - \log b_t^i(s_N).$$

Agent i reports her estimates to her neighbors, and simultaneously receives their reports.[16] She then observes signal $x_t^i \in X^i$ from nature, and period t ends. The timing is summarized in the timeline below. Note that agent i's estimates $\boldsymbol{\beta}_t^i$ are based on the reports and signals she observed prior to period t. We will formally introduce $\boldsymbol{\beta}_t^{ij}, \boldsymbol{\beta}_t^{ijk}, \ldots$, in Section 3.

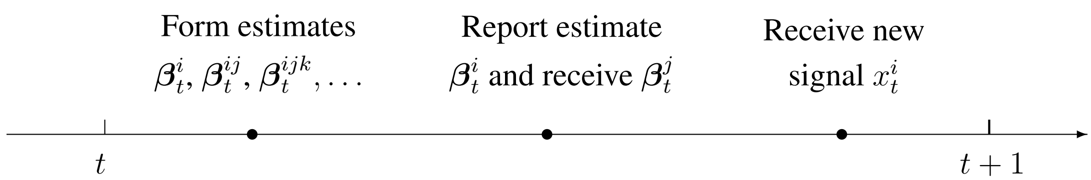

Figure 2: Timeline

Before showing how locally Bayesian agents learn, we remark on two aspects of our model. First, our behavioral assumption, Assumption 1, 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 naive learning. This is a necessary step toward modeling how agents avoid forming wrong beliefs due to repeated and distorted information from the network. Allowing the agents to have perfect memory does add significantly to the complexity of characterizing the agents' short-run learning dynamics and long-run learning outcomes. Specifically, the agents' beliefs do not satisfy the memory-less properties of Markov chains, and thus classic results such as the Perron-Frobenius theorem do not apply.[17] This motivates us to develop a tractable learning rule to implement locally

[16] We do not model a utility function formally, but each agent's report (or action) is consistent with her maximizing a 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: $\mathbb{E}_{\mathbf{b}_t^i}\left[-\sum_{s_n}\left(r_t^i(s_n) - \mathbb{1}_{s_n=s^*}\right)^2\right]$, 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, $\mathbf{r}_t^i = \mathbf{b}_t^i$.

[17] In 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. Thus one can use the Markov chain theory to study 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 her period-$(t-1)$ beliefs and her beliefs based on the earlier information shared in her local network. 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, there are no simple sufficient conditions for convergence.

Bayesian learning and to derive its useful properties, both of which we do in the next section.

Second, locally Bayesian agents can in principle form beliefs about the entire network and about all her neighbors' information structures, but it is not necessary to include these beliefs explicitly in our analysis. To see this, note that Assumption 1 removes an important component—learning about the network structure—from Bayesian learning. Bayesian agents with non-degenerate priors should learn about the network as well as the true state from their neighbors' reports. But it is well-known that the cognitive and computational cost of Bayesian learning about an unknown network is very high.[18] Assumption 1 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 believes there is no outside network and thus behave *as if* all the information from outside their local networks is due to exogenous signals. In addition, while our agent can update her beliefs about her neighbors' information structures as in many standard models, doing so does not affect her learning about the true state, which is the focus of this paper. Intuitively, we will show in Section 3.1 that a locally Bayesian agent i believes that she can extract each of neighbor j's private signals using all the reports agent i can see. Because she does not rely on her beliefs about agent j's information structure to learn his signals, we do not include it in our analysis.

3 The locally Bayesian learning rule

3.1 Extracting new signals using higher-order beliefs

Agent i updates her belief in each period based on all the past reports from her local network and her most recent private signal.[19] Formally, her belief at $t = 1$ is based on x_0^i, and for all $t \geq 2$, her belief is based on $\left\{(\mathbf{b}_\tau^h)_{1\leq\tau\leq t-1, h\in g_i}, x_{t-1}^i\right\}$. The key to locally Bayesian learning is how agent i extracts new information contained in the reports she observes. We now define her higher-order beliefs and illustrate how she uses them to extract new information.

Recall that the underlying uncertainty among the agents is the true state in S, and agent i's first-order belief is $\mathbf{b}_t^i \in \Delta(S)$. Agent i's second-order belief is her belief over the space

[18] An 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.

[19] As defined in Section 2, each agent reports the log-likelihood ratios of her belief every period. One can think of an agent's report as her belief, because there is a one-to-one mapping between them given our assumption that no state is ruled out by any signal.

of S and all her neighbors' first-order beliefs, that is, her second-order belief belongs to $\Delta(S \times (\Delta(S))^{|g_i|-1})$. Next, agent i's *marginal second-order belief* of each neighbor j's first-order belief $\mathbf{b}_t^j$ can be formed by taking expectations of i's second-order beliefs over S and all the other neighbors' (except for agent j's) beliefs. Each of agent i's marginal second-order belief belongs to the space $\Delta(\Delta(S))$; that is, it is a distribution over agent j's belief. Agent i knows that agent j's belief is formed using all the information that agent j has received, including his signal x_{t-1}^j which agent j has received after they exchanged reports. This marginal second-order belief is often hard to compute, because agent i needs to form a belief about x_{t-1}^j. Instead, we introduce a set of simpler statistics which we will show is sufficient for locally Bayesian learning. Specifically, agent i only needs to form her marginal second-order belief in one event: when x_{t-1}^j is uninformative. We call this agent i's *second-order** *belief* from now on.

We can define agent i's higher-order beliefs, her marginal higher-order beliefs, and most importantly, her *higher-order** *beliefs* in a similar way. For instance, her *marginal third-order belief* is her belief about neighbor j's belief about another neighbor k's first-order belief $\mathbf{b}_t^k$, which belongs to the space $\Delta(\Delta(\Delta(S)))$. Her *third-order** *belief* is her belief about agent j's belief about agent k's belief in the event that x_{t-1}^k is uninformative. The log-likelihood ratios of agent i's higher-order* beliefs are her *higher-order** *estimates*. It is worth noting that these higher-order* beliefs are degenerate and easy to compute.

OBSERVATION **1.** *Under Assumption 1, all higher-order** *beliefs are degenerate. For each clique* $\{i, j, \ldots, l\}$*, when* x_{t-1}^l *is uninformative, agent* i *believes with probability* 1 *that agent* j *believes with probability* $1 \ldots$ *that agent* l*'s belief is some probability distribution, denoted as* $\mathbf{b}_t^{ij\ldots l} \in \Delta(S)$.

To see this, start with agent i's second-order* belief about agent j's belief in period t. By definition, agent i forms $\mathbf{b}_t^{ij}$ in the event that x_{t-1}^j is uninformative. Thus $\mathbf{b}_t^{ij}$ contains all the information (agent i believes that) agent j has learned from his neighbors' reports prior to period t. By Assumption 1, agent i believes that she can see agent j's entire local network, which she believes is (g_{ij}, G_{ij}). Thus, in the event that x_{t-1}^j is uninformative, agent i believes that she has access to all the reports that agent j has learned. Therefore she can make the same inferences using these reports and form the same first-order belief as agent j. That is, she believes with probability 1 that j's belief is $\mathbf{b}_t^{ij} \in \Delta(S)$. This argument applies to agent i's all other higher-order* beliefs.

It follows immediately that when agent i hears agent j's report $\mathbf{b}_t^j$, she attributes any difference between agent j's report and her second-order* beliefs to his private signal x_{t-1}^j.

From agent i's perspective, this is the only new information that agent j has but she does not.[20] To differentiate the actual signal x_{t-1}^{j} from what agent i believes to be this signal, we denote the latter as x_{t-1}^{ij}. Formally, agent i recovers x_{t-1}^{ij} by using $\mathbf{b}_t^{ij}$ as her prior and agent j's belief $\mathbf{b}_t^{j}$ as her posterior. By Bayes' rule, for any $s_n \in S$,

$$b_t^{j}(s_n) \;=\; \frac{b_t^{ij}(s_n)\Pr\left(x_{t-1}^{ij} \mid s_n\right)}{\sum_{n'=1}^{N} b_t^{ij}(s_{n'})\Pr\left(x_{t-1}^{ij} \mid s_{n'}\right)}.$$

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

$$\log\frac{b_t^{j}(s_n)}{b_t^{j}(s_N)} \;=\; \log\frac{b_t^{ij}(s_n)}{b_t^{ij}(s_N)} + \log\frac{\Pr\left(x_{t-1}^{ij} \mid s_n\right)}{\Pr\left(x_{t-1}^{ij} \mid s_N\right)}.$$

Using the definition of $\boldsymbol{\beta}_t^{j}$ and $\boldsymbol{\beta}_t^{ij}$, we have:

$$\beta_t^{j}(s_n) \;=\; \beta_t^{ij}(s_n) + \log\frac{\Pr\left(x_{t-1}^{ij} \mid s_n\right)}{\Pr\left(x_{t-1}^{ij} \mid s_N\right)}.$$

Let $\boldsymbol{\alpha}_t^{ij}$ be the log-likelihood ratios of the conditional probability of x_{t-1}^{ij}, we have

$$\alpha_t^{ij}(s_n) \equiv \log\frac{\Pr\left(x_{t-1}^{ij} \mid s_n\right)}{\Pr\left(x_{t-1}^{ij} \mid s_N\right)} = \beta_t^{j}(s_n) - \beta_t^{ij}(s_n). \tag{1}$$

Intuitively, agent i extracts the new signal by removing old information from agent j's report as in the right hand side of (1). From now on, we abuse notations slightly and refer to the log-likelihood ratios $\boldsymbol{\alpha}_t^{ij}$—instead of x_t^{ij}—as the signal agent i extracts from j.

In a similar way, agent i makes inferences about what signal each neighbor in a clique may extract from another neighbor. For instance, consider a triangle $\{ijk\}$. By definition, $\mathbf{b}_t^{ijk}$ is what agent i believes about agent j's second-order* belief about agent k's belief when x_{t-1}^{k} is uninformative. Agent i believes that agent j attributes any difference between agent k's report $\mathbf{b}_t^{k}$ and $\mathbf{b}_t^{ijk}$ to agent k's private signal x_{t-1}^{k}. As above, to differentiate the actual signal x_{t-1}^{k} from what agent i believes j believes to be this signal, we denote the latter as x_{t-1}^{ijk}. In agent i's mind, agent j uses $\mathbf{b}_t^{ijk}$ as the prior and $\mathbf{b}_t^{k}$ as the posterior to extract agent k's private signal. Similar derivations show that the log-likelihood ratios of the conditional

[20] In reality, this difference could be a combination of agent j's signal and what agent j has learned from his neighbors who are not connected to agent i.

probability of x_{t-1}^{ijk} is:

$$\boldsymbol{\alpha}_t^{ijk} = \boldsymbol{\beta}_t^k - \boldsymbol{\beta}_t^{ijk}.$$

Similarly, in any clique $\{i, j, \ldots, l\}$, agent i believes that agent j believes $\ldots$ that agent l extracts $\boldsymbol{\alpha}_{t-1}^{ij\ldots lh}$ from agent $h \in g_{ij\ldots l}$, where

$$\boldsymbol{\alpha}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^h - \boldsymbol{\beta}_t^{ij\ldots lh}. \tag{2}$$

3.2 How do locally Bayesian agents learn?

Agent i learns using a *locally Bayesian learning rule*, which maps the reports she observes in her local network and her private signal into a probability distribution of the states. More specifically, for each agent i and each period t, $LB_t^i(\cdot)$ maps all the reports she observed $\left(\{\boldsymbol{b}_\tau^h\}_{1\le\tau\le t-1, h\in g_i}\right)$ and x_{t-1}^i into (the log-likelihood ratios of) a point in $\Delta(S)$. Similarly, for each clique $\{i, j, \ldots, l\}$, $LB_t^{ij\ldots l}(\cdot)$ maps what she observed into (the log-likelihood ratios of) a point in $\Delta(S)$. Thus, this locally Bayesian learning rule is iterative and self-contained.[21]

We now describe how agent i learns period-by-period. To be consistent with the other signals agent i extracts, let $\boldsymbol{\alpha}_t^{ii} = \{\alpha_t^{ii}(s_1), \ldots, \alpha_t^{ii}(s_N)\}$ be the log-likelihood ratios based on the conditional distribution of her signal from nature x_t^i, that is, for each s_n, $\alpha_t^{ii}(s_n) \equiv \log \Pr(x_t^i \mid s_n) - \log \Pr(x_t^i \mid s_N)$.

Initial values. At the beginning of $t = 1$, agent i learns only from her initial signal. Let $LB_1^i(\cdot) = \boldsymbol{\alpha}_0^{ii}$. Also, let the initial values $LB_1^{ij\ldots l}(\cdot) = LB_1^{ij\ldots lh}(\cdot) = \mathbf{0}$, where $h \in \{i, j, \ldots, l\}$ for each clique $\{i, j, \ldots, l\}$.

At the beginning of each period $t \geq 2$, agent i learns from the most recent reports in her local network and her own signal x_{t-1}^i. Then, agent i forms $LB_t^i(\cdot)$ in two steps:

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

$$\boldsymbol{\alpha}_{t-1}^{ij} = \boldsymbol{\beta}_{t-1}^j - \boldsymbol{\beta}_{t-1}^{ij}. \tag{3}$$

Similarly, she extracts the signal she believes that agent j believes that $\ldots$ agent l extracts from agent h, $h \in g_{ij\ldots l}$. That is, she extracts $\boldsymbol{\alpha}_{t-1}^{ij\ldots lh}$ according to expression (2).

[21] Agent i forms and reports her estimates $\boldsymbol{\beta}_t^i$ in each period as in our timeline (Figure 2). She also simultaneously calculates $LB_t^i(\cdot)$, which may in principle differ from $\boldsymbol{\beta}_t^i$. But we will show in the next subsection that they are the same, and thus the function $LB_t^i(\cdot)$ fully describes the formation of $\boldsymbol{\beta}_t^i$.

Step 2: Updating. Agent i then updates $LB^i_t(\cdot)$ using the signals extracted from each neighbor and from nature.

$$LB^i_t(\cdot) = \boldsymbol{\beta}^i_{t-1} + \sum_{h \in g_i} \boldsymbol{\alpha}^{ih}_{t-1}. \tag{4}$$

In an analogous fashion, for every clique $\{i, j, \ldots, l\}$, agent i updates $LB^{ij\ldots l}_t(\cdot)$ using the signals agent i believes that j believes ... that agent l extracted:

$$LB^{ij\ldots l}_t(\cdot) = \boldsymbol{\beta}^{ij\ldots l}_{t-1} + \sum_{h \in g_{ij\ldots l}} \boldsymbol{\alpha}^{ij\ldots lh}_{t-1}. \tag{5}$$

To complete the learning rule, agent i sets $LB^{ij\ldots lh}_t(\cdot) = LB^{ij\ldots l}_t(\cdot)$ for each $h \in \{i, j, \ldots, l\}$, where h shows up for the second time in this sequence. Agent i does not use the locally Bayesian learning rule for any other sequence of agents involving repeated agents. ‖

We hasten to add that agents who are not locally Bayesian can still use part of this learning rule (expression (3) and (4)), except that they may form their (pseudo) second-order* estimates differently. 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 (pseudo) second-order* estimates about agent j to be (the likelihood ratios of) the uninformative prior: $\tilde{\boldsymbol{\beta}}^{ij}_{t-1} = \mathbf{0}$ for any $t \geq 2$. This implies that, at period t, she does not recall the reports in period $1, \ldots, t-2$. Then by expression (3), $\tilde{\boldsymbol{\alpha}}^{ij}_{t-1} = \tilde{\boldsymbol{\beta}}^{j}_{t-1}$. That is, she treats each neighbor's entire report at period $t-1$ as a new signal, and then she can compute her estimates according to expression (4).

3.3 Implementing locally Bayesian learning

We now show agents who follow the learning rule above form locally Bayesian beliefs, and thus our learning rule is an algorithm to implement locally Bayesian learning.

PROPOSITION **1.** *If agent i follows expression* (2)*,* (3)*,* (4) *and* (5)*, then for all i, t and clique $\{i, j, \ldots, l\}$, $LB^i_t(\cdot) = \boldsymbol{\beta}^i_t$, and $LB^{ij\ldots l}_t(\cdot) = \boldsymbol{\beta}^{ij\ldots l}_t$.*

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. As shown in Section 3.1, agent i believes that her second-order* estimates of agent j's estimates include all the information

j has learned, except for his most recent private signal x^j_{t-1}. She also believes that she can correctly extract x^j_{t-1} after hearing agent j's report containing that signal. Thus, agent i believes that all these signals extracted from her neighbors are independent, and she should update her estimates using them by Bayes' rule, which is expression (4). This implies that $LB^i_t(\cdot)$ formed using the procedure is indeed her estimates $\boldsymbol{\beta}^i_t$. The same argument also applies to all higher-order* estimates.[22]

This result also implies that it is without loss for agents to form higher-order* estimates involving only distinct agents. In Appendix A.1, we show that the agents' learning outcomes do not change even if they form all the (infinitely many) higher-order* estimates. In practice, our learning rule in Section 3.2 reduces the agents' computations significantly.

3.4 Properties of the locally Bayesian learning rule

We now illustrate how our learning rule works and showcase some of its properties.

EXAMPLE **1.** *The network has three agents connected in a line:* $g = \{1, 2, 3\}$ *and* $G = \{12, 23\}$. *The states are binary:* $S = \{s_1, s_2\}$. *The set of signals is* $X^i = \{x^0, x^{i,1}, x^{i,2}\}$, *where* x^0 *is uninformative. Let agent* 1 *receive* $x^1_0 = x^{1,1}$, *agent* 2 *receives* x^0, *and agent* 3 *receives* $x^3_0 = x^{3,1}$. *The corresponding log-likelihood ratios given the two informative signals are* $\log\left(\Pr(s_1 \mid x^{1,1})/\Pr(s_2 \mid x^{1,1})\right) = \varphi^1$ *and* $\log\left(\Pr(s_1 \mid x^{3,1})/\Pr(s_2 \mid x^{3,1})\right) = \varphi^3$.

Throughout our examples, we use the special case of binary states and binary informative signals. Also, we only show $\beta^i_t(s_1)$ when we describe the agents' reports $\boldsymbol{\beta}^i_t$. Since the states are binary and the estimates are in log-likelihood ratios, all $\beta^i_t(s_2) = 0$. The agents' learning dynamics are summarized in the following table.

At $t = 0$, agent 1 and 3 observe x^1_0 and x^3_0 respectively. At $t = 1$, agent 1 reports her estimates based on x^1_0: $\beta^1_1(s_1) = \varphi^1$. Agent 2 has no informative signal and reports $\beta^2_1(s_1) = 0$. Agent 3 reports her estimates based on x^3_0: $\beta^3_1(s_1) = \varphi^3$. The initial second-order* estimates are all $\mathbf{0}$. This is summarized in the first row of Table 1.

At $t = 2$, agent 2 extracts $\alpha^{21}_1(s_1) = \beta^1_1(s_1) - \beta^{21}_1(s_1) = \varphi^1$ from agent 1 and extracts $\alpha^{23}_1(s_1) = \beta^3_1(s_1) - \beta^{23}_1(s_1) = \varphi^3$ from agent 3, both by expression (3). By expression (4),

[22] It will become clear using the results in the next section that Proposition 1 can be generalized. It continues to hold if we use a weaker version of Assumption 1 such that 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. Also, each agent holding this belief is common knowledge. These types of beliefs are consistent with Fainmesser and Goldberg (2016) who show in a random network in which the number of each agent's neighbors is bounded, as the population gets large, each agent believes asymptotically that the network is a random tree where she is the root agent.

	$\beta_t^1(s_1)$	$\beta_t^2(s_1)$	$\beta_t^3(s_1)$
$t=1$	φ^1	0	φ^3
$t=2$	φ^1	$\varphi^1+\varphi^3$	φ^3
$t\geq 3$	$\varphi^1+\varphi^3$	$\varphi^1+\varphi^3$	$\varphi^1+\varphi^3$

Table 1: A three-agent line

$\beta_2^2(s_1)=\varphi^1+\varphi^3$. Agent 1 and 3 do not learn from agent 2: $\alpha_1^{12}(s_1)=\alpha_1^{32}(s_1)=0$.

At $t=3$, agent 1 extracts $\alpha_2^{12}(s_1)=\varphi^3$ and agent 3 extracts $\alpha_2^{32}(s_1)=\varphi^1$. Thus their estimates are those in the third row of the table. Note that agent 2 expects 1 and 3 to learn from her and does not change her estimates. For all $t\geq 4$, no agent changes her estimates and their beliefs are the correct Bayesian posterior given the two informative signals. $\diamond$

Two nice properties of our locally Bayesian learning rule greatly simplify our analysis. First, locally Bayesian updating 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 $\boldsymbol{\beta}_t^i$ is agent i's estimates when X_{t-1} is the set of signals from nature. Let $\boldsymbol{\beta}_t^{\mu,i}$ and $\boldsymbol{\beta}_t^{\nu,i}$ 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 $t\geq 1$,*

$$\boldsymbol{\beta}_t^{i}=\boldsymbol{\beta}_t^{\mu,i}+\boldsymbol{\beta}_t^{\nu,i}; \tag{6}$$

$$\boldsymbol{\beta}_t^{ij}=\boldsymbol{\beta}_t^{\mu,ij}+\boldsymbol{\beta}_t^{\nu,ij}. \tag{7}$$

Lemma 1 shows 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}. It allows us to study one signal at a time: if the agents' learning outcomes are correct under every signal, their learning outcomes are also correct under any sequence of these signals. The intuition can be seen from Example 1: divide the two signals into $X_{t-1}^{\mu}=\{x_0^1\}$ and $X_{t-1}^{\nu}=\{x_0^3\}$. Under X_{t-1}^{μ}, everyone's estimates are $\beta_3^{\mu,i}(s_1)=\varphi^1$ at $t=3$. Similarly, under X_{t-1}^{ν}, everyone's estimates are $\beta_t^{\nu,i}(s_1)=\varphi^3$ at $t=3$. When nature sends both signals, even to different agents (or in different periods), the agents' learning about one signal is independent of the other one. At $t=3$, their estimates are the sum of their estimates under X_{t-1}^{μ} and X_{t-1}^{ν}.

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 (3)) and to incorporate it into her own estimates (expression (4)), and 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 due to what agent i did not observe.

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

$$\boldsymbol{\alpha}_t^{ij} = \sum_{l \in (g_j \setminus g_i) \cup \{j\}} \boldsymbol{\alpha}_{t-1}^{jl} + \sum_{h \in g_{ij} \setminus \{j\}} \left(\boldsymbol{\alpha}_{t-1}^{jh} - \boldsymbol{\alpha}_{t-1}^{ijh} \right). \tag{8}$$

We can decompose $\boldsymbol{\alpha}_t^{ij}$, the signal agent i extracts from neighbor j at the beginning of period $t+1$, into two parts according to equation (8). The first part consists of what agent j has just learned from nature ($\boldsymbol{\alpha}_{t-1}^{jj}$) and from his neighbors who are not connected to agent i ($\boldsymbol{\alpha}_{t-1}^{jl}$ for $l \in g_j \setminus g_i$). In Example 1, the signal agent 1 extracts from agent 2 at $t = 3$ is the new signal agent 2 extracted from agent 3 at $t = 2$: $\boldsymbol{\alpha}_2^{12} = \boldsymbol{\alpha}_1^{23}$. Moreover, this part also shows that agent i does not mistakenly learn old information from her local network again, unlike in models with imperfect memory such as DeGroot (1974).[23]

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 ($\boldsymbol{\alpha}_{t-1}^{jh} - \boldsymbol{\alpha}_{t-1}^{ijh}$) is the difference between what agent j extracted from h and what agent i believes agent j extracted from h. The second term is zero in certain networks such as the three-agent line in Example 1. It is not zero in other networks in which some agents know they learn from the same source while others don't, which is the failure of local connection symmetry described in the introduction. For example, in the diamond with a link in Figure 1(b), agent 2 and 4 know any signals they extract from 1 are perfectly correlated, but agent 3 believes they are independent. Therefore what agent 3 believes 2 extracts from 4 could differ from what 2 truly extracts from 4: $\boldsymbol{\alpha}_{t-1}^{24} \neq \boldsymbol{\alpha}_{t-1}^{324}$. We discuss this type of learning error in more details in Section 5.2.

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 these central questions, we lay out our notions of correct learning. Our strongest notion of correct learning is for each agent to learn correctly in every period *given the travel paths of signals*. To define it, we

[23]To see this, note that the first part of expression (8) does not include what j has learned from i (no $\boldsymbol{\alpha}_{t-1}^{ji}$). It also does not include what agent j has learned from a common neighbor k (no $\boldsymbol{\alpha}_{t-1}^{jk}$, $k \in \mathrm{N}_i \cap \mathrm{N}_j$). Agent i does not mistakenly treat these old information in her local network as new.

begin with the set of signals that can reach agent i in period t. Recall that X_t is the union of X_t^i, the set of signals agent i receives up to and including period t from nature. Since T is the period at or after which the agents receive no informative signal, X_T contains 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 $l \in g$, 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 $X^l_{t-d(il)-1}$, where $X^l_{t-d(il)-1} = \emptyset$ if $t < d(il) + 1$. Suppose that agent i can correctly identify and incorporate every signal that has reached her at the beginning of t once and only once, then for every $s_n \in S$, her Bayesian posterior is:

$$q_t^i(s_n) = \Pr\left(s_n \mid X^1_{t-d(i1)-1}, \ldots, X^I_{t-d(iI)-1}\right). \tag{9}$$

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

Strong efficiency is the strongest notion of correct learning in the network context. It implies that when T is finite, all agents form the correct posterior at or before period $T+D$.[24] 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. But it is not appropriate when the agents only receive a finite number of signals because the correct Bayesian posterior is bounded away from 0 and 1. In this case, we use efficient learning which requires the agents' estimates in the long run to match (the log-likelihood ratios of) the Bayesian posterior.

[24]Even when the network is common knowledge and all agents are Bayesian, it often takes much longer than the diameter of the network for agents to learn (see Mossel, Olsman, and Tamuz (2016)).

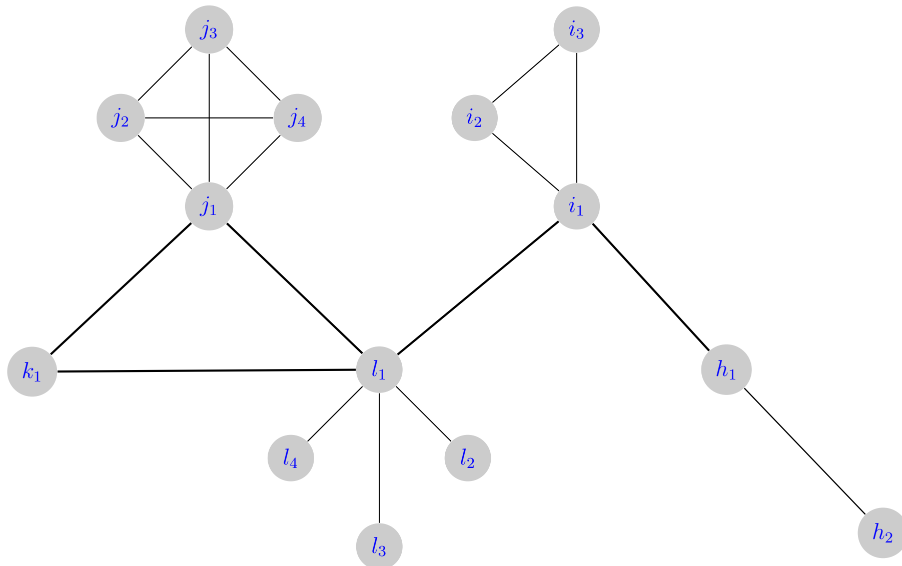

Figure 3: A social quilt

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, she should 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 a particular type of networks, *a social quilt*, and only this type of networks, meets these conditions. Recall a path $(i_1 \dots i_l)$ is a circle if $i_1 i_l \in G$. Also, a graph is a *tree* if it contains no circle.

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 3 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.[25] Our main result is an intuitive and clean relationship between social quilts and strongly efficient learning outcomes.

[25]The 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

PROPOSITION **2.** *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.*

Proposition 2 shows that in social quilts, the agents do not suffer from correlation neglect, a well-known and commonly observed learning error. To see why, observe that a locally Bayesian agent treats any "unexpected" change in her neighbors' reports as new, independent information. But this approach has two pitfalls. First, if information travels through a large circle (beyond those embedded in an agent's local network), she cannot identify this as old information and thus will double count it. This cannot happen in a social quilt because it is a tree globally (connecting all cliques), and thus no information can travel back and reach an agent a second time. Second, if her neighbors have a common neighbor who she cannot observe, then her neighbors' reports can be correlated, but she does not know that and still treats these reports as independent. This mistake also cannot happen in a social quilt because each agent is in a local clique, and if two of her neighbors share a common neighbor, she must know that common neighbor. Together, a social quilt—a global tree of local cliques—ensures that any unexpected change in a neighbor's report is truly due to new signals, and thus the locally Bayesian agents' learning outcomes are strongly efficient.

We now define two features that jointly characterize a social quilt before examining their respective roles in Proposition 2 in more depth.

LEMMA **3.** *Network (g, G) is a social quilt if and only if*

1. *it contains no* ***simple circle****, which is a circle that contains at least four agents and each agent has exactly two links to other agents in the circle, and*
2. *every agent i's local network satisfies* ***local connection symmetry****: g_{ij} is a clique for every $j \in \mathrm{N}_i$.*

By definition, a social quilt has no simple circles. Whenever a network has simple circles, there are multiple paths between one agent and another. As a result, each signal could travel along these different paths and reach an agent repeatedly. For example, (1234) in Figure 1(a) is a simple circle. If agent 1 has a signal, it reaches agent 2 and 4 first, and then agent 3 will double count it as she learns from both of her neighbors.

with the core-periphery structure are social quilts, which are important for financial markets (Babus and Kondor (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.

Next, local connection symmetry for agent i (LCS$_i$ from now on) holds if for any neighbor $j \in \mathrm{N}_i$, $\mathrm{N}_i \cap \mathrm{N}_j = \emptyset$, which is the case when they are part of a simple circle or a line. For example, in the simple circle (1234) in Figure 1(a), LCS$_1$ holds because agent 1 and 2, as well as 1 and 4, do not have any common neighbor. LCS$_i$ also holds if each pair of agent i and j's common neighbors k and l are connected, for instance if the network is a clique. In contrast, in the diamond with a link network in Figure 1(b), LCS$_2$ does not hold because 2 and 4 have two common neighbors 1 and 3 who are not connected. We say that a network satisfies *local connection symmetry* (LCS from now on) if LCS$_i$ holds for all $i \in g$. Given these definitions, it is easy to show that if a network satisfies LCS and contains no simple circle, then any circle must be in a clique and the network must be a social quilt.

LCS ensures that agents have symmetric knowledge about information correlation in their local networks, which is crucial for the agents' higher-order* estimates to be well-behaved.[26] To prove Proposition 2, we need to show that there is cross-agent consistency: agent i's estimates of j's estimates of their common neighbor h's estimates are exactly j's estimates of h's estimates, and so on for all higher-order* estimates. To see why it matters, recall the iterative rule characterizing a signal's travel from Lemma 2. The second part of expression (8) is

$$\sum_{h \in g_{ij} \setminus \{j\}} (\boldsymbol{\alpha}_{t-1}^{jh} - \boldsymbol{\alpha}_{t-1}^{ijh}). \tag{10}$$

If the network satisfies LCS, then we show that $\boldsymbol{\beta}_{t-1}^{ijh} = \boldsymbol{\beta}_{t-1}^{jh}$ in the appendix, and thus all these differences in (10) are zero. That is, there is no local learning errors because the signal agent i believes j has extracted from h is exactly what agent j extracted from h. The same arguments apply to all the higher-order* estimates.

The above two features imply that every agent learns a 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 sequences of realized 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. Proposition 2 also shows that social quilts are necessary for the agents to have strongly efficient learning outcomes for all realized sequences of signals. When a network is not a social quilt, it must either contain simple circles or fail LCS. Each of them

[26]To see this, note that when LCS fails such as for agent 2 and 4 in Figure 1(b), agents then have asymmetric knowledge about information correlation. In particular, agent 2 and 4 know any signal they learn from agent 1 is perfectly correlated, but agent 3 thinks they are independent.

leads to a specific type of learning error we study in the next section.

5 When efficient learning is impossible

5.1 Repetitions due to echo chambers

To isolate the learning error caused by simple circles, we consider a network that satisfies LCS, but is not a social quilt. By Lemma 3, 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 extracting "new" signals from her neighbors when it is the same signal reaching her again and again through the simple circle(s).

EXAMPLE **2.** *Consider the four-agent simple circle in Figure 1(a). Recall that* $S = \{s_1, s_2\}$*. Let* $X = \{x^0, x^1, x^2\}$ *and the informative signals be symmetric:* $\Pr(x^1 \mid s_1) = \Pr(x^2 \mid s_2) = \phi$*. Agent* 1 *receives the only informative signal* $x_0^1 = x^1$*. The corresponding log-likelihood ratio is* $\log\left(\Pr(s_1 \mid x^1)/\Pr(s_2 \mid x^1)\right) = \varphi$*.*

	$\beta_t^1(s_1)$	$\beta_t^2(s_1) = \beta_t^4(s_1)$	$\beta_t^3(s_1)$
$t = 1$	φ	0	0
$t = 2$	φ	φ	0
$t = 3$	φ	φ	2φ
$t = 4$	φ	2φ	2φ
$t = 5$	3φ	2φ	2φ

Table 2: Learning in a simple circle

The signal x_0^1 travels from agent 1 in both directions. 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 (3) yields $\alpha_3^{23}(s_1) = \alpha_3^{43}(s_1) = \varphi$. That is, agent 2 (and agent 4) extracts a second copy of the signal from agent 3 because he only expects agent 3 to learn one copy from himself, but agent 3 reports 2φ instead. At $t = 5$, agent 1 learns two new copies from agent 2 and agent 4, and thus she believes there are three copies of the signal (the first five periods are summarized in Table 2). Similarly, in every four periods, the agents learn two additional copies of the signal. In each period $t = 4\tau + 1$, $\tau = 0, 1, \ldots$, agent 1 believes in $2\tau + 1$ copies of the signal and all other 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.

PROPOSITION **3.** *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 ($t \to \infty$), every agent is wrong, because they believe in the state that is most likely given x_0^i with probability 1. 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. In the case of one simple circle, Proposition 3 shows locally Bayesian agents have the wisdom of the crowds when they receive infinitely many signals. 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 extracts 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 extracted 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 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: the Law of Large Numbers can fail. Specifically, each signal travels both within a simple circle, and back and forth from one simple circle to another. Agents in one simple circle keep extracting more and more new signals from all the other simple circles, and passing their own repeatedly extracted signals to them. The number of copies of each signal grows exponentially. Thus in any network with two or more simple circles, there exists a period after which 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 a wrong state. This is because each of the correct new signals arrives too late, and is dominated by the exponentially growing existing signals.

Proposition 3 suggests that fake news—propaganda and disinformation pretending to be real news—may thrive in networks containing multiple simple circles ("echo chambers").[27] 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 Figure 4.

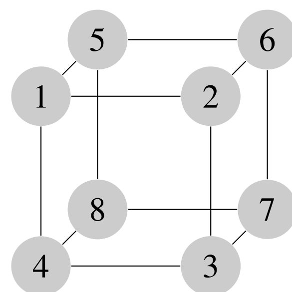

Figure 4: A cube of eight agents

EXAMPLE **3.** *Eight agents are connected in a cube as in Figure 4. The information structure is the same as in Example 2. 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 $t \to \infty$, everyone believes the true state is s_2 with a probability arbitrarily close to 1.*

	$t = 1$	$t = 2$	$t = 3$	$t \geq 4$
$\beta_t^i(s_1)$	$-\varphi$	-3φ	-5φ	$-(2t-1)\varphi$

Table 3: Learning in a cube

All agents are symmetric in this example and their estimates are updated as in Table 3. Why do they believe in state s_2 despite so many opposing (and correct) signals from $t = 1$ onward? Observe that at $t = 1$, each agent reports $\beta_1^i(s_1) = -\varphi$ which is based on the initial

[27]This 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.

signal x^2. At $t = 2$, each agent extracts three signals of x^2 from their neighbors, in addition to her own signal of x^1 . Therefore her count of copies of x^2 increases by two and she reports $\beta_2^i(s_1) = -3\varphi$. Her estimates of each neighbor j's estimates are $\beta_2^{ij}(s_1) = -2\varphi$, because she believes that agent j learns a signal of x^2 from herself plus his own signal of x^2. Therefore at $t = 3$, each agent extracts another x^2 from each neighbor, $\alpha_2^{ij}(s_1) = \beta_2^j(s_1) - \beta_2^{ij}(s_1) = -\varphi$, net of one copy of x^1 from nature. Therefore she believes in two more copies of x^1, just like in period 2. The agents' learning in each ensuing 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 failure of local connection symmetry

Proposition 2 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. We first illustrate this learning error with an example.

EXAMPLE **4.** *Consider the diamond with a link network in Figure 1(b). The information structure is the same as in Example 2. Let $x_0^1 = x^1$ be the only informative signal. The corresponding log-likelihood ratio remains* $\log(\Pr(s_1 \mid x^1)/\Pr(s_2 \mid x^1)) = \varphi$.

The agents' learning outcomes are summarized in Table 4. Recall that agent 2's and agent 4's local network fails LCS_2 and LCS_4. At $t = 1$, agent 1 reports $\beta_t^1(s_1) = \varphi$. At $t = 2$, agent 2 and 4 learn the signal from agent 1, and thus $\beta_2^2(s_1) = \beta_2^4(s_1) = \varphi$. Since agent 2 and 4 know the entire network, they form the correct posterior. So does agent 1 since he will not learn new information from 2 and 4: $\beta_t^1(s_1) = \beta_t^2(s_1) = \beta_t^4(s_1) = \varphi$ for $t \geq 2$.

	$\beta_t^1(s_1)$	$\beta_t^2(s_1) = \beta_t^4(s_1)$	$\beta_t^3(s_1)$	$\alpha_{t-1}^{32}(s_1) = \alpha_{t-1}^{34}(s_1)$
$t = 1$	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.

At $t = 3$, agent 3 extracts two signals, one from agent 2 and one from agent 4, so $\beta_3^3(s_1) = 2\varphi$. Also, agent 3 believes that agent 2 and 4 should learn from each other because he believes these two signals are independent. That is, $\beta_3^{32}(s_1) = \beta_3^{34}(s_1) = 2\varphi$. Interestingly, at $t = 4$, agent 3 compares $\beta_3^2(s_1) = \varphi$ with $\beta_3^{32}(s_1) = 2\varphi$, and extracts $\alpha_3^{32}(s_1) = -\varphi$, a signal negatively correlated with the initial signal. He extracts another negatively correlated copy from agent 4, and thus $\beta_4^3(s_1) = 0$. Intuitively, 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: 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 agent 3 to report 2φ in odd periods and 0 in even periods, because they know agent 3 does not know agent 1 exists. Therefore their own estimates are unaffected by agent 3's opinion swings. $\diamond$

In the example above, 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 **4.** *Consider a network with no simple circles, but fails LCS. Then there exists a sequence of signals X_T, $T = \infty$, such that at least one agent's learning outcomes are not efficient (and not converging).*

If the network fails LCS, then we can find at least one diamond with a link embedded in the network. That is, some agent l (like agent 3 in Example 4) has two (or more) neighbors who share a common neighbor, whom agent l does not know. Proposition 4 shows that the oscillation of agent l as found in Example 4 persists when the four of them are embedded in a larger network that contains no simple circles. To show it, we use a key feature of learning in the networks without simple circles: a signal travels sequentially away from the agent who receives it and never travel backwards. If agent i receives a signal, we can classify the agents by their distance to agent i, $\mathrm{N}_i^d = \{h \in g : d(ih) = d\}$. Then this feature says that no agent in N_i^d extracts any new signal from her successors in N_i^{d+1}. Therefore, when agent i receives the only (correct) signal, agent l's oscillation persists because he does not learn any signal back from his successors. If agent i receives more correct signals, it could exacerbate 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 also fails LCS. This type of learning error may lead to unreliable poll results and unstable experimental outcomes.

If a network has simple circles and fails LCS, both repetition and belief oscillations occur locally. For any such network, our learning rule provides an algorithm to calculate the learning dynamics. But we are unable to fully characterize the agents' learning outcomes because this problem lacks structure in general. Note that whenever a signal reaches a subnetwork that fails LCS, some agent in the subnetwork extracts signals negatively correlated with the original signal. Unlike in Proposition 4, 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 characterize the net number of signals for any network.[28]

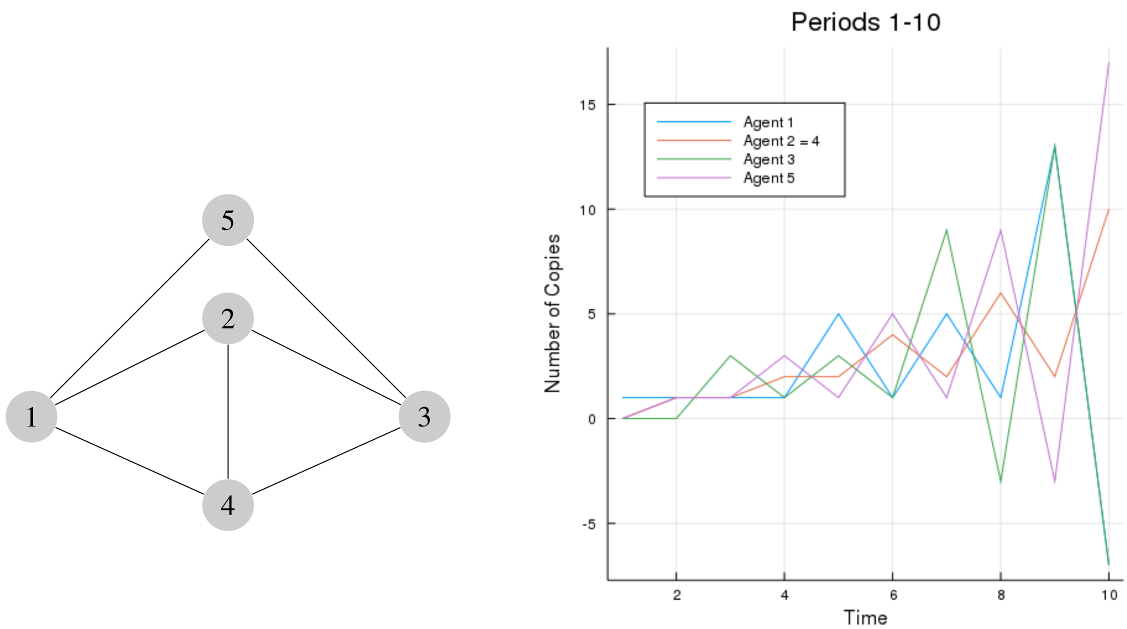

Figure 5: (a) Expanded diamond with a link (b) learning dynamics in this network

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 in Figure 5(a), 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 extract as many negatively correlated copies. In short, to every positively correlated signal there is always an equal negatively

[28]While 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.

correlated signal. Figure 5(b) 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 types of 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 model can be extended in several directions. First, we can relax the behavioral assumption which makes agents believe information from outside their local networks is independent. Suppose agents account for repeated information from outside their local networks by a simple rule-of-thumb: Dismiss any signal they have already extracted as old information. We can show that with this simple rule, their learning outcomes are strongly efficient in any network if all signals reach the *same* agent initially. Therefore a policymaker may want to disseminate information through one central agent over time. Second, one may argue that the locally Bayesian learning still demands a high level of cognitive and computational ability from agents. In Li and Tan (2019), we study how agents with cognitive constraints learn in local networks. We show there exists a critical level of cognitive ability (which can be very low) above which the agents' learning outcomes will be correct.

References

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 (2018): “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 Infuence, 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 (2017): "Correlation neglect in belief formation," *The Review of Economic Studies*, 86(1), 313–332.

EYSTER, E., AND M. RABIN (2010): "Nave Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2(4), 221–43.

——— (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 (2018): "Experiments on belief formation in networks," *Journal of the European Economic Association*.

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 (2018): "Information diffusion in networks with the Bayesian Peer Influence heuristic," *Games and Economic Behavior*, 109, 262–270.

LI, W., AND X. TAN (2019): "Cognitively constrained learning in local networks," mimeo.

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 (2018): “A Theory of Non-Bayesian Social Learning,” *Econometrica*, 86(2), 445–490.

MOSSEL, E., N. OLSMAN, AND O. TAMUZ (2016): “Efficient bayesian learning in social networks with gaussian estimators,” in *the 54th annual Allerton conference on Communication, control, and computing*. IEEE Press.

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.

XIE, D., AND S. WANG (2012): “Consensus of second-order discrete-time multi-agent systems with fixed topology,” *Journal of Mathematical Analysis and Applications*, 387(1), 8–16.

YANG, T., Z. MENG, D. V. DIMAROGONAS, AND K. H. JOHANSSON (2014): “Global consensus for discrete-time multi-agent systems with input saturation constraints,” *Automatica*, 50(2), 499–506.

A Appendix: An extension and proofs

A.1 A general learning rule allowing for any sequence of agents

In our locally Bayesian learning rule described in Section 3, agent i forms $LB_t^{ij\ldots l}(\cdot)$ for each clique $\{i, j, \ldots, l\}$ within her local network. Moreover, she directly sets values (instead of forming them through the learning rule) when the last agent is a repeated agent, that is, for $h \in \{i, j, \ldots, l\}$, she sets $LB_t^{ij\ldots lh}(\cdot) = LB_t^{ij\ldots l}(\cdot)$. One may wonder whether our learning rule is with loss because agent i does not apply the learning rule to all other sequences of agents involving repeated agents. In this section, we show that the answer is no.

To do so, we first describe a *complete locally Bayesian learning rule*, denoted as $CLB_t^i(\cdot)$. We say a sequence of agents is *fully-connected* if it contains at least two distinct agents, and every pair of distinct agents in the sequence is connected. We allow agent i to apply $CLB_t^i(\cdot)$ to all sequences of fully-connected agents in her local network (we drop $(\cdot)$ for simplicity in the rest of Appendix A). Then we show that the learning outcomes of these two rules are the same. Clearly, the learning rule in Section 3 economizes on computation.

Initial values. At the beginning of $t = 1$, agent i learns from her initial signal. Let $CLB_1^i = \boldsymbol{\alpha}_0^{ii}$. Also, let the initial values $CLB_1^{i \ldots r} = \mathbf{0}$ for every sequence of fully-connected and possibly repeated agents $(i \ldots r)$.

At the beginning of each period $t \geq 2$, agent i learns from the most recent reports in her local network and her own signal x_{t-1}^i. Then, agent i forms CLB_t^i in two steps:

Step 1: Extracting new information. Agent i extracts a new signal $\boldsymbol{\alpha}_{t-1}^{ij}$ from each neighbor j. This is the same as expression (3),

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

Similarly, she extracts the signal she believes that . . . agent r extracts from agent h, $h \in g_{i \ldots r}$. That is, she extracts $\boldsymbol{\alpha}_{t-1}^{i \ldots rh}$ as follows:

$$\boldsymbol{\alpha}_{t-1}^{i \ldots rh} = \boldsymbol{\beta}_{t-1}^{h} - \boldsymbol{\beta}_{t-1}^{i \ldots rh}, \tag{11}$$

which is the counterpart of expression (2) in the text.

Step 2: Updating. Agent i then updates CLB_{t-1}^i using the signals extracted from each neighbor and from nature. This is the counterpart of expression (4):

$$CLB_t^i = \boldsymbol{\beta}_{t-1}^{i} + \sum_{h \in g_i} \boldsymbol{\alpha}_{t-1}^{ih}. \tag{12}$$

In an analogous fashion, agent i updates $CLB_{t-1}^{i \ldots r}$ using the signals agent i believes . . . that agent r extracted. This is the counterpart of expression (5):

$$CLB_t^{i \ldots r} = \boldsymbol{\beta}_{t-1}^{i \ldots r} + \sum_{h \in g_{i \ldots r}} \boldsymbol{\alpha}_{t-1}^{i \ldots rh}, \tag{13}$$

for each sequence of fully-connected (possibly repeated) agents $(i \ldots r)$. $\|$

Agent i applies the complete locally Bayesian learning rule to infinitely many sequences of agents in her local network, which involves a large amount of computation. We now show that only sequences of distinct agents matter. Therefore the much simpler locally Bayesian learning rule in the text yields the same learning outcomes.

LEMMA **4.** *Let the set of distinct agents in a sequence of fully-connected agents* $(l_1 \ldots l_z)$ *be* $\{i, j, \ldots, l\}$. *Then,* $CLB_t^{l_1 \ldots l_z} = CLB_t^{ij\ldots l} = LB_t^{ij\ldots l}$ *for all* $t \geq 1$.

Proof of Lemma 4: At $t = 1$, by definition, $CLB_1^{l_1 \ldots l_z} = CLB_1^{ij\ldots l} = LB_1^{ij\ldots l} = \mathbf{0}$. Next, consider any period $t \geq 2$. To begin with, because $\{i, j, \ldots, l\}$ is the set of distinct agents in the sequence $(l_1 \ldots l_z)$, the shared local networks include the same set of agents: $g_{ij\ldots l} = g_{l_1 \ldots l_z}$. By Assumption 1, agent i believes that agent j believes $\ldots$ that agent l believes the set of agents in the network is $g_{ij\ldots l}$. Agent i forms her higher-order* estimates $\boldsymbol{\beta}_t^{ij\ldots l}$ in the event that x_{t-1}^l is uninformative; that is, agent i only uses the reports in the shared local network. The same is true when agent l_1 forms her higher-order* estimates $\boldsymbol{\beta}_t^{l_1 \ldots l_z}$. Thus, the higher-order* estimates $\boldsymbol{\beta}_t^{ij\ldots l}$ and $\boldsymbol{\beta}_t^{l_1 \ldots l_z}$ are the same, because they are formed based on the same set of reports $\{\boldsymbol{\beta}_\tau^h\}_{1 \leq \tau \leq t-1, h \in g_{ij\ldots l}}$. Then by expression (2) and (11), we have for any $h \in g_{ij\ldots l}$,

$$\boldsymbol{\alpha}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^h(s_n) - \boldsymbol{\beta}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^h - \boldsymbol{\beta}_t^{l_1 \ldots l_z h} = \boldsymbol{\alpha}_t^{l_1 \ldots l_z h}. \tag{14}$$

Then, using expression (5), (13) and (14), we have:

$$CLB_{t+1}^{ij\ldots l} = \boldsymbol{\beta}_t^{ij\ldots l} + \sum_{h \in g_{ij\ldots l}} \boldsymbol{\alpha}_t^{ij\ldots lh} = \boldsymbol{\beta}_t^{l_1 \ldots l_z} + \sum_{h \in g_{l_1 \ldots l_z}} \boldsymbol{\alpha}_t^{l_1 \ldots l_z h} = CLB_{t+1}^{l_1 \ldots l_z}.$$

Lastly, by expression (5) and (13), it is easy to see that,

$$CLB_t^{ij\ldots l} = \boldsymbol{\beta}_{t-1}^{ij\ldots l} + \sum_{h \in g_{ij\ldots l}} \boldsymbol{\alpha}_{t-1}^{ij\ldots lh} = LB_t^{ij\ldots l}.$$

Thus, the two learning rules yield the same learning outcomes. ▮

A.2 Proofs

Proof of Proposition 1: Recall that for all $t \geq 2$, agent i forms $LB_i^t(\cdot)$ from the entire history of reports $(\{\boldsymbol{\beta}_\tau^h\}_{1 \leq \tau \leq t-1, h \in g_i})$ and her latest private signal x_{t-1}^i. Appendix A.1 above shows that it is without loss for agents to apply the locally Bayesian learning rule only to sequences of distinct agents. We now show $LB_t^i(\cdot) = \boldsymbol{\beta}_t^i$, and $LB_t^{ij\ldots l}(\cdot) = \boldsymbol{\beta}_t^{ij\ldots l}$ for all i, t

and clique $\{i, j, \ldots, l\}$. For simplicity, we drop the $(\cdot)$ from the learning rule in the rest of the proof.

At $t = 1$, agent i only has her initial signal x_0^i. The log-likelihood ratio of her Bayesian posterior is $\boldsymbol{\beta}_1^i = \boldsymbol{\alpha}_0^{ii}$ by definition. All her higher-order* estimates $\boldsymbol{\beta}_1^{ij\ldots l} = \mathbf{0}$, because they are formed in the event agent l has no informative signals. By definition, the initial values $LB_1^i = \boldsymbol{\alpha}_0^{ii}$ and $LB_1^{ij\ldots l} = \mathbf{0}$.

For all $t \geq 2$, by expression (3), agent i extracts $\boldsymbol{\alpha}_{t-1}^{ij} = \boldsymbol{\beta}_{t-1}^{j} - \boldsymbol{\beta}_{t-1}^{ij}$ from each $j \in \mathrm{N_i}$, which is the log-likelihood ratio of signal x_{t-2}^{ij} as described in Section 3.1. By Assumption 1, agent i believes these are the log-likelihood ratios of x_{t-2}^{j}, and believes that they are all the new signals the other agents received since the previous set of reports. Recall that $\mathbf{b}_{t-1}^i$ is her Bayesian posterior belief given all her information $(\{\boldsymbol{\beta}_\tau^h\}_{1\leq\tau\leq t-2, h\in g}, x_{t-2}^i)$ up to the end of period $t-2$. She only incorporates what she believes to be new information into her estimates. That is, she uses $\mathbf{b}_{t-1}^i$ as her prior and incorporates all the signals she extracted (x_{t-2}^{ij}) and her own signal (x_{t-1}^i) into $\mathbf{b}_t^i$ by Bayes' rule. For every $s_n \in S$, we have

$$b_t^i(s_n) \propto b_{t-1}^i(s_n)\mathrm{Pr}(x_{t-1}^i \mid s_n) \prod_{j\in \mathrm{N}_i} \mathrm{Pr}(x_{t-2}^{ij} \mid s_n).$$

Take the log-likelihood ratios, and we have $\boldsymbol{\beta}_t^i = \boldsymbol{\beta}_{t-1}^i + \sum_{h\in g_i} \boldsymbol{\alpha}_{t-1}^{ih}$. This is exactly expression (4), and thus $LB_t^i = \boldsymbol{\beta}_t^i$.

Next, recall that $\boldsymbol{\beta}_{t-1}^{ij\ldots l}$ is her higher-order* estimates given all her information up to the end of period $t-2$ when agent l receives an uninformative signal. Similar to above, agent i believes that $(g_{ij\ldots l}, G_{ij\ldots l}) = (g_{j\ldots l}, G_{j\ldots l})$ by Assumption 1, and thus she knows all the reports agent j believes that ... agent l can observe. Therefore she can extract all the signals one neighbor can extract from another using expression (2). Specifically, for every $s_n \in S$, by Bayes' rule,

$$b_t^{ij\ldots l}(s_n) \propto b_{t-1}^{ij\ldots l}(s_n) \prod_{h\in g_{ij\ldots l}} \mathrm{Pr}(x_{t-2}^{ij\ldots lh} \mid s_n).$$

Take the log-likelihood ratios, and agent i's updated higher-order* estimates $\boldsymbol{\beta}_t^{ij\ldots l}$ follows expression (5) exactly. Thus $LB_t^{ij\ldots l} = \boldsymbol{\beta}_t^{ij\ldots l}$. ∎

Proof of Lemma 1: Recall the definition of the disjoint sets (X_t^μ, X_t^ν). For each agent i, let $\{x_t^{\mu,i}, x_t^{\nu,i}\} = \{x_t^i, x^0\}$, where x^0 is the uninformative signal. That is, agent i is uninformed in one and learns x_t^i in the other. In addition to equations (6) and (7) in the lemma, we claim that for any clique, $\{i, j, \ldots, l\}$ and $t \geq 1$,

$$\boldsymbol{\beta}_t^{ij\ldots l} = \boldsymbol{\beta}_t^{\mu,ij\ldots l} + \boldsymbol{\beta}_t^{\nu,ij\ldots l}. \tag{15}$$

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

By the definition of $\{x_0^{\mu,i}, x_0^{\nu,i}\}$, we have $\{\boldsymbol{\beta}_1^{\mu,i}, \boldsymbol{\beta}_1^{\nu,i}\} = \left\{\boldsymbol{\beta}_1^{i}, \mathbf{0}\right\}$. Also, all the higher-order* estimates are $\mathbf{0}$ by definition since there has been no previous report. Thus equations (6), (7) and (15) hold at $t = 1$.

Next, suppose equations (6), (7) and (15) hold at time t. We now show they also hold at time $t+1$. Recall that agent i's extracted signals under X_t^μ and X_t^ν are respectively

$$\boldsymbol{\alpha}_t^{\mu,ij} = \boldsymbol{\beta}_t^{\mu,j} - \boldsymbol{\beta}_t^{\mu,ij}, \text{ and } \boldsymbol{\alpha}_t^{\nu,ij} = \boldsymbol{\beta}_t^{\nu,j} - \boldsymbol{\beta}_t^{\nu,ij}.$$

Further, by the induction hypothesis, from (6) and (7), we have:

$$\boldsymbol{\alpha}_t^{ij} = \boldsymbol{\beta}_t^{j} - \boldsymbol{\beta}_t^{ij} = \left(\boldsymbol{\beta}_t^{\mu,j} + \boldsymbol{\beta}_t^{\nu,j}\right) - \left(\boldsymbol{\beta}_t^{\mu,ij} + \boldsymbol{\beta}_t^{\nu,ij}\right) = \boldsymbol{\alpha}_t^{\mu,ij} + \boldsymbol{\alpha}_t^{\nu,ij}. \tag{16}$$

Since $\{x_t^{\mu,i}, x_t^{\nu,i}\} = \{x_t^i, x^0\}$, $\{\boldsymbol{\alpha}_t^{\mu,ii}, \boldsymbol{\alpha}_t^{\nu,ii}\} = \{\boldsymbol{\alpha}_t^{ii}, \mathbf{0}\}$ which implies $\boldsymbol{\alpha}_t^{ii} = \boldsymbol{\alpha}_t^{\mu,ii} + \boldsymbol{\alpha}_t^{\nu,ii}$.

$$\boldsymbol{\beta}_{t+1}^{i} = \boldsymbol{\beta}_t^{i} + \sum_{h \in g_i} \boldsymbol{\alpha}_t^{ih} = \boldsymbol{\beta}_t^{\mu,i} + \boldsymbol{\beta}_t^{\nu,i} + \sum_{h \in g_i} \left(\boldsymbol{\alpha}_t^{\mu,ih} + \boldsymbol{\alpha}_t^{\nu,ih}\right) = \boldsymbol{\beta}_{t+1}^{\mu,i} + \boldsymbol{\beta}_{t+1}^{\nu,i}.$$

The second equality holds by (6) and (16), and the last equality holds because it is expression (4) of the learning rule under X_t^μ and X_t^ν respectively. Thus, (6) 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 (16). Specifically,

$$\boldsymbol{\alpha}_t^{ijh} = \boldsymbol{\alpha}_t^{\mu,ijh} + \boldsymbol{\alpha}_t^{\nu,ijh} \text{ and } \boldsymbol{\alpha}_t^{ij\ldots lh} = \boldsymbol{\alpha}_t^{\mu,ij\ldots lh} + \boldsymbol{\alpha}_t^{\nu,ij\ldots lh}.$$

Then we can show that:

$$\boldsymbol{\beta}_{t+1}^{ij} = \boldsymbol{\beta}_t^{ij} + \sum_{h \in g_{ij}} \boldsymbol{\alpha}_t^{ijh} = \boldsymbol{\beta}_t^{\mu,ij} + \boldsymbol{\beta}_t^{\nu,ij} + \sum_{h \in g_{ij}} \left(\boldsymbol{\alpha}_t^{\mu,ijh} + \boldsymbol{\alpha}_t^{\nu,ijh}\right) = \boldsymbol{\beta}_{t+1}^{\mu,ij} + \boldsymbol{\beta}_{t+1}^{\nu,ij}.$$

In a similar way, we can show for all cliques $\{i, j, \ldots, l\}$, $\boldsymbol{\beta}_{t+1}^{ij\ldots l} = \boldsymbol{\beta}_{t+1}^{\mu,ij\ldots l} + \boldsymbol{\beta}_{t+1}^{\nu,ij\ldots l}$. Thus (7) and (15) also hold at time $t+1$. ∎

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

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

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

Proof of Proposition 2: We first show several 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, $(i_ki_{k+1}\dots i_hj_{h-1}\dots j_{k+1})$ 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.

Second, by Lemma 3 (which we will prove next), a social quilt contains no simple circles and satisfies LCS. We now show a property of social quilts which highlight the role of no simple circles. Specifically, if agent i's signal travels from agent l to k, and then extracted by h who connected to k but not 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, $d(ih) = d(ik) + 1$.

Next, we show that because a social quilt satisfies LCS, the agents' higher-order* estimates have cross-agent consistency, which is important for efficient learning.

LEMMA **5.** *For any agent $j \in \mathrm{N}_i$, $\boldsymbol{\beta}_t^{ij} = \boldsymbol{\beta}_t^{ji}$. Moreover, if (g_i, G_i) satisfies LCS$_i$, then for any clique $\{i, j, k, \ldots, l\}$,*

$$\boldsymbol{\beta}_t^{ij} = \boldsymbol{\beta}_t^{ik} = \ldots = \boldsymbol{\beta}_t^{il}, \text{ and } \boldsymbol{\beta}_t^{ij} = \boldsymbol{\beta}_t^{ijk} = \ldots = \boldsymbol{\beta}_t^{ijk\ldots l}. \tag{17}$$

Proof of Lemma 5: First, $\boldsymbol{\beta}_t^{ij} = \boldsymbol{\beta}_t^{ij}$ 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 are based on uninformative signals. Thus for any clique $\{i, j, k, \ldots, l\}$, $\boldsymbol{\beta}_1^{ij} = \boldsymbol{\beta}_1^{ik} = \boldsymbol{\beta}_1^{ijk} = \ldots = \boldsymbol{\beta}_1^{ijk\ldots l} = \mathbf{0}$.

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 $g_{ij} = g_{ik}$ for all k such that agent $\{i, j, k\}$ form a triangle. By the induction hypothesis, for any $h \in g_{ij} = g_{ik}$,

$$\boldsymbol{\alpha}_t^{ijh} = \boldsymbol{\beta}_t^{h} - \boldsymbol{\beta}_t^{ijh} = \boldsymbol{\beta}_t^{h} - \boldsymbol{\beta}_t^{ikh} = \boldsymbol{\alpha}_t^{ikh}.$$

Then, using expression (5), we have:

$$\boldsymbol{\beta}_{t+1}^{ij} \;\; = \;\; \boldsymbol{\beta}_t^{ij} + \sum_{h \in g_{ij}} \boldsymbol{\alpha}_t^{ijh} = \boldsymbol{\beta}_t^{ik} + \sum_{h \in g_{ik}} \boldsymbol{\alpha}_t^{ikh} = \boldsymbol{\beta}_{t+1}^{ik}.$$

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

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

Then, using expression (5),

$$\boldsymbol{\beta}_{t+1}^{ij} \;\; = \;\; \boldsymbol{\beta}_t^{ij} + \sum_{h \in g_{ij}} \boldsymbol{\alpha}_t^{ijh} = \boldsymbol{\beta}_t^{ijk\ldots l} + \sum_{h \in g_{ijk\ldots l}} \boldsymbol{\alpha}_t^{ijk\ldots lh} = \boldsymbol{\beta}_{t+1}^{ijk\ldots l}.$$

Thus, $\boldsymbol{\beta}_{t+1}^{ij} = \boldsymbol{\beta}_{t+1}^{ik} = \boldsymbol{\beta}_{t+1}^{ijk} = \ldots = \boldsymbol{\beta}_{t+1}^{ijk\ldots l}$. ▮

We now proceed to prove the proposition. By Lemma 1, if we can show the 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. 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 extracts the signal at $t = d(ih) + 1$ from this neighbor k (who can be agent i), and this is the only signal agent h extracts from his neighbors at any time. Specifically, for any $k' \in \mathrm{N}_h$ and any time t, $\boldsymbol{\alpha}_t^{hk'} = \boldsymbol{\alpha}_0^{ii}$ 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 $h \in \mathrm{N}_i$, then agent h extracts the signal from agent i's report $\boldsymbol{\beta}_1^i$ such that $\boldsymbol{\alpha}_1^{hi} = \boldsymbol{\alpha}_0^{ii}$. No other agents (including agent i) extract any new signal from their neighbors. On the other hand, if $\boldsymbol{\alpha}_1^{hk} = \boldsymbol{\alpha}_0^{ii}$, then clearly $k = i$ and $d(ik) = 0$, $d(ih) = 1$.

Next, suppose this holds at time t, we show that it also holds at time $t + 1$. First, if $\boldsymbol{\alpha}_t^{hk} = \boldsymbol{\alpha}_0^{ii}$ at time $t + 1$, then using the iterative relationship between extracted signals in equation (8) and the fact that the second term is zero by Lemma 5, 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 extract the signal from someone (say l) outside g_h in the previous period, so $hl \notin G$. By the induction hypothesis, since $\boldsymbol{\alpha}_{t-1}^{kl} = \boldsymbol{\alpha}_0^{ii}$, 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 $\boldsymbol{\alpha}_{t-1}^{kl} = \boldsymbol{\alpha}_0^{ii}$ 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, $\boldsymbol{\alpha}_t^{hk} = \boldsymbol{\alpha}_0^{ii}$. Thus $\boldsymbol{\alpha}_t^{hk} = \boldsymbol{\alpha}_0^{ii}$ if and only if $d(ih) = t$ and $d(ik) = t - 1$. Since agent h incorporates signal x_0^i exactly once at period $d(ih) + 1$, $\boldsymbol{\beta}_t^h = \boldsymbol{\alpha}_0^{ii}$ if $t > d(ih)$ and $\boldsymbol{\beta}_t^h = \mathbf{0}$ 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. To see this, note that Lemma 3 shows that when a network is not a social quilt, it must either contain simple circles or violate LCS. We show in Proposition 3 and 4 that both lead to learning errors. ∎

Proof of Lemma 3: 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 $j \in \mathrm{N}_i$, if there exist agents k and k' such that $k, k' \in \mathrm{N}_i \cap \mathrm{N}_j$, 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 at least three agents must be a clique. In a circle of three agents, a triangle, clearly all three agents are connected. So we start with any four-agent circle. 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. ∎

Proof of Proposition 3: 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 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 extracted signals are $\boldsymbol{\alpha}_1^{i_1 i} = \boldsymbol{\alpha}_1^{i_{k-1} i} = \boldsymbol{\alpha}_0^{ii}$. Recall that LCS holds, and thus the second term of the iterative relationship between extracted signals in equation (8) is zero. Also, by assumption, $\boldsymbol{\alpha}_t^{ll} = \mathbf{0}$ for any $t > 0$, $l \in g$. Then equation (8) can be rewritten as

$$\boldsymbol{\alpha}_{t+1}^{jh} = \sum_{l \in g_h \setminus g_j} \boldsymbol{\alpha}_t^{hl}. \tag{18}$$

At period $t = k+1$, the signal finishes traveling around the simple circle in both directions, and thus $\boldsymbol{\alpha}_k^{ii_{k-1}} = \boldsymbol{\alpha}_0^{ii}$ and $\boldsymbol{\alpha}_k^{ii_1} = \boldsymbol{\alpha}_0^{ii}$. 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 extract 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 extract 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 $\boldsymbol{\alpha}_t^{i_k j} = \boldsymbol{\alpha}_0^{ii}$ for some $j \in \mathrm{N}_{i_k}$. 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 $\boldsymbol{\alpha}_{t+1}^{i_{k-1} i_k} = \boldsymbol{\alpha}_0^{ii}$, 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 extract two more copies of x_0^i every k periods. Recall that D is the diameter of network. These newly extracted 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 $t \to \infty$. 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 $t \in \big[\tau(D + \lceil \kappa/2 \rceil) + 1, (\tau+1)(D + \lceil \kappa/2 \rceil)\big]$, any agent l in a simple circle believes there are at least two copies of x_0^i if $\tau = 1$; and at least

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

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 extracting more and more new signals from all the other $\kappa_{sc} - 1$ simple circles, and passing their own repeatedly extracted signals to them. As $t \to \infty$, 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 ($T < \infty$), and then let it go to infinity. 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:

$$\boldsymbol{\beta}_{T+D}^{i_k} = \sum_{l \in g, t \leq T} \left(\eta_{T+D}^{i_k}(x_t^l) \cdot \boldsymbol{\alpha}_t^{ll} \right). \tag{20}$$

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,

$$\boldsymbol{\beta}_{T+D+ok}^{i_k} = \sum_{l \in g, t \leq T} \left(\left(\eta_{T+D}^{i_\kappa}(x_t^l) + 2o \right) \cdot \boldsymbol{\alpha}_t^{ll} \right). \tag{21}$$

Given the agents' information structure, let $s^* = \arg\max_{s_n \in S} \Pr(s_n \mid X_T)$ since the probability that there are multiple states that maximizes $\Pr(s_n \mid X_T)$ is zero. Thus for any given T, as $o \to \infty$, the agents believe that only s^* can be the true state. 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 s^* with probability arbitrarily close to 1 as $t \to \infty$. When each agent in the network receives an infinite number of signals, by the Law of Large Numbers, $s^* = \arg\max_{s_n \in S} \Pr(s_n \mid X_T)$ is the true state if $T = \infty$.

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 ($T = \infty$). Let the true state be $s = s^*$. Recall that the set of all possible signals that agents can receive from nature is $X = \cup_i X^i$, which is randomly drawn by nature. Fix a (possibly large) value B; consider the set of all realizations of X such that $\Pr(s_n \mid x)/\Pr(s_{n'} \mid x) \leq B$ for all $x \in X$, $s_n \neq s_{n'}$. That is, for any signal $x \in X$, the ratio of the conditional probability of any pair of states is bounded by B. Denote this set as $\mathbb{X}$. Clearly, this set $\mathbb{X}$ occurs with a positive probability. We focus on the case that $X \in \mathbb{X}$ from now on. Given the agents' information structure, with probability 1, there exists a possible signal $x^{i,m}$ belonging to some agent i such that some other state $s' \neq s^*$ is the most likely state given $x^{i,m}$, that is, $s' = \arg\max_{s_n} \Pr(s_n \mid x^{i,m})$. Denote $x^{i,m}$ as x'. Clearly, $\Pr(s' \mid x') > \Pr(s^* \mid x')$.

Consider the following sequence of signals. Let nature send signal x' to agent i in every period from $t = 0$ to $t = t^*$ ($t^* \geq k$). Recall that the largest simple circle has k agents. This interval is set to insure that starting from some finite time, each simple circle receives new copies of x' from every other simple circle in every ensuing period. This interval also allows each signal x' to reach every other simple circle and travels back to the initial simple circle. It takes two steps to determine t^*. In the first step, we identify the integer k' such that

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

Here x^* is the signal most in favor of s^* relative to s'. To avoid carrying this likelihood ratio for the rest of the proof, for any signal x (or set of signals), we introduce

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

In the second step, we require that in each period from period $t^* - k$, the repetition must be strong enough such that every signal one simple circle extracts from any other simple circle includes at least $(2k + D + 1)Ik'$ copies of x' (excluding other later exogenous signals), where $I = |g|$ 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 extracted such a strong signal.

Next, we claim that regardless of the signals agents receive from nature after period t^*, all agents believe s' is increasingly more likely than s^* over time. That is, $\lim_{t\to\infty} \beta_t^h(s') - \beta_t^h(s^*) = \infty$ for all $h \in g$. We consider the signal one simple circle (for instance the largest one, $C = (i_1 i_2 \dots i_k)$) extracts 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,

$$\alpha_t^{i_1 j}(s') - \alpha_t^{i_1 j}(s^*) \geq \beta\left(s', s^* \mid (2k + D + 1)Ik' \text{ copies of } x'\right). \tag{23}$$

That is, the signal i_i extracts from j should favor s' over s^* by at least as many as $(2k + D + 1)Ik'$ copies of x' since period t^* (excluding other later exogenous signals).

Next, $\boldsymbol{\alpha}_t^{i_1 j}$ 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 period $t + 1$, using equation (8), agent i_2 extracts $\boldsymbol{\alpha}_{t+1}^{i_2 i_1}$ from agent i_1 such that

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

This is because agent i_1 gets fewer than I exogenous signals most favorable to s^* from 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

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

And the same is true for $\boldsymbol{\alpha}^{i_1 i_2}_{t+k}$. Use equation (8) again for the next period, we have

$$\alpha^{ji_1}_{t+k+1}(s') - \alpha^{ji_1}_{t+k+1}(s^*) \geq \beta\left(s', s^* \mid (2k+2D+1)Ik' \text{ copies of } x'\right).$$

That is, the signal agent j extracts from agent i_1 includes $\boldsymbol{\alpha}^{i_1 i_k}_{t+k}$ and $\boldsymbol{\alpha}^{i_1 i_2}_{t+k}$ (net of the exogenous signals reaching agent i_1 in time $t+k$). Then this signal $\boldsymbol{\alpha}^{ji_1}_{t+k+1}$ 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

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

This shows that the initial condition about the signal one simple circle extracts 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 extracted 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 $t \to \infty$.

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

Proof of Proposition 4: 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^l_0, which is the only informative signal. We can classify all agents based on their distance to l, that is, $\mathrm{N}^d_l = \{h \in g : d(lh) = d\}$, and $\mathrm{N}^1_l = \mathrm{N}_l$. To begin with, we claim that if agent a and $c \in \mathrm{N}^d_l$ are both linked to some agent h in N^{d+1}_l, then $ac \in G$. To see why, find a's connection to some agent f in N^{d-1}_l, then agent f and h are not linked, because their distance must be 2. Similarly the agent who is linked to c in N^{d-1}_l, say f', cannot be linked to h. If agent a and c are not linked, then there exists a simple circle consisting of agent f, a, h and c (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^d_l never extract new signals from their neighbors in N^{d+1}_l. Suppose to the contrary, the first time some agent extracts from her successor is agent a in N^d_l extracts a new signal from h in N^{d+1}_l. Notice that in the previous period, h does not extract new signal from her successors, so the new signal a extracts must come from h's neighbors in either N^d_l or N^{d+1}_l. Suppose that the new information a extracts comes from some c 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 extracts new information from h. The other possibility is that the new information a extracts 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 extracts 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 c in N_l^d. To make sure no simple circle exists, ch must be linked, so h would have learned it at the same time as h' from c. So we are back to the first possibility where the new information goes from c 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 $\boldsymbol{\beta}_t^l = \boldsymbol{\alpha}_0^{ll}$ (which reflects her initial signal x_0^l). Moreover, the estimates of agents in N_l must remain at $\boldsymbol{\alpha}_0^{ll}$. This is because, first, they cannot extract 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 extract new information from each other.

Lastly, we claim that there must exist some agent $l' \in \mathrm{N}_l^2$, who is linked to at least two agents in N_l but does not extract 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 $i, j \in \mathrm{N}_l$ and $k \in \mathrm{N}_l^2$ 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 extract 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 extracting 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 extracts new information from one of his peers. The first case is that he learns from agent $h \in \mathrm{N}_l^2$, whose new signal comes from some agent $j' \in \mathrm{N}_l$ 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 extracts 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 $h' \in \mathrm{N}_l^2$. 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 $j' \in \mathrm{N}_l$. 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 to believe in a unique state with probability 1. By assumption, the signal x_0^l uniquely favors one state almost surely, and the Bayesian posterior under an arbitrarily large number of x_0^l 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 $T = \infty$. That is, $x_\tau^l = x_0^l$ for all even τ. Recall from above that each such signal x_τ^l makes some agent k^* extracts 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 $t \to \infty$. ▮