Learning in Local Networks*

Wei Li
University of British Columbia

Xu Tan
University of Washington

December 9, 2016

Abstract

Agents in a network want to learn the true state of the world from their own signals and reports from their neighbors. Each agent only knows her local network, consisting of her immediate neighbors and any connections among them. In each period, every agent updates her own estimates about the state distribution based on perceived new information. She also forms estimates about each neighbor's estimates given the new information she thinks the neighbor has received. Whenever a neighbor's report differs from what the agent thinks he should report, the agent attributes the difference to new information. Agents learn correctly in any network if their information structures are partitional. They can also do so for more general information structures if the network is a social quilt, a tree-like union of fully connected subnetworks. Otherwise, agents may fail to learn despite an arbitrarily large number of correct signals.

JEL: D03, D83, D85

Keywords: Bayesian learning, local networks, iterated learning procedure.

*Preliminary draft, all comments welcome. We thank Andreas Blume, Vitor Farinha Luz, Drew Fudenberg, Fahad Khalil, Jacques Lawarrée, Stephen Morris, Edward Schlee, Wing Suen, Joel Watson, Quan Wen, and especially Matt Jackson, Li, Hao and Mike Peters for insightful comments and extensive discussions. We also thank the seminar participants at University of British Columbia, University of Washington, Peking University, University of International Business and Economics, Shanghai University of Finance and Economics, Hong Kong University, Emory University, Simon Fraser University, University of New South Wales, University of Technology Sydney, the Second Annual Conference on Network Science and Economics, the 2016 Canadian Economic Theory Conference and the 2016 North American Summer Meeting of the Econometric Society for great feedback. Wei Li thanks the Hampton Established Scholar Grant for financial support.

Email: wei.li@ubc.ca; tanxu@uw.edu.

1 Introduction

We often learn from those we interact with, who in turn talk to and learn from their neighbors, and so on.[1] To use information learned from networks, we need to account for possible repetitions and distortions. Learning errors may contribute to entrenched poverty, political polarization, and other financial and personal failures. Consider, for instance, the plight of those who live in a poor, isolated neighborhood. Lack of good information of the wider network means that a poor kid may not bother to try, because he believes that the underlying true state is so unfavorable—"the system is against us"—that the chance of success is negligible.[2] He learned this from his neighbors, who in turn reached such a conclusion because they may have failed themselves, or they know a neighbor who has a neighbor who tried and failed. A few instances of failures, however, may reach the poor kid through multiple neighbors via different paths. Because he only knows his immediate neighbors, he fails to account for the repetition in information. Consequently he believes erroneously—and increasingly so if the same information travels back to him again—that the underlying true state is far less favorable than it actually is.

Systematic learning errors are not uncommon. In recent years, mounting evidence from the lab and field has shown that people often struggle with distinguishing new information from old, existing ones when they learn from their neighbors. Subjects often err by treating correlated information as independent (see Chandrasekhar, Larreguy and Xandri 2012, Grimm and Mengel 2014, Enke and Zimmermann 2015, Golosov, Qian and Kai 2015, among others). For instance, they may double-count information when they fail to realize that their neighbors have learned from the same sources farther away in the network. Along with evidence from psychology and sociology, failure to fully account for correlation in information from one's network is a robust feature of social learning. But it is worth noting that this body of research also finds that people are not oblivious to the problem. Some of them, especially those who perform well in the experiments, try to reduce the severity of

[1]There is a vast empirical literature showing that we learn from our social networks. In technology adoption, Conley and Udry (2001) show that pineapple farmers in Ghana will begin to use more fertilizer after a neighbor uses high amounts of fertilizer and achieves surprisingly high profits. In personal finance, Duflo and Saez (2002) show that participation in retirement savings plans by employees in a university is strongly influenced by their peer groups (along gender, service, status, age lines). In financial market, Hong, Kubik and Stein (2005) show that a mutual fund manager is more likely to buy (or sell) a particular stock in any quarter if other managers in the same city are buying (or selling) that same stock.

[2]Wilson, Quane and Rankin (1998) show that, using data from Chicago inner-city residents, low social-economic status residents of ghetto neighborhoods know almost two fewer employed people, one fewer college educated person, and nearly three more welfare recipients in their social network than those in the low-poverty neighborhoods. More ominously, "only 61% of the youth in ghetto neighborhoods reported the most of their friends attended school regularly, compared to 89% in low-poverty neighborhoods." See also Mobius and Rosenblat (2001) and Ioannides and Loury (2004) for more discussions.

this problem in rudimentary ways (see Celen and Kariv 2004, Grimm and Mengel 2014).

In this paper, we develop a tractable learning procedure incorporating several realistic features of how people actually learn. First, agents do not need to know the structure of the entire network they are in. Specifically, we do not assume agents have common knowledge of their network; nor do we assume they have common prior over the network. Instead, each agent only needs to know her *local network*, consisting of her immediate neighbors and the links among them.[3] Second, each agent does not need to know any neighbor's information structure, which is that neighbor's private information.

Specifically, there are finitely many states, and agents want to learn the true state. Each agent learns by forming and updating her estimates—what she thinks is the most up-to-date distribution of the states given available information.[4] Time is discrete. In the initial period, each agent receives a signal. In every ensuing period, each agent first forms her current estimates, and then simultaneously reports them to all her neighbors.[5] Each agent infers the new information contained in her neighbors' reports, and then receives a signal from the nature. Using all the inferred information, each agent updates her estimates in the next period. The innovation of our model lies in how each agent infers the new information using second (and higher)-order estimates. That is, she forms estimates of a neighbor's estimates of the state distribution based on the reports they both have observed, which she thinks is *all* the information that neighbor has learned so far. Similarly, she forms estimates of one neighbor's estimates of another's estimates based on the reports all three have observed, and so on. Because the agent thinks she has accounted for all sources of a neighbor's information, if any neighbor's reported estimates in a period differ from her estimates of that neighbor's estimates, she attributes the difference to a "new" signal. This inferred signal is a composite of the neighbor's new signal last period and any information he has received from his local network unbeknownst to her. This procedure continues iteratively until no one learns anything new.

An important feature of the above procedure is that agents avoid repetition *within their local networks*, but they behave *as if* all the information from outside their local networks is new and independent. This reflects the lab findings that people may be able to reduce repetition when it is easy to detect, but they neglect the correlation of information otherwise.

[3]This is consistent with evidence from surveys on subjects' knowledge of the network. For instance, Krackhardt (1990) finds that the accuracy of knowing other people's connections is between 15% and 48% in a small startup with 36 people, Casciaro (1998) finds that the accuracy is around 45% in a research center of only 25 people, and Chandrasekhar, Breza and Tahbaz-Salehi (2016) find similar patterns.

[4]Because our agents do not form and update their beliefs about the network beyond their local networks, we use estimates to differentiate them from the standard Bayesian beliefs. They are identical for an agent whose local network is the entire network, or who believes so (see Section 4.4 for further details).

[5]We refer to the generic agent as "she" and each of her neighbors as "he."

To fix ideas, suppose agent 1, 2 and 3 are connected in a triangle. If agent 1 infers a new signal from agent 2's report, she expects agent 3 to infer the same signal, and she knows that agent 3 expects her to do the same. Therefore neither agent 1 or 3 double counts agent 2's information. Similarly when she infers a new signal from agent 3. If the triangle is the entire network, then agent 2 and 3's new information can only come from the nature, and are thus truly independent. If instead, agent 2 and 3 are linked to another agent not observable to agent 1, then agent 1 still treats the inferred signals from agent 2 and agent 3 as independent. But they are actually correlated if 2 and 3 learn the information from their common neighbor. Our way of modeling reflects the heavy burden agents face to properly account for correlations in the information they receive when they don't know the network. To do so, an agent has to first form beliefs about the total number of agents in the network and consider every possible network for each given number.[6] For each given network, she then assigns probabilities to all the possible signals and travel paths through which each signal may have traveled to reach her. Moreover, she needs to update all these beliefs as she receives more information. In contrast, agents can follow our procedure easily when they do not have the computational or cognitive ability to undertake these calculations, or when they find such endeavors too costly.[7] Viewed in this light, we provide a lower bound on how well agents can learn given such limited information.

Our first contribution is to show this procedure has several useful properties, making it a potentially portable component for other network models. First, each agent only needs to form estimates up to the order of the number of agents in the largest fully connected subset of her local network, which can be far lower than the number of her neighbors. In fact, in many commonly studied networks such as trees, circles and cliques, second-order estimates suffice.[8] Thus our procedure is far simpler than it appears. Second, signals can be decomposed. That is, fix the sequence of realized signals a network receives, the agents' learning outcomes are the same if we analyze their estimates under each signal separately and then combine them by Bayes' rule. This implies that if the agents can learn each signal correctly, they can learn correctly for all of them.

Our second contribution contains two positive results characterizing when the agents can learn correctly—their estimates agree with the correct Bayesian posterior given all the signals the network receives—under this procedure despite limited knowledge of the network. First, agents can do so in *any* network when their information structures are partitional as

[6]For example, in a L-agent undirected network, there are $L(L-1)/2$ number of possible links. Because each link may or may not exist for a given network, the number of total possible networks is $2^{L(L-1)/2}$.

[7]In fact, subjects in the lab are unable to make inferences correctly even when the network is small and commonly known. We further discuss and generalize this feature in Section 2.2 and Section 5.3.

[8]A clique is a fully connected network in which every pair of agents are connected.

in Aumann (1976). Since each agent's signal informs her the element of the partition the true state is in, treating correlated information as independent does not lead to learning errors. Suppose an agent infers the same new signals from her neighbors who have learned the information from a common source. The agent's estimates are unaffected in that she eliminates the same set of states as she would have with one such signal. Not only the agents cannot disagree forever, they agree as soon as every agent has learned all the signals. They believe all the states in the intersection of every agent's elements of partition are equally likely to be the true state, and assign zero probability to the rest.

For more general information structures, our second positive result is that agents learn correctly if the network is a social quilt, a tree-like union of cliques.[9] Because each clique is fully connected, when information arrives at one member (say agent i) of the clique, all other members can identify this as new information. More importantly, all members correctly expect that all the others in the clique have learned this information from i. Thus they avoid learning the same information repeatedly. Moreover, the overall tree structure of a social quilt ensures that each signal travels through the network once and only once. Thus agents form the correct estimates after learning all the signals.

Next, we show what may be viewed as a negative result: the agents fail to learn correctly for some sequence of realized signals if the network is not a social quilt. Since many networks in reality are not social quilts, our third contribution is to identify two network features leading to correlated inferred signals. First, a *simple circle* contains at least four agents, each of whom has exactly two links, one to the neighbor on each side. Consider one with four agents $1, 2, 3$ and 4. Agent 1 receives the only signal. Agent 2 and 4 learn it first, and then agent 3, not knowing the existence of agent 1, must think that there are two copies of the signal. More importantly, she passes her estimates on to agent 2 and 4, who know 3 learned one copy from themselves and infer the other copy as new signal. In the end, everyone thinks they have learned an infinite number of the signal. The problem is exacerbated with multiple simple circles—the number of repeated inferred signals grows at an exponential rate—and the arrival time of each signal matters. In fact, the Law of Large Numbers may fail in that all agents believe in a wrong initial signal even if they receive an arbitrarily large number of correct signals later.

In a *diamond with a link*, one pair of non-adjacent agents in a four-agent simple circle is connected, say agent 2 and 4 in the example above. Then, agent 3's inferred signal may be

[9] Empirical analysis found cliques and clustering—a measure of the likelihood that one agent's neighbors are connected with one another—to be much higher than that predicted in a random network (see for example MacRae (1960), Adamic (1999), and Goyal, van der Leij and Moraga-Gonzlez (2006)). In this context, our paper joins the recent attempts to provide micro foundations for locally tightly connected subgroups such as Jackson, Rodriguez-Barraquer and Tan (2012) and Ali and Miller (2013).

negatively correlated with agent 1's original signal. Because agent 2 and 4 know each other has learned the signal from agent 1, they do not attribute each other's changed estimates to new information. But in addition to inferring two copies of this signal, agent 3 also expects 2 and 4 to infer a copy from each other. Since 2 and 4 do not change their estimates, agent 3 infers that each has received an opposite, offsetting signal. She incorporates these two "new" offsetting signals and her estimates return to her prior beliefs. Her estimates oscillate forever between her prior and the posterior given two copies of 1's signal. We show that some agent's learning outcomes can not be Bayesian for some sequence of realized signals.

We extend our model in several directions. First, we expand each agent's local network to all her neighbors within a certain distance (and the links among them). Surprisingly, knowing more about one's indirect neighbors may not help because an agent may infer more copies of the same signal, not knowing all these direct and indirect neighbors have learned from one common source. Second, agents may weigh their neighbors' reports differently. In this case, agents may disagree when they stop learning, and polarization of opinions can appear among agents further apart. In a third extension, we let agents account for some correlated information from outside their local networks in a simple way. Namely, each agent maintains a list of signals already inferred. Whenever she infers any signal that is identical to one of her stored signals, she thinks that the signal has been duplicated somewhere in the network and treats it as old information. Following this procedure, agents can learn correctly in any network if all signals reach the same agent initially. But doing so cannot completely remove information repetition in general.

A vast theoretical literature has studied the question of learning in social networks. One strand of the existing literature shows that agents can form the correct Bayesian beliefs (asymptotically) if everyone knows the structure of their social network (see Gale and Kariv (2003), Mueller-Frank (2013), Mossel, Sly and Tamuz (2015), among many others), or if the agents can communicate in complex ways.[10] Indeed, in our model, if the agents have common knowledge of the network, then they form the correct Bayesian beliefs within a finite number of periods because they can account for any correlation in their information eventually. The other strand of the literature does not assume common knowledge of the network. It eschews the complexity of Bayesian learning by assuming that agents learn by following reasonable rules of the thumb (see DeGroot (1974), DeMarzo, Vayanos and Zwiebel (2003), Golub and Jackson (2010), Molavi, Tahbaz-Salehi and Jadbabaie (2016), among many others.) In Ellison and Fudenberg (1993, 1995), agents use only currently available social information

[10] For instance, they can report the sources and travel paths of all information as in the tagged information system of Acemoglu, Bimpikis and Ozdaglar (2014). But agents are unlikely to communicate in such a complicated way in real life networks.

and disregard historical data (including their own past experience) in making decisions. The classic model of DeGroot (1974) assumes that agents treat their neighbors' information in each period as new and incorporate it into their beliefs. Often called myopic (or naive) learning, this influential model may feature high levels of information repetition, because the same information is repeated even within one's local network.

More closely related, several papers consider *semi-Bayesian* learning in networks. In Bala and Goyal (1998), each agent chooses an action not knowing the true payoffs of her action. She observes an outcome every period and then updates her belief about the optimal action rationally based on the outcomes in her local network, but she does not infer information from the choices of her neighbors. They focus on the long-run convergence in the network, while we explicitly model how agents infer new information from neighbors and characterize how their learning depends on the network structure. In Alatas et al. (2016), agents report their most recent signal to their neighbors, and everyone knows their social distance to others in the network. Each agent treats all signals received as independent. Our agents know less about the network and they account for correlations of signals within their local network.

Our paper is also related to the large literature on herding and observational learning (Banerjee (1992), Bikhchandani, Hirshleifer and Welch (1992), Lee (1993), Lones and Sorensen (2000), Acemoglu et al. (2011), Harel et al. (2014), among many others). The classical papers on herding and information cascades focus on a directed, linear chain, in which every agent observes all predecessors' actions and chooses her only action optimally given her observations and her own private signal. If the agents communicate their posterior distributions as in our model, it is easy to see that they learn correctly since the linear chain is a special case of a social quilt. More closely related to our paper, Eyster and Rabin (2014) consider both a rich message space and a non-linear network structure. Their agents can directly observe the actions of some of their predecessors, who in turn observe the actions of some of their predecessors. One of their results shows that, similar to our model, all rational agents can learn correctly if the implied network structure for all agents are such that the aforementioned information repetition problem does not arise.[11]

Many papers use experiments to study observational learning in networks (see Anderson and Holt (1997), Celen and Kariv (2004), Alevy, Haigh and List (2007), Cai, Chen and Fang (2009), Mobius, Phan and Szeidl (2015) among many others). More recently, Chandrasekhar, Larreguy and Xandri (2012) compare Bayesian learning with myopic learning when the

[11] They address the information repetition problem by assuming a "shield" structure. That is, if agent k observes actions of agent i and j, who share a common predecessor k', then agent k must directly observe the action of k'. Therefore agent k knows what agent i and j learn from k', and will not double count it. Applying their "shield" structure for directed communication to our two-way communication, i, j, k, k' must form a clique.

network is common knowledge. They find that while the myopic learning model performs better, it can only explain 76% of the actions taken. Grimm and Mengel (2014) show that while myopic learning fits better in some information environment, their subjects do account for correlated information by putting a lower weight on their neighbors' information when it is expected to be highly correlated.

We introduce our learning procedure in Section 2 and characterize its properties in Section 3. Section 4 shows when agents can learn correctly, and the types of learning errors when they cannot. Section 5 extends our model and Section 6 concludes. All proofs are contained in the Appendix.

2 Model

Consider a network (g, G): $g = \{1, 2, \ldots, L\}$ 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. The network is undirected, so information flows both ways. That is, $ij \in G$ if and only if $ji \in G$. It is also path-connected so that information can diffuse to everyone. That is, for any $i, h \in g$, there is a path $\{i_0 i_1 \ldots i_k\}$ such that $i_0 = i$, $i_k = h$ and $i_l i_{l+1} \in G$ for all $l < k$. Denote $\mathrm{N}_i = \{j : ij \in G\}$ as the set of agent i's neighbors. Let $(g_i, G_i) \subseteq (g, G)$ be agent i's *local network*, where $g_i = \mathrm{N}_i \cup \{i\}$ and $G_i = \{jk : j, k \in g_i \text{ and } jk \in G\}$. That is, her local network consists of herself, all her neighbors, and all the links among them in the original network. Also, let (g_{ij}, G_{ij}) be agent i and a neighbor j's *shared local network*, where $g_{ij} = g_i \cap g_j$ and $G_{ij} = G_i \cap G_j$. That is, their shared network consists of the agents, their common neighbors, and all the links among them. Similarly, Let $\{i, i_1, \ldots, i_l\}$[12] be a fully connected subset of g_i such that all the agents in the subset are distinct, and every pair of agents in $\{i, i_1, \ldots, i_l\}$ are linked. Let $(g_{ii_1\ldots i_l}, G_{ii_1\ldots i_l})$ be the shared local network of agents $\{i, i_1, \ldots, i_l\}$, where $g_{ii_1\ldots i_l} = g_i \cap g_{i_1} \cap \ldots \cap g_{i_l}$, and $G_{ii_1\ldots i_l} = G_i \cap G_{i_1} \cap \ldots \cap G_{i_l}$.

Agents in the network face a learning problem. There is a finite set of states: $s \in S = \{s_1, s_2, s_3, ..., s_n, ..., s_N\}$. All the states are a priori equally likely: $\Pr(s = s_n) = 1/N$ for all s_n. The true state is realized before learning begins. Agents receive signals from nature about the state of the world. The support of each agent i's signals is also finite and includes an uninformative signal. Specifically, $x^i \in X^i = \{x^i_\emptyset, x^i_1, x^i_2, \ldots, x^i_m, ..., x^i_{M_i}\}$, where $x^i_\emptyset$ is uninformative. Let the conditional probability of receiving signal x^i_m if the state is s_n be $\phi^i_{mn} = \Pr(x^i_m | s_n)$. Time is discrete with an infinite horizon: $t = 0, 1, 2, \ldots$. In each period

[12] To clarify, we use $\{ij \ldots k\}$ to denote a sequence of agents in which the order matters such as in a path, and we use $\{i, j, \ldots, k\}$ to denote a set of agents in which the order does not matter.

prior to T, which can be arbitrarily large, agent i observes a signal x_t^i.[13] For simplicity, we assume all agent i's realized signals are generated according to the same signal distribution defined above. The signals are independent across agents and time conditional on the state. No informative signal arrives at or after period T.[14]

Let X_T^i be the signals agent i receives up to T, and let $X_T = \bigcup_{i \in g} X_T^i$. An agent's learning outcomes are *Bayesian* if her posterior belief of the state distribution at the end of the learning process is the Bayesian posterior given X_T. This paper aims to show whether and when agents can achieve the Bayesian learning outcomes when each agent only knows her local network (g_i, G_i), her information structure, and her realized signals.

2.1 An iterative learning procedure

We begin with a formal description of our learning procedure and defer discussions of this procedure to Section 2.2. Every agent $i \in g$ learns according to the following procedure.

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

At $t = 1$, agent i's first-order estimates of the state distribution are $\mathbf{p}_1^i = \{p_1^i(1), \ldots, p_1^i(N)\}$, where $p_t^i(n) = \Pr(s_n | x_0^i)$ for each s_n.[15] By assumption, her second-order estimates of each neighbor $j \in \mathrm{N}_i$'s estimates of the state distribution are $\mathbf{p}_1^{ij} = \{p_1^{ij}(1), \ldots, p_1^{ij}(N)\} = \{1/N, \ldots, 1/N\}$. Similarly, if $\{i, i_1, i_2\}$ are fully connected, then i's estimates of agent i_1's estimates of agent i_2's estimates of the state distribution $\mathbf{p}_1^{ii_1i_2} = \{1/N, \ldots, 1/N\}$, and so on for all her higher-order estimates in each of the fully connected subsets of g_i. Agent i then reports her first-order estimates $\mathbf{p}_1^i$ to all her neighbors and simultaneously hears all her neighbors' reports $\mathbf{p}_1^j$. She then observes her signal from nature x_1^i. Period 1 ends.

For all $t \geq 1$, agent i first forms all her (higher-order) estimates. She then reports her first-order estimates $\mathbf{p}_t^i$ and simultaneously hears all her neighbors' reports $\mathbf{p}_t^j$. Next, she receives x_t^i and period t ends. She updates all her estimates in $t+1$ in three steps:

Step 1: Identify new information. She first infers new information each neighbor j has learned from outside her local network during period $t-1$. Let y_{t-1}^{ij} be an *inferred*

[13] To ease exposition, we abuse the notation between possible and realized signals, which should be clear from the context. More specifically, x_m^i is agent i's m-th informative signal only if the subscript is m, otherwise x_t^i is agent i's realized signal at period t.

[14] We introduce T so that there are only a finite number of informative signals, and thus we can study whether the agents' learning stops and characterize the outcomes when it does. We assume agents don't know the existence of T for simplicity. Otherwise agents should increasingly believe that all information learned at a time sufficiently later than T contains old signals only.

[15] We use boldface letters to denote vectors throughout the paper.

signal agent i thinks j has learned. Formally, from agent i's perspective, for each state s_n,

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

That is, the numerator is the joint probability agent i believes agent j receives y_{t-1}^{ij} and the state is s_n; and the denominator is the total probability agent i believes agent j receives signal y_{t-1}^{ij}. Let $\boldsymbol{\alpha}_t^{ij} = \{\alpha_t^{ij}(1), \ldots, \alpha_t^{ij}(N)\}$ denote the (normalized) distribution of y_{t-1}^{ij} conditional on the state, where $\alpha_t^{ij}(n) = \Pr\left(y_{t-1}^{ij}|s_n\right) / \sum_{n'} \Pr\left(y_{t-1}^{ij}|s_n'\right)$. It is easy to see that

$$\alpha_t^{ij}(n) = \frac{p_t^j(n)}{p_t^{ij}(n)} \Bigg/ \sum_{n'} \frac{p_t^j(n')}{p_t^{ij}(n')}. \tag{2}$$

Clearly, $\sum_{n'} \alpha_t^{ij}(n') = 1$, and $\alpha_t^{ij}(n) = 1/N$ for all s_n if $\mathbf{p}_t^j = \mathbf{p}_t^{ij}$ (the equality holds component wise).[16] Similarly, the new information agent i believes that agent j infers from herself and another agent $k \in \mathrm{N}_i \cap \mathrm{N}_j$ are respectively:

$$\alpha_t^{iji}(n) = \frac{p_t^i(n)}{p_t^{iji}(n)} \Bigg/ \sum_{n'} \frac{p_t^i(n')}{p_t^{iji}(n')} \text{ and } \alpha_t^{ijk}(n) = \frac{p_t^k(n)}{p_t^{ijk}(n)} \Bigg/ \sum_{n'} \frac{p_t^i(n')}{p_t^{ijk}(n')}. \tag{3}$$

The higher-order new information is defined similarly. For each fully connected subset $\{i, i_1, \ldots, i_{l-1}, i_l\}$ of g_i, agent i believes that agent i_1 believes...that agent i_{l-1} infers $\boldsymbol{\alpha}_t^{ii_1\ldots i_l} \equiv \{\alpha_t^{ii_1\ldots i_l}(1), \ldots, \alpha_t^{ii_1\ldots i_l}(N)\}$ from agent i_l, where

$$\alpha_t^{ii_1\ldots i_{l-1}i_l}(n) = \frac{p_t^{i_l}(n)}{p_t^{ii_1\ldots i_{l-1}i_l}(n)} \Bigg/ \sum_{n'} \frac{p_t^{i_l}(n')}{p_t^{ii_1\ldots i_{l-1}i_l}(n')}. \tag{4}$$

Finally, $\boldsymbol{\alpha}_t^{ii_1\ldots i_l k}$, where $k \in \{i, i_1, \ldots, i_l\}$ is similarly defined. Let the number of agents in the largest fully connected subset of g_i be $\widehat{L}_i$. Then the highest-order new information agent i infers is of order $\widehat{L}_i + 1$.[17]

Step 2: Update first-order estimates. In period t, agent i may observe an informative new signal $x_t^i = x_m^i$. To simplify the notation, let $\alpha_t^{ii}(n) = \phi_{mn}^i / (\sum_{n'} \phi_{mn'}^i)$ be agent i's inferred signal from nature such that $\boldsymbol{\alpha}_t^{ii} \equiv \{\alpha_t^{ii}(1), ..., \alpha_t^{ii}(N)\}$ for any $t \geq 0$.[18] Using her inferred signals from each neighbor j and her own signal, agent i updates her first-order

[16] We show in Section 2.2 why $\boldsymbol{\alpha}_t^{ij}$ is the part of the inferred signal that can be learned by agent i. Also, if $p_t^{ij}(n) = 0$ for some state s_n, then $p_{t-1}^h(n)$ must be 0 for some $h \in g_{ij}$, which implies that $p_\tau^j(n)$ must be 0 for all $\tau \geq t$. In this case, we define the ratio $p_\tau^j(n)/p_\tau^{ij}(n) = 0$.

[17] The only repetition allowed in agent i's higher-order new information is the last agent.

[18] Observe that despite the same time index, $\boldsymbol{\alpha}_t^{ii}$ for all $t \geq 0$ reflects agent i's realized signal at t, but $\boldsymbol{\alpha}_t^{ij}$ for all $t \geq 1$ is the signal agent i thinks agent j has learned at $t-1$.

estimates such that for each s_n:

$$p_{t+1}^{i}(n)=\frac{p_t^i(n)\prod_{h\in g_i}\alpha_t^{ih}(n)}{\sum_{n'=1}^{N}p_t^i(n')\prod_{h\in g_i}\alpha_t^{ih}(n')}. \tag{5}$$

Step 3: Update higher-order estimates. When updating her estimates of neighbor j's estimates, agent i starts with agent j's latest report $\mathbf{p}_t^j$. She then incorporates the new information she thinks that j learned from their shared local network. Agent i's second-order estimates are formed by Bayes' rule such that for each s_n:

$$p_{t+1}^{ij}(n)=\frac{p_t^j(n)\prod_{h\in(g_{ij}\setminus\{j\})}\alpha_t^{ijh}(n)}{\sum_{n'=1}^{N}p_t^j(n')\prod_{h\in(g_{ij}\setminus\{j\})}\alpha_t^{ijh}(n')}. \tag{6}$$

Agent i's higher-order estimates are formed similarly:

$$p_{t+1}^{ii_1\dots i_l}(n)=\frac{p_t^{i_l}(n)\prod_{h\in\left(g_{ii_1\dots i_l}\setminus\{i_l\}\right)}\alpha_t^{ii_1\dots i_l h}(n)}{\sum_{n'=1}^{N}p_t^{i_l}(n')\prod_{h\in\left(g_{ii_1\dots i_l}\setminus\{i_l\}\right)}\alpha_t^{ii_1\dots i_l h}(n')}, \tag{7}$$

for all sequences of fully connected and distinct agents $\{ii_1\dots i_l\}$. The highest order of estimates agent i needs to form is thus $\widehat{L}_i$.

Finally, let $\mathbf{p}_{t+1}^{ii_1\dots i_l k}=\mathbf{p}_{t+1}^{ii_1\dots i_l}$ for all $k\in\{i,i_1,\dots,i_l\}$.[19] These are degenerate estimates because they are set to be equal to the estimates one order lower. Agent i uses them to infer her $\widehat{L}_i+1$-order new information in period $t+2$.

Learning stops if no agent in the network infers any (higher-order) new information from the next period onward. ‖

To be concrete, we illustrate the above procedure with a simple example.

Example **1.** *The network (g,G) is a triangle: $G=\{12,13,23\}$. Suppose both the states and the signals are binary: $s\in S=\{0,1\}$, $x^1\in X^1=\{0,1\}$, and $\Pr(x^1=1|s=1)=\Pr(x^1=0|s=0)=\phi^1$. Agent 1 receives the only informative signal $x_0^1=1$.*

At $t=0$, agent 1 observes $x_0^1=1$. We describe only agent 1's learning in details first.

At $t=1$, $p_1^1(1)=\phi^1$ by Bayes' rule.[20] Agent 1's higher-order estimates are the symmetric prior by assumption. Agent 2 and 3 report $p_1^2(1)=p_1^3(1)=\frac{1}{2}$, and their higher-order estimates are also the symmetric prior.

At the beginning of $t=2$:

[19]We show in Appendix A.1 that there is no loss in generality to restrict attention to distinct agents only. That is, for any sequence of fully connected and possibly repeated agents $\{ij\dots k\}$ such that $\{i,i_1,\dots,i_l\}$ is the set of all the distinct agents it contains, we can set $\mathbf{p}_{t+1}^{ij\dots k}=\mathbf{p}_{t+1}^{ii_1\dots i_l}$.

[20]Because the state is binary, we only keep track of $p_t^i(1)$.

Step 1: given the reports at $t = 1$, agent 1 notices that $p_1^2(1) = p_1^{12}(1)$ and $p_1^3(1) = p_1^{13}(1)$, and thus $\alpha_1^{12}(1) = \alpha_1^{13}(1) = \frac{1}{2}$. That is, agent 1 does not learn any new information from 2 or 3. Next, agent 1 believes that 2 and 3 learn new information from herself: $\alpha_1^{121}(1) = \alpha_1^{131}(1) = \phi^1$. But she believes agent 2 does not infer any new information from 3, or vice versa: $\alpha_1^{132}(1) = \alpha_1^{123}(1) = \frac{1}{2}$. Finally, because the largest fully connected subset of g_1 consists of $\{1, 2, 3\}$, agent 1 also infers $\alpha_1^{1231}(1) = \alpha_1^{1321}(1) = \phi^1$ and $\alpha_1^{1232}(1) = \alpha_1^{1323}(1) = \frac{1}{2}$.

Step 2: by expression (5), $p_2^1(1) = p_1^1(1) = \phi^1$.

Step 3: by expression (6), we have $p_2^{12}(1) = p_2^{13}(1) = \phi^1$. Using expression (7), agent 1's third-order estimates are $p_1^{123}(1) = p_1^{132}(1) = \phi^1$. The highest-order estimates agent 1 forms are $p_1^{1231}(1) = p_1^{1232}(1) = p_1^{1321}(1) = p_1^{1323}(1) = p_1^{123}(1) = \phi^1$.

Agent 2 and 3 are symmetric, and thus we only describe how agent 2 learns. Given the reports at $t = 1$. agent 2 notices that $p_1^1(1) \neq p_1^{21}(1)$. By expression (3), $\alpha_1^{21}(1) = \phi^1$. Agent 2 learns nothing from agent 3: $\alpha_1^{23}(1) = \frac{1}{2}$. By expression (5), $p_2^2(1) = \phi^1$. Next, it is easy to see that $\alpha_1^{213}(1) = \frac{1}{2}$, $\alpha_1^{212}(1) = \frac{1}{2}$, and thus by expression (6), $p_2^{21}(1) = \phi^1$, $p_2^{23}(1) = \phi^1$. Finally, it is easy to see that $p_2^{213}(1) = \phi^1$ and $p_2^{231}(1) = \phi^1$. The highest-order estimates agent 2 forms is $p_2^{2131}(1) = p_2^{213}(1) = \phi^1$, and so are all her other fourth-order estimates.

At $t = 2$, every agent's first and higher-order estimates that the state is 1 are ϕ^1. No one learns anything or thinks that anyone else learns anything new from then on, and thus learning stops. Their learning outcomes are Bayesian. $\diamond$

2.2 Remarks on the model

Inferred signals. Each agent wants to learn the true state based on information available to her. She can use our procedure when she does not know the outside network, her neighbors' information structures or their realized signals. To do so, agent i infers and learns new information from each neighbor j based on the reports from their shared local network. To interpret expression (1), agent i uses her own estimates of j's estimates of the state distribution as her prior, and agent j's actual report as her posterior in her inference of the new signal. She then applies the Bayes' rule *as if* agent j's new information comes from one new signal, whether it is from nature or from agents not connected to i. Because agent i knows neither the distribution of j's signals nor j's local network, she cannot differentiate the sources of j's information. Therefore $\boldsymbol{\alpha}_t^{ij}$ is the part of the inferred signal relevant to agent i's learning. She then updates her own estimates $\mathbf{p}_{t+1}^i$ using all these inferred signals.

The innovation of our procedure is that each agent uses higher-order estimates to keep track of old and existing information in her local network. To see this, observe that $\mathbf{p}_t^{ij}$ contains all the information agent i believes that agent j has, namely, $\{\mathbf{p}_\tau^h : \tau \leq t-1, h \in g_{ij}\}$.

Therefore, she only updates her estimates about j's estimates based on new reports at period t. This is clearly true at $t = 1$. Suppose this is true at t. Then at $t + 1$, using expression (1), we can rewrite expression (6) such that:

$$p_{t+1}^{ij}(n) = \frac{p_t^{ij}(n)\alpha_t^{ij}(n)\prod_{h\in(g_{ij}\setminus\{j\})}\alpha_t^{ijh}(n)}{\sum_{n'=1}^{N} p_t^{ij}(n')\alpha_t^{ij}(n')\prod_{h\in(g_{ij}\setminus\{j\})}\alpha_t^{ijh}(n')}. \tag{8}$$

Starting from $\mathbf{p}_t^{ij}$, agent i first incorporates her inferred signal from j, $\boldsymbol{\alpha}_t^{ij}$, which measures what j learns from outside agent i's local network since $\mathbf{p}_{t-1}^{j}$. She then incorporates all the new information agent j should have learned from her local network in period t.[21] If she believes there is no new information at all, $\mathbf{p}_{t+1}^{ij} = \mathbf{p}_t^{ij} = \mathbf{p}_t^{j}$, which is simply agent j's report at t. But if there is new information, she expects j to incorporate it and thus will not double count this information from $t + 1$ onward. In a similar way, agent i keeps track of all the information available to each of the fully connected subset of g_i using $\mathbf{p}_t^{ii_1\dots i_l}$.

No learning about the network outside one's local network. To follow our learning procedure, agents neither need to know, nor do they update their beliefs about, the network beyond their local networks. That is, they behave *as if* all their inferred signals from outside their local networks are independent. To see this, observe from expression (5) that agent i treats all her inferred signals from her neighbors and her signal from nature in the same way: They enter her updating rules multiplicatively. In particular, she ignores possible correlations in these inferred signals. If, for instance, her two neighbors have inferred the same signal from an agent unobserved by agent i in period $t - 1$, agent i would have double counted this signal according to our procedure.

Although we show in Section 4.4 that our learning procedure is consistent with fully Bayesian learning when each agent believes that her local network is the entire network, and such beliefs are common knowledge, this is not at all our motivation. Rather, our primary motivation is to develop, based on the existing evidence, a tractable learning procedure in which agents learn rationally within her local network, but ignore possible correlations in her information from the wider network. Agents can simply follow our procedure when they are unable to handle the computational and cognitive burden of Bayesian learning. We also generate clear predictions of when agents can still learn correctly, and identify systematic errors they make when they fail to learn. These predictions can potentially be taken to the data and tested. In Section 5.3, we allow agents to account for some type of information

[21] We use $p_t^{j}(n)$ as agent i's prior in expression (6) because j's report is observable and it makes the expression simpler. But it is clear from the above discussion that we can use $p_t^{ij}(n)$ as her prior as well, which is more standard. We can similarly rewrite any of the higher-order estimates $p_{t+1}^{ii_1\dots i_l}(n)$ in expression (7), using $p_t^{ii_1\dots i_l}(n)$ as the prior instead of $p_t^{i_l}(n)$.

correlation by a simple rule of the thumb to improve their learning outcomes.

Knowledge of the local network only. For agents to follow our procedure, we need every agent to know her local network, and every agent knows every agent in her local network knows his local network, and so on up to the order of the number of agents in each agent's local network. An agent does not need to know a neighbor's information structure because agent i, for example, only infers a posterior distribution $\left\{\Pr\left(y_{t-1}^{ij}|s_1\right), \ldots, \Pr\left(y_{t-1}^{ij}|s_N\right)\right\}$ from each neighbor j every period.

This local network knowledge assumption explains why agents need to form higher-order estimates. Observe from Step 1 of the procedure that $\boldsymbol{\alpha}_t^{ik}$, which is based on i and k's shared information $\{\mathbf{p}_\tau^h : \tau \leq t, h \in g_{ik}\}$, is generally different from $\boldsymbol{\alpha}_t^{jk}$, the new information agent j learns from k, which is based on j and k's shared information $\{\mathbf{p}_\tau^h : \tau \leq t, h \in g_{jk}\}$. Therefore, agent i can only estimate what she believes j learns from k, $\boldsymbol{\alpha}_t^{ijk}$. It also explains why agent i's higher-order estimates and inferred information are restricted to fully connected subsets of g_i. If an agent is not connected to some of the agents in $\{i, i_1, \ldots, i_l\}$, then they don't know his existence, and thus cannot form any estimates involving him. In Section 5.1, we expand the agents' local networks to include their indirect neighbors and the links among them. Surprisingly, the agents' learning outcomes may become worse.

Distribution of the state and signals. We illustrate our learning procedure with a model of finitely many states and finitely many signals, but it is applicable to other information structures. First, the influential information partition model is a special case of our model. Recall that S is the state space and agent i's information structure can be represented by a mapping $\mathcal{P}^i : S \to 2^S \setminus \emptyset$. $\mathcal{P}^i$ associates each state s_n with a non-empty element $\mathcal{P}^i(s_n)$ such that at s_n, agent i considers $\mathcal{P}^i(s_n)$ to be the set of possible states. Moreover, $\mathcal{P}^i$ induces an information partition over the state space if (1) for any $s_n \in S$, $s_n \in \mathcal{P}^i(s_n)$; and (2) for any $s_n, s_{n'} \in S$, $s_n \in \mathcal{P}^i(s_{n'})$ implies $\mathcal{P}^i(s_n) = \mathcal{P}^i(s_{n'})$. In our context, each signal x_m^i informs agent i of an element $\mathcal{P}^i(s_n)$ of her partition. For all $s_{n'} \in \mathcal{P}^i(s_n)$, $\phi_{mn'}^i = 1$; and $\phi_{mn'}^i = 0$ otherwise. The number of possible signals agent i has, M_i, corresponds to the number of elements in her partition. Second, our learning procedure can be easily adapted to the widely-used model with uniformly distributed states and normally distributed signals. That model is simpler because each agent's estimates contain only the expected value of the true state and the precision of the associated distribution.

Communication protocols. Our agents report their most up-to-date estimates of the state distribution similar to Lee (1993) and Eyster and Rabin (2014). Although these reports contain more information than those in the observational learning literature, which typically involve agents' actions or payoffs (see for example Bala and Goyal (1998) and Mossel, Sly and Tamuz (2015)), they are *not* without loss of generality. Each agent can in principle report the history through which she receives her information. That is, agent i reports, in addition to her estimates, "I have heard this report from agent j who has heard it from agent k," and so on. We use the posterior distributions because the messages above are too complex for agents to remember and to use in reality.

3 Main properties

We now characterize several useful properties of our learning procedure. To begin with, separate g_i into (overlapping) subsets within which all agents are fully connected, our learning procedure restricts agent i's higher-order estimates to those within each such subset. For example, in the network depicted in Figure 1 below, referred to as *diamond with a link*, agent i has three neighbors, and the fully connected subsets in g_i are $\{i, j, k\}$, $\{i, j, k'\}$ and $\{i, j\}, \{i, k'\}, \{i, k\}$. Since $\widehat{L}_i = 3$, agent i forms up to third-order estimates.

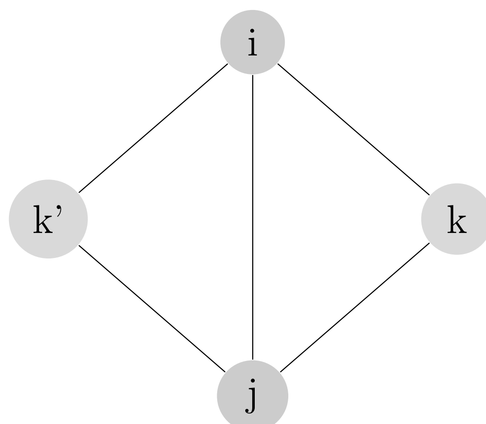

Figure 1: Diamond with a link

We first show that the order of agents in higher-order estimates does not matter.

LEMMA **1.** *Consider any fully connected subset* $\{i_1, \ldots, i_l\}$ *of* g_i*. Let* $\{\beta(1), \ldots, \beta(l)\}$ *be a permutation of* $\{1, \ldots, l\}$*. Then* $\mathbf{p}_t^{i_{\beta(1)} \cdots i_{\beta(l)}}$ *is the same for all* $t \geq 1$*.*

Intuitively, agent i_1 can only form her higher-order estimates using reports from the relevant shared local network. Since the set of distinct agents is the same, $g_{i_1 \ldots i_l} = g_{i_1} \cap \ldots \cap g_{i_l} = g_{i_{\beta(1)}} \cap \ldots \cap g_{i_{\beta(l)}} = g_{i_{\beta(1)} \ldots i_{\beta(l)}}$. That is, agent i_1 and $i_{\beta(1)}$'s estimates are based on the same information $\{\mathbf{p}_\tau^h : \tau < t, h \in g_{i_1 \ldots i_l}\}$, and thus must be identical. The fact that order does not matter allows us to compare two neighbors' estimates about each other easily. By Lemma 1, even if agent i and j have different estimates in a period, $\mathbf{p}_t^i \neq \mathbf{p}_t^j$, agent i's estimates of j's estimates always agree with agent j's estimates of i's: $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ji}$. This pairwise agreement shows that our learning procedure is internally consistent: Each agent always knows the signal a neighbor infers from herself. That is, what i thinks j infers from her is exactly what j infers, $\boldsymbol{\alpha}_t^{iji} = \boldsymbol{\alpha}_t^{ji}$ for all t, because $\mathbf{p}_t^{iji} = \mathbf{p}_t^{ji}$.[22] Similarly, agent i and j also agree with what each other infers from a common neighbor.

Our learning procedure may still seem to impose a heavy computational burden on the agents because it requires them to form higher-order estimates. But agent i's higher-order estimates can be further simplified in certain type of local networks. We say (g_i, G_i) satisfies *local connection symmetry* (LCS from now on) if g_{ij} is fully connected for every $j \in \mathrm{N}_i$. Intuitively, this is the case if either $\mathrm{N}_i \cap \mathrm{N}_j = \emptyset$, or for every $k \in \mathrm{N}_i \cap \mathrm{N}_j$, there does not exist another agent k' such that $k' \in \mathrm{N}_i \cap \mathrm{N}_j$, but $kk' \notin G$. For instance, in Figure 1, agent k's local network clearly satisfies this property since $g_{ik} = g_{jk} = \{i, j, k\}$, and similarly for agent k'. But the local networks for agent i and j fail this property because $kk' \notin G$. When i's local network satisfies LCS, agent i only needs to form second-order estimates.

Lemma **2.** *For every agent $j \in \mathrm{N}_i$ and every period t,*

(1) Individual agreement: if (g_i, G_i) satisfies LCS, then $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik}$ and $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ijk} = \mathbf{p}_t^{ij\ldots k}$ for distinct agents $ij \ldots k \in g_{ij}$.

(2) Local-network agreement: if (g_l, G_l) satisfies LCS for every agent $l \in g_i$, then $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$ for any $k \in \mathrm{N}_i \cap \mathrm{N}_j$.

Intuitively, for agent i, if agents in g_{ij} are fully connected for every j, then $g_{ij} = g_{ik}$ if her neighbors j and k are connected. From agent i's perspective, information agents in g_{ij} share is the same as information agents in g_{ik} share, implying $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik}$. Moreover, in forming higher-order estimates, agent i thinks everyone in g_{ij} also thinks each other has access to the same set of information, and thus her second-order estimates suffice. The second part of Lemma 2 shows that if the local network of every agent in g_i has this property, the learning procedure can be simplified as agent i directly use i's inferred signal from k to replace j's inferred signal from k since $\mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$. In step 3 of our learning procedure, there is no need to evaluate equation (7), and equation (6) becomes:

$$\alpha_t^{ji}(n) = \frac{p_t^i(n)}{p_t^{ij}(n)} \Bigg/ \sum_{n'} \frac{p_t^i(n')}{p_t^{ij}(n')}, \text{ and } \; p_{t+1}^{ij}(n) = \frac{p_t^j(n)\alpha_t^{ji}(n)\prod_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\alpha_t^{ik}(n)}{\sum_{n'=1}^{N} p_t^j(n')\alpha_t^{ji}(n')\prod_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\alpha_t^{ik}(n')}. \tag{9}$$

Similarly, we say that a network (g, G) satisfies *global connection symmetry* (GCS from now

[22] One implication is that when agents only know their neighbors, but not the links among them, they can still learn correctly in some networks such as a tree. See Section 5.1 for more details.

on) if property LCS holds for every agent $i \in g$. In these networks, every agent forms only her own estimates and her second-order estimates for each neighbor.

So far we focus on learning from agent i's perspective. Since each agent learns from her neighbors who in turn learns from their neighbors, the agents' inferred signals can be rewritten using expression (5) and (6) as follows.

LEMMA **3.** *For every agent $i \in g$, $t \geq 1$,*

$$\alpha_{t+1}^{ij}(n) = \prod_{l \in (g_j \setminus g_i) \cup \{j\}} \alpha_t^{jl}(n) \prod_{h \in g_{ij} \setminus \{j\}} \frac{\alpha_t^{jh}(n)}{\alpha_t^{ijh}(n)} \Bigg/ \left(\sum_{n'} \prod_{l \in (g_j \setminus g_i) \cup \{j\}} \alpha_t^{jl}(n') \prod_{h \in g_{ij} \setminus \{j\}} \frac{\alpha_t^{jh}(n')}{\alpha_t^{ijh}(n')} \right) \quad (10)$$

Focusing only on the numerator in expression (10), we can see that what agent i learns from j consists of two parts. The first part is, as expected, what agent j learned in the previous period from nature or his neighbors unobserved by agent i. That is, all the information from outside agent i's local network. The second part is more subtle: It reflects the difference between what j actually learned from his neighbor h and what i thinks that j learned from h. We begin with the case of networks that satisfy GCS. In this case, only the first part matters because $\boldsymbol{\alpha}_t^{jh} = \boldsymbol{\alpha}_t^{ijh}$ for all $h \in g_{ij} \setminus \{j\}$, and thus the second part becomes 1. Note that Lemma 3 does not mean that agent i is able to learn all the other agents' signals correctly. If a signal reaches multiple neighbors of agent i, she may infer there are multiple copies of that very signal according to equation (10). For instance, remove the link between agent i and j in Figure 1, so that the network becomes a diamond $\{k'ijk\}$, in which agent k' receives the only initial signal at $t = 0$. Then at $t = 2$, agent i and j each infers the same signal. At $t = 3$, agent k infers a product of these two signals, namely his inferred signal contains two copies of the original signal received by agent k', and is thus incorrect. But if each agent in a network does learn each true signal only once, she will be able to form the correct Bayesian posterior despite not knowing the entire network.

An agent's inferred signals can also be wrong due to the second part in the numerator of expression (10). Even if there is no new information from outside agent i's local network previously (the first part $\alpha_t^{jl}(n) = 1/N$ for all $l \in (g_j \setminus g_i) \cup \{j\}$), she may nonetheless infer "new" signals because i and j have different local networks. We will illustrate in Example 2 at the end of this section that in Figure 1, agent k may infer new signals from i because what k thinks i learns from j, $\boldsymbol{\alpha}_t^{kij}$, is different from what i truly learns from j, $\boldsymbol{\alpha}_t^{ij}$.

Although informative signals may travel through network (g, G) via many different paths, a very convenient property of our learning procedure is that for any given sequence of realized signals, signals travel independently. Therefore we can analyze learning under each signal separately. Divide the full sequence of observed signals in the network X_T into any two disjoint sets of signals X_T^a and X_T^b such that $X_T = X_T^a \cup X_T^b$ and $X_T^a \cap X_T^b = \emptyset$. Recall that $\mathbf{p}_t^i$ is agent i's estimates of the true state under X_T, and let $(\mathbf{p}_t^{a,i}, \mathbf{p}_t^{b,i})$ be her estimates under X_T^a and X_T^b respectively. We say signals can be *decomposed* if the agent's estimates under X_T is equal to the combination of her estimates under X_T^a and X_T^b using Bayes' rule. That is, for all $t \geq 1$, i and s_n,

$$p_t^i(n) = \frac{p_t^{a,i}(n)p_t^{b,i}(n)}{\sum_{n'=1}^N p_t^{a,i}(n')p_t^{b,i}(n')}; \tag{11}$$

$$p_t^{ij}(n) = \frac{p_t^{a,ij}(n)p_t^{b,ij}(n)}{\sum_{n'=1}^N p_t^{a,ij}(n')p_t^{b,ij}(n')}. \tag{12}$$

Moreover, for any fully connected subset $\{ii_1 \ldots i_l\}$ of g_i,

$$p_t^{ii_1\ldots i_l}(n) = \frac{p_t^{a,ii_1\ldots i_l}(n)p_t^{b,ii_1\ldots i_l}(n)}{\sum_{n'=1}^N p_t^{a,ii_1\ldots i_l}(n')p_t^{b,ii_1\ldots i_l}(n')}. \tag{13}$$

Lemma 4. *Signals can be decomposed under the proposed learning procedure. If $X_T = X_T^a \cup X_T^b$, $X_T^a \cap X_T^b = \emptyset$, then equations (11), (12) and (13) hold for all i, s_n and $t \geq 1$.*

Lemma 4 does not mean that agents can learn each signal correctly. Rather, it means that given a sequence of realized signals, one signal travels, with all its possible repetitions and distortions, independently from another signal. Consider two informative signals observed by different agents at different times, x_0^l and x_τ^h, which are inferred by i from her neighbors in N_i at time $t+1$ in the form of y_{t-1}^{ij} and y_{t-1}^{ik} respectively. Then agent i's estimates at $t+1$ are just the combination of y_{t-1}^{ij} and y_{t-1}^{ik} by Bayes' rule. They are the same if we study agent i's estimates when x_0^l or x_τ^h was the only signal, and then combine these estimates at time $t+1$ by Bayes' rule.

Lemma 4 holds because agent i's inferred signal from j is independent from her other inferred signals, and agent j knows it. Therefore even though j does not know all the sources of agent i's information, j can identify the part of i's updated estimates that is due to his information (and the information from their common neighbors). Agent j considers the rest as new information.[23] An important implication of Lemma 4 is that if agents can learn each separate signal correctly, they can do so for many signals. Therefore it is possible to provide general results for *all* sequences of realized signals.

We end this section with an example to illustrate the properties of our learning procedure.

[23]If this were not the case, for example, if agent i incorporates information from g_{ij} differently depending on her other inferred signals, agent j may not be able to form his estimates of i's estimates (and any higher-order estimates) in a consistent way. See Section 5.3 for more details.

EXAMPLE **2.** *In the diamond with a link depicted in Figure 1, agent i and her neighbors have three shared local networks: $g_{ij} = g_i \cap g_j = \{i, j, k', k\}$, $g_{ik'} = \{i, j, k'\}$ and $g_{ik} = \{i, j, k\}$. We continue to use the binary state and signal case described in Example 1. There are two informative signals $x_0^{k'} = 1$ and $x_1^j = 1$.*

At $t = 0$, agent k' learns $x_0^{k'} = 1$. At $t = 1$, the agents' estimates are $p_1^{k'}(1) = \phi^{k'}$ and $p_1^i(1) = p_1^j(1) = p_1^k(1) = \frac{1}{2}$. All their higher-order estimates are the symmetric prior.

At $t = 2$, agent i observes that $p_1^{k'}(1) \neq p_1^{ik'}(1) = \frac{1}{2}$, and so is j. They each infer a signal $\alpha_1^{ik'}(1) = \alpha_1^{jk'}(1) = \phi^{k'}$ from k', and nothing from agent k. Agent i needs to form up to third-order estimates: $p_2^i(1) = p_2^{ij}(1) = p_2^{ijk'}(1) = \phi^{k'}$ and $p_2^{ijk}(1) = \frac{1}{2}$. Agent j updates using both his inferred signal from k' and his signal x_1^j: $p_2^j(1) = \Pr(s = 1|x_0^{k'}, x_1^j) = \frac{\phi^{k'}\phi^j}{\phi^{k'}\phi^j + (1-\phi^{k'})(1-\phi^j)}$. By Lemma 1, the order of the agents in the estimates does not matter, and thus $p_2^{ji}(1) = p_2^{ij}(1) = \phi^{k'}$, $p_2^{ijk'}(1) = p_2^{jik'}(1) = \phi^{k'}$ and $p_2^{ijk}(1) = p_2^{jik}(1) = \frac{1}{2}$.

Agent k' learns nothing from his neighbors: $\alpha_1^{k'i}(1) = \alpha_1^{k'j}(1) = \frac{1}{2}$, and thus $p_2^{k'}(1) = \phi^{k'}$. Moreover, it is easy to see that $p_2^{k'i}(1) = p_2^{k'j}(1) = p_2^{k'ij}(1) = p_2^{k'ji}(1) = \phi^{k'}$. Because the local network of agent k' satisfies LCS, by Lemma 2, the second-order estimates suffice. As to agent k, he still learns nothing and thus all his estimates remain the symmetric prior.

At $t = 3$, agent i and k' observes that $p_2^j(1) \neq \phi^{k'}$ and infers that $\alpha_2^{ij}(1) = \alpha_2^{k'j}(1) = \phi^j$, while agent j learns nothing from k', i and k. It is easy to see that in the shared local network of i, j, k', their estimates agree: $p_3^{k'}(1) = p_3^i(1) = p_3^j(1) = p_3^{k'i}(1) = p_3^{ij}(1) = p_3^{jk'}(1) = p_2^j(1)$. From now on, we focus on the shared local network of i, j, k.

Agent k infers $\alpha_2^{ki}(1) = \phi^{k'}$ and $\alpha_2^{kj}(1) = p_2^j(1)$. Agent k's local network satisfies LCS, and by first part of expression (10) in Lemma 3, $\alpha_2^{ki}(1) = \alpha_1^{ik'}(1)$ and $\alpha_2^{kj}(1) = \frac{\alpha_1^{jk'}(1)\alpha_1^{jj}(1)}{\sum_{s\in\{0,1\}} \alpha_1^{jk'}(s)\alpha_1^{jj}(s)}$. Also, $\alpha_2^{kij}(1) = p_2^j(1)$ and $\alpha_2^{kji}(1) = \phi^{k'}$ because he thinks that i and j should infer new signals from each other. By Bayes' rule, $p_3^k(1) = \frac{(\phi^{k'})^2\phi^j}{(\phi^{k'})^2\phi^j + (1-\phi^{k'})^2(1-\phi^j)}$. Agent i and j know what k learns from them: $\alpha_2^{iki}(1) = \alpha_2^{jki}(1) = \phi^{k'}$ and $\alpha_2^{ikj}(1) = \alpha_2^{jkj}(1) = p_2^j(1)$.

At $t = 4$, agent i, j and k' do not learn anything new, and thus their estimates remain unchanged: $p_4^i(1) = p_4^j(1) = p_4^{k'}(1) = p_2^j(1)$. But agent k again infers new signals: $\alpha_3^{kj}(1) = \alpha_3^{ki}(1) = 1 - \phi^{k'}$. That is, he infers two signals offsetting $x_0^{k'}$ because he does not know i and j have a common neighbor k'. This can be seen from Lemma 3. The first part of expression (10) is 1, and the second part is $\alpha_4^{kj}(1) = \frac{\alpha_3^{ji}(1)}{\alpha_3^{kji}(1)} \Big/ \sum_{n'} \frac{\alpha_3^{ji}(n')}{\alpha_3^{kji}(n')} = 1 - \phi^{k'}$, and similarly $\alpha_4^{ki}(1) = 1 - \phi^{k'}$. The difference in his third-order inferred information $\boldsymbol{\alpha}_3^{kij}$ and the actual inferred signal $\boldsymbol{\alpha}_3^{ij}$ causes his inferred signal to be wrong. Thus $p_4^k(1) = \phi^j$.

From now on, agent k', i, j's estimates remain at $p_2^j(1)$. But for all the odd periods $t \geq 3$, agent k's estimates are the same as $p_3^k(1)$; and for all the even periods $t \geq 4$, his estimates are $p_4^k(1)$. His estimates oscillate and never converge.

Finally, under Lemma 4, the agents' estimates are the same if they have received the signals separately, and we combine their estimates via Bayes' rule. It is easy to see that since everyone observes j, if x_1^j is the only signal, then at $t = 3$, everyone's estimates are ϕ^j and do not change. Next, if $x_0^{k'}$ is the only informative signal, from $t = 2$ onward, k', i, j's estimates remain constant at $\phi^{k'}$. But from $t \geq 3$, agent k's estimates keep oscillating between $\frac{(\phi^{k'})^2}{(\phi^{k'})^2+(1-\phi^{k'})^2}$ in odd periods and $\frac{1}{2}$ in even periods. The combination by Bayes' rule is exactly the same as above. $\diamond$

4 Bayesian learning outcomes

We now provide sufficient conditions for the agents to learn correctly in two benchmark cases. We also characterize two network features that lead to systematic learning errors.

Recall that the agents' learning outcomes are *Bayesian* if there exists some period t in which all agents' first-order estimates agree with the Bayesian posterior given X_T, and remain constant afterwards.[24] We sometimes also use a stronger notion of correct learning, in which the agents' estimates agree with the Bayesian posterior in every period *given the travel paths of signals*.[25] Let $d(il)$ be the distance, or the length of the shortest path between i and l; and let $d(ii) = 0$ for convenience. The diameter of network (g, G) is then

$$D = \max_{i,l \in g} d(il).$$

A signal (or information contained in this signal) then takes at most D periods to reach every agent $\{1, \ldots, L\}$ in (g, G). Recall that $X_t^i = \{x_\tau^i : \forall \tau \leq t\}$ if $t \geq 0$ and $X_t^i = \emptyset$ otherwise. Because it takes $d(il)$ periods for a signal to travel from l to i, agent i's posterior based on the signals agent l received up to period $t - d(il)$ for all $l \in g$ is:

$$q_{t+1}^i(n) = \Pr\left(s_n \middle| X_{t-d(i1)}^1, \ldots, X_{t-d(iL)}^L\right). \tag{14}$$

We call the agents' learning outcomes *strongly Bayesian* if $\mathbf{p}_t^i = \mathbf{q}_t^i$ for all i and t.

[24]Note that this implies that agents infer no new information since t, and thus each agent's higher-order estimates also agree with her first-order estimates $\mathbf{p}_t^i$.

[25]Agents may still make mistakes even when their learning outcomes are Bayesian. For instance, in Example 2. Suppose the informative signals are $x_0^{k'} = 1$ and $x_2^{k'} = 0$. It is easy to show that every agent agrees at $t = 4$: $p_4^h(1) = \frac{1}{2}$, $h \in \{k', i, j, k\}$, which are the correct posterior given that agent k''s signals cancel out each other.

4.1 Bayesian learning under information partition model

For the rest of this section, let s_1 be the true state. Recall from Section 2.2 that the information partition model is a special case of our setup, in which each agent's initial signal x_0^i informs her of the element $\mathcal{P}^i(s_1)$, which contains the set of states she cannot distinguish from the true state. For simplicity only, assume the agents receive no further signals.

Suppose every agent has information partitions $(\mathcal{P}^i_{i\in g}, S)$ as defined in Section 2.2. We now turn to the question of whether they can all agree by following our learning procedure and if so, what they can agree on. This question has been studied in the literature on knowledge and consensus (see Aumann (1976), Geanakoplos and Polemarchakis (1982), Parikh and Krasucki (1990), Mueller-Frank (2013), among many others). It analyzes under which conditions and what reporting protocols, repeated communication among a finite set of individuals leads to consensus. Our procedure generalizes the reporting protocol to more than pairwise communication, and the message space to the posterior distribution of the states. Consider the following example from Geanakoplos and Polemarchakis (1982).

Example **3.** *There are two agents* 1 *and* 2*. The state space is* $S = \{s_1, s_2, ..., s_9\}$*. The states are equally likely. Agent* 1*'s partition is* $\mathcal{P}^1 = \{(s_1, s_2, s_3), (s_4, s_5, s_6), (s_7, s_8, s_9)\}$*; and agent* 2*'s partition is* $\mathcal{P}^2 = \{(s_1, s_2, s_3, s_4), (s_5, s_6, s_7, s_8), s_9\}$*. Thus their initial signals are* $\mathcal{P}^1(s_1) = \{s_1, s_2, s_3\}$ *and* $\mathcal{P}^2(s_1) = \{s_1, s_2, s_3, s_4\}$*. Let event* $A = \{s_3, s_4\}$*.*

Geanakoplos and Polemarchakis (1982) allow the agents to know each other's information partitions and to announce and revise their posteriors of how likely event A is true.[26] At $t = 1$, agent 1 knows A can only be true at $s = s_3$, while agent 2 thinks both $s = s_3$ and $s = s_4$ can be true. So they announce 1/3 and 1/2 respectively. But this is also consistent with the true state being s_4 and $\mathcal{P}^1(s_4) = \{s_4, s_5, s_6\}$ and $\mathcal{P}^2(s_4) = \{s_1, s_2, s_3, s_4\}$. At $t = 2$, agent 1 announces 1/3 again because she has learned nothing. Next, noticing that agent 1 does not change her posterior to 1, agent 2 realizes that the true state cannot be s_4, and thus changes his posterior to 1/3.[27] From $t = 3$ onwards, they agree and the learning is over.

In our model, we do not need the partitions to be common knowledge. Instead, agents directly learn their neighbors' element containing the true state from their reports. At $t = 1$ agent 1 announces her estimates of the state distribution as $\mathbf{p}_1^1 = \{1/3, 1/3, 1/3, 0, 0, 0, 0, 0, 0\}$. Agent 2 announces $\mathbf{p}_1^2 = \{1/4, 1/4, 1/4, 1/4, 0, 0, 0, 0, 0\}$. By our Step 1, agent 2 infers the signal $\boldsymbol{\alpha}_1^{21} = \{1/3, 1/3, 1/3, 0, 0, 0, 0, 0, 0\}$. Using updating rule (5), we can see that

[26] The posterior of an event A given agent i's information $\mathcal{P}^i(s_1)$ is simply $\frac{\Pr(A\cap\mathcal{P}^i(s_1))}{\Pr(\mathcal{P}^i(s_1))}$.

[27] To see this, note that agent 2's report makes $\mathcal{P}^2(s_1)$ common knowledge. Knowing agent 2's element of partition, agent 1 should have changed her posterior to 1 if s_4 is the true state.

$\mathbf{p}_2^2 = \mathbf{p}_2^1 = \{1/3, 1/3, 1/3, 0, 0, 0, 0, 0, 0\}$. Learning stops because no one has any new information. The correct learning takes one, instead of two periods of communication. $\diamond$

Under our learning procedure, not only agents in any network with information partition will agree, they agree as soon as the signals finish traveling.

PROPOSITION **1.** *If all agents in* (g, G) *have information partitions* $\left(\mathcal{P}^i_{i\in g}, S\right)$*, then their learning outcomes are strongly Bayesian and they reach consensus at* $t = D + 1$*.*

Proposition 1 holds for all networks because of a special feature of the information partition model: Agents do not make mistakes even if they treat correlated information as independent. To see this, note that upon the first time each signal reaches agent i, she eliminates the states from her current set of possible states according to the signal. Even if agent i has inferred the same signal from multiple neighbors unknowingly, her estimates are unaffected in that she removes the same set of states as she would given one such signal, and then assigns equal probabilities to the remaining states. At $t = D + 1$, the agents have received all signals. Therefore their estimates are simply $p_t^i(n) = 1/|\mathcal{P}^g(s_1)|$ if $s_n \in \mathcal{P}^g(s_1)$, where $\mathcal{P}^g(s_1) \equiv \cap\{\mathcal{P}^i(s_1)\}_{i\in g}$ is the intersection of all agents' elements of partition containing state s_1, and 0 otherwise. More generally, the same logic shows that if any agent believes that a state is the true state with probability 0 given a signal, all agents in the network can learn this within $D + 1$ periods after the arrival of this signal. In comparison, Geanakoplos and Polemarchakis (1982) show that using their communication protocol, agents may never form the posterior they would have if they pool their information. That is, their learning outcomes may not be Bayesian eventually.

Proposition 1 no longer holds when the information partition model is perturbed.[28] There is a discontinuity in that if agents have any doubt about the mapping from the signals to their elements of partition, their estimates may depend on the network structure. Let us revisit the diamond-with-a-link network in Figure 1.

EXAMPLE **4.** *The state space* $S = \{s_1, s_2, s_3, s_4\}$*. Partition of agent* k' *is* $\mathcal{P}^{k'} = \{(s_1, s_2), (s_3, s_4)\}$*. All the other agents' partitions are simply* S*. Only agent* k' *sees an informative signal, which indicates the correct element with probability* $1 - \varepsilon$*, and the wrong element with probability* ε*. That is, when* $x_0^{k'} = (s_1, s_2)$*,* $\mathbf{p}_1^{k'} = (\frac{1-\varepsilon}{2}, \frac{1-\varepsilon}{2}, \frac{\varepsilon}{2}, \frac{\varepsilon}{2})$*.*

At $t = 1$, the estimates are $\mathbf{p}_1^i = \mathbf{p}_1^j = \mathbf{p}_1^k = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$ and $\mathbf{p}_1^{k'} = \left(\frac{1-\varepsilon}{2}, \frac{1-\varepsilon}{2}, \frac{\varepsilon}{2}, \frac{\varepsilon}{2}\right)$. At $t = 2$, i and j infer the signal from agent k', and thus $\mathbf{p}_2^i = \mathbf{p}_2^j = \mathbf{p}_2^{k'} = \left(\frac{1-\varepsilon}{2}, \frac{1-\varepsilon}{2}, \frac{\varepsilon}{2}, \frac{\varepsilon}{2}\right)$ and

[28] One way to perturb the information partition model is to introduce the possibility of errors. For instance, each signal informs an agent of her true element of partition with probability close to 1, but inform her of some other element(s) with the complementary probability.

$\mathbf{p}_2^k = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$. For all $t \geq 3$, $\mathbf{p}_3^i$, $\mathbf{p}_3^j$ and $\mathbf{p}_3^{k'}$ remain unchanged. But agent k learns two copies of $x_0^{k'}$, so $\mathbf{p}_3^k = \left(\frac{(1-\varepsilon)^2}{2((1-\varepsilon)^2+\varepsilon^2)}, \frac{(1-\varepsilon)^2}{2((1-\varepsilon)^2+\varepsilon^2)}, \frac{\varepsilon^2}{2((1-\varepsilon)^2+\varepsilon^2)}, \frac{\varepsilon^2}{2((1-\varepsilon)^2+\varepsilon^2)}\right)$. At $t = 4$, k infers two opposite signals, so $\mathbf{p}_4^k = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$. Hereafter agent k oscillates between $\mathbf{p}_3^k$ and $\mathbf{p}_4^k$. At $\varepsilon = 0$, all agents estimates are Bayesian. But for any $\varepsilon > 0$, the agents never agree. $\diamond$

4.2 Bayesian learning in social quilts

Given the result in Section 4.1, from now on, we restrict attention to the case of non-partitional signals, where $\phi_{mn}^i \in (0, 1)$ for all the agents. That is, every signal can be observed with a positive probability in every state. Since signals can be decomposed by Lemma 4, we often start with the case of one signal only.

It is easy to see that in fully connected networks (cliques) such as the triangle in Example 1, agents are able to learn any signal correctly. Whenever a signal reaches agent i in a clique, all her neighbors learn this signal from her report in one period. Moreover, every agent thinks that everyone else has learned the signal from agent i, and thus they will not double count this information. In the next period, everyone believes that there is one and only one copy of this signal and forms the correct estimates. We now generalize this intuition.

To begin with, a *tree* is a graph in which any two agents $i, h \in g$ are connected by a unique path. A *circle* is a path from an agent i_0 to herself through distinct agents. That is, $c = \{i_0, ..., i_k\}$, $i_l i_{l+1} \in G$ for any $l < k$, $i_0 i_k \in G$, and $i_l \neq i_h$ for any $l, h \leq k$. By definition, a tree contains no circles. Because there are at least two paths for any pair of agents in a circle to reach each other: one clockwise and one counterclockwise.

Definition 1. *Network (g, G) is a social quilt if $ij \in G$ for any agent i and j in the same circle.*

Intuitively, a social quilt is a tree-like union of cliques, such that any circle within a social quilt must be embedded in a clique. Let c be a *simple circle* if it contains more than three agents, and $\mathrm{N}_{i_l} \cap c = \{i_{l-1}, i_{l+1}\}$ for any $1 \leq l \leq k-1$, $\mathrm{N}_{i_0} \cap c = \{i_1, i_k\}$, and $\mathrm{N}_{i_k} \cap c = \{i_0, i_{k-1}\}$. That is, other than the two adjacent neighbors, agent i_l has no connections to any other agent in a simple circle. Then a social quilt can be characterized by the following property.

Lemma 5. *Network (g, G) is a social quilt if and only if it satisfies GCS and does not contain a simple circle.*

Agents reach consensus if their first-order estimates agree in some period and remain constant afterwards. Recall that no informative signal arrives at or after T. Then we have:

Proposition 2. *If (g, G) is a social quilt, then the agents' learning outcomes are strongly Bayesian and they reach consensus at period $T + D$.*

Intuitively, two important features of social quilts enable agents to learn correctly: local cliques and a global tree. In each of the local cliques, every agent is able to infer a new signal correctly, and to infer that everyone else in the clique also infers the same signal as she does. In this way, they agree that there is one and only one copy of each signal. Cliques are then connected in a tree, and thus there is a unique path to go from one clique to another. This implies that every signal arrives at each clique only once, and when each signal reaches the "terminal" cliques of the tree, the signal stops traveling.

More precisely, suppose there is only one signal, x_0^i. We show that the learning is strongly Bayesian by proving that each agent h learns this signal at, and only at, period $d(ih) + 1$. First, there exists a unique shortest path between any two agents i and h in a social quilt, which is the path the signal travels from i to h. Intuitively, the path can be constructed by taking a direct cross of any clique and then taking the unique shortest path along the tree. Second, once agent h learns the signal, he will not see this signal again. More specifically, suppose the shortest path between i and h is $\{i, \ldots, k, h\}$, and agent l is a neighbor of agent h who learns the signal from h one period later. Then, l must be further away from i, $d(il) = d(ih) + 1$. In other words, the unique shortest path from agent i to agent l must go through h.[29] As the signal is learned by agents further away from i, it cannot reach agent h again. Similarly, with multiple signals, agent i's estimates at $t + 1$ include signals observed by each agent h from period 0 to period $t - d(ih)$, and thus the learning outcome is strongly Bayesian. Moreover, because the last signal arrives at period $T - 1$ and it takes at most $D + 1$ periods to reach all agents, all learning stops by the end of period $T + D$.

Next, if T becomes sufficiently large, even though the agents may receive uninformative signals in each period, the network receives a large number of signals, a positive fraction of which is informative.[30] One immediate implication of Proposition 2 is that if for any $s_{n'} \neq s_n$, there exist agent i and signal x_m^i such that $\phi_{mn}^i \neq \phi_{mn'}^i$, then as $T \to \infty$, $p_T^l(s_1) \to 1$ for all $l \in g$. Intuitively, since the agents' estimates agree with the Bayesian posterior in a social quilt, they learn the true state eventually.

[29] If to the contrary, there are two paths from i to l, then there must be a circle involving (but not limited to) agent k, h and l while $kl \notin G$, which is impossible in a social quilt.

[30] More precisely, we assume that at least one agent receives an informative signal with a probability strictly between 0 and 1.

4.3 When Bayesian learning is impossible

If a network is not a social quilt, then by Lemma 5, it must feature simple circles or networks that do not satisfy GCS (or both). Both features impede correct learning. First, in simple circles, because each agent only knows her local network, she may keep inferring "new" signals from her neighbors when it is the same signal reaches her repeatedly. Second, when GCS does not hold, agents have asymmetric knowledge of their local networks, and thus even neighbors may never form consensus as in Example 2.

We first isolate the problems caused by simple circles by assuming that the network satisfies GCS. The following result provides a lower bound on how often a signal is repeatedly learned in such a network as a function of time. In particular, it shows that when there is only one simple circle, the repetition increases linearly in time; but if there are multiple simple circles, the repetition grows exponentially.

PROPOSITION **3.** *Suppose a network satisfies GCS and contains $k_{sc} \geq 1$ simple circles. Let k be the number of agents in the largest simple circle and x_0^i be an informative signal. Then at any $t \in \left[\tau(D + \lceil k/2 \rceil) + 1, (\tau+1)(D + \lceil k/2 \rceil)\right]$,*[31] *any agent l in a simple circle believes there are at least two copies of x_0^i if $\tau = 1$; and at least*

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

copies of signal x_0^i if τ is an integer larger than 1.

Begin with only one simple circle. Once the signal reaches any agent in the simple circle, it travels around in both directions. After k periods, every agent infers two new copies of this signal, one from each neighbor. This continues every k periods, and thus the repetition of the signal grows linearly. With multiple simple circles, each agent also learns the new information from all the other simple circles, and thus the repetition grows faster. Moreover, since all the agents are path-connected, all "new" information will reach all the agents, even those outside any simple circle within D periods. In particular, the number of copies of x_0^i that any agent has learned at period $t + D$ must be (weakly) higher than the maximal number of copies any agent has learned at period t. Therefore we can easily give a lower bound for all agents using Proposition 3.[32]

[31] For any $x \in \mathbf{R}$, $\lceil x \rceil$ is the smallest integer that is greater or equal to x.

[32] This result provides only a lower bound because it takes at most D period for any signal to reach every agent in the network, and at most another $\lceil k/2 \rceil$ periods for this signal to reach every agent in a simple circle again and becomes double counted. Therefore even though we only double the number of the signals agents in a simple circle learn every $D + \lceil k/2 \rceil$ periods, signals may travel through some simple circles multiple

Using Proposition 3, we can see right away that if a network contains any simple circle, repeated learning of each signal is inevitable. This immediately implies that if there is only one informative signal, then all the agents form the wrong estimates in the long run. Let $S(X_T)$ be the set of states that are most likely to be the true state given the sequence of signals X_T a network received. More precisely, $\tilde{s}(X_T) \in S(X_T)$ if $\tilde{s}(X_T) \in \arg\max_{s \in S} \Pr(s|X_T)$, and $|S(X_T)|$ is the number of states in the set. If $X_T = \{x_0^i\}$, the agents believe they have received an infinite number of copies of this signal. That is: $p^h_{t\to\infty}(\tilde{s}(X_T)) = \frac{1}{|S(X_T)|}$; and $p^h_{t\to\infty}(s_{n'}) = 0$ if $s_{n'} \notin S(X_T)$. Another implication is as follows.

Corollary **1.** *Suppose a network satisfies GCS and contains exactly one simple circle, then if $\tilde{s}(X_T)$ is unique, all agents' estimates converge: $p^h_{t\to\infty}(\tilde{s}(X_T)) = 1$.*

Corollary 1 shows the order of signal's arrival does not matter because with only one simple circle, all signals are repeated at the same linear rate by Proposition 3. After period $T + D$, all the informative signals have reached all agents. No matter how many copies of each signal the agents have at this point, the number would increase by two after every k periods. Because the signals grow at the same rate, the first-order effect is the ratio of how likely an agent receives these signals given two different states: $\Pr(X_T|s_{n'})/\Pr(X_T|s_n)$. When $\tilde{s}(X_T)$ is unique, $\Pr(X_T|s_{n'})/\Pr(X_T|\tilde{s}(X_T)) < 1$ for all $s_{n'} \neq \tilde{s}(X_T)$. As t increases, these terms all go to zero, and thus all agents think that $\tilde{s}(X_T)$ is the true state.

By Lemma 4 and Proposition 3, an agent's estimates at any given time are determined by the number of copies of each signal she has inferred so far. Because the signals are repeated at different rates in any given period depending on their arrival times and locations, in general, the agents' estimates may not converge if the network has multiple simple circles. Furthermore, if the network has multiple simple circles, agents may unknowingly put higher and higher weights on earlier signals at the expense of later ones. The next example shows that the earlier signals may grow so fast that agents cannot be persuaded by an arbitrarily large number of later signals. Therefore they may be led astray by a false signal in the beginning and fail to learn the true state, as in the inner city neighborhood example in the introduction.

Example **5.** *Failure of the Law of Large Numbers. Consider eight agents connected in a cube, as in Figure 2. Continue with the binary state and signal setting, with two symmetric signals satisfying $\Pr(x_t^i = 1|s = 1) = \Pr(x_t^i = 0|s = 0) = \phi > \frac{1}{2}$ for all i. The true state is $s = 1$. Suppose that each agent observes a signal of $x_0^i = 0$ at $t = 0$; and a signal of*

times and generate more copies of this signal in the meantime. The bound in proposition 3 is tight if and only if the network itself is a simple circle.

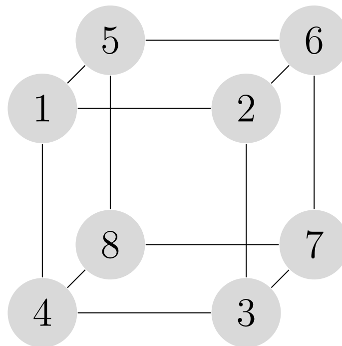

Figure 2: A cube with 8 agents

$x_t^i = 1$ *for* $t \in \{1, ..., T-1\}$*. Then the agents believe the state is* 0 *as* T *approaches infinity:* $\lim_{T\to\infty} p_T^i(0) = 1$*.*

Why do agents believe the state is 0 even when they receive so many opposing (and correct) signals from $t = 1$ onward? Observe that at $t = 1$, each agent reports $p_1^i(0) = \phi$. At $t = 2$, each agent infers three signals of 0 from their neighbors (plus her own signal of 1 received by the end of $t = 1$), so her own estimates are $p_2^i(0) = \frac{\phi^3}{\phi^3+(1-\phi)^3}$. Her estimates of her neighbors' estimates are $p_2^{ij}(0) = \frac{\phi^2}{\phi^2+(1-\phi)^2}$, because she thinks that each of her neighbors learns a signal of 0 from herself plus his own signal of 0. Therefore at $t = 3$, each agent again infers three new signals of 0, plus one signal of 1 from the nature. The agents' learning in all later periods up to T is identical to that in period 2. Therefore, all agents think they are learning more and more signals of 0 in net and believe the state is 0 in the limit.[33] $\diamond$

We now consider the problems caused by networks that fail GCS, by assuming the network contains no simple circles. Let $\tilde{g}$ be the set of agents whose neighbors fail property LCS due to their presence. More specifically, agent $l \in \tilde{g}$ if there exist some agents i, j and k such that $l \in g_{ij}, k \in g_{ij}$, but $lk \notin G$. We show that some agent must fail to learn correctly.

Proposition 4. *Suppose* (g, G) *contains no simple circles but property GCS does not hold. If any agent* $l \in \tilde{g}$ *receives* x_0^l*, the only informative signal, then* $p_t^l = \Pr(s_n|x_0^l)$ *for all* $t \geq 1$*, but there exists an agent whose learning outcomes are never Bayesian.*

Classify the agents by their distance to agent l, $\mathrm{N}_l^d = \{h \in g : d(hl) = d\}$. Because simple circles are the channel for signals to travel back to its source, we can show that without simple circles, no agent in N_l^d infers any new signals from her successors in N_l^{d+1}. Suppose to the contrary, agent i in N_l^d infers a new signal from j in N_l^{d+1}, which is the first

[33]It is unlikely that all signals from period 1 to period $T-1$ are 1 for all agents $i = 1, \ldots, 8$ even if the true state is 1. But since the agents put a probability arbitrarily close to 1 on state 0 in this case, they would put a even higher probability on state 0 in all other cases.

time someone learns from a successor. Then information travels to i through two different paths, one from l through the shortest path to i, and another from l to j to i. As i treats it as new information, she cannot be connected to her peer, the agent in N_l^d who passed this information to j. Otherwise she would have learned the new information from this agent directly. But then there is a simple circle involving these agents, which is impossible.

Since no signal travels back to the source, agent l's estimates are Bayesian. Similarly, learning outcomes from agents in N_l are also Bayesian since they only learn from agent l.[34] But some agent in N_l^2 must form the wrong estimates like agent k in Example 2. This is because the only possibility for an agent in N_l^2, say agent k, to stop her oscillation is to learn a specific type of information from a peer in N_l^2. But for her to infer any new information, this peer must have more connections to those in N_l than agent k does. Then we can show this peer's estimates must oscillate unless he in turn learns from another peer. There are a finite number of agents in N_l^2, and thus there must be some agent in N_l^2 who are connected to multiple agents in N_l, but does not learn from any peer. This agent's estimates keep oscillating and cannot be Bayesian.

Even for a network that is not a social quilt, it is easy to find some sequence of realized signals such that the learning outcomes are Bayesian.[35] Given Proposition 3 and 4, however, we show that social quilts are also necessary for correct learning in the following sense.

Corollary **2.** *If a network is not a social quilt, then there exists some sequence of realized signals such that the agents' learning outcomes are not strongly Bayesian.*

Since simple circles result in repetition of signals while networks without GCS may lead to estimates oscillations, agents must make one or the other type of errors for some sequence of realized signals. Thus their learning outcomes are not strongly Bayesian as long as the network is not a social quilt. It is more difficult to characterize the agents' eventual learning outcomes in a network that both contains simple circles and fails to satisfy GCS. Consider the case with one informative signal only. If this signal reaches a simple circle, then the repeatedly learned signals are positively correlated, which will reach agents in the other simple circles and non-GCS networks. If the initial signal reaches a non-GCS local network, then agents may infer offsetting copies of the initial signal, just like agent k in Example 2. These inferred signals are negatively correlated with the original signal and can be repeatedly

[34] They cannot learn from each other because if they are connected, they know each other has inferred the same signal from l.

[35] For instance, consider network (g, G) in which agent 1, 2 and 3 are connected in a triangle, and agent 3 is connected to a set of agents in a subnetwork which is not a social quilt. In the symmetric binary setting, suppose the only informative signals are $x_0^1 = 1$ for agent 1 and $x_0^2 = 0$ for agent 2. At $t = 2$, the signals offset each other and agents 1, 2 and 3 believe the state is 1 with probability 1/2, which is both the prior and the correct posterior given the signals. All other agents maintain their prior, which are also correct.

learned by those in the other subnetworks. In an arbitrary network, the numbers of positively and negatively correlated copies vary and evolve in each period. Thus we are unable to show there always exists some sequence of realized signals such that the agents' estimates are not Bayesian in networks with both features.[36]

4.4 When is our procedure consistent with Bayesian learning?

Under our learning procedure, agents behave as if all the newly inferred signals from outside their local networks are independent. We now show that, although this procedure is *not* motivated by agents holding any particular prior beliefs, it is consistent with Bayesian learning if the agents hold certain priors over the network. That is, the estimates defined in Section 2 are the correct Bayesian posterior *from each agent's perspective* given these priors.[37]

PROPOSITION **5.** *If every agent i believes her local network (g_i, G_i) is the entire network with probability* 1*, and this belief is common knowledge, then every agent learns according to our learning procedure.*

Intuitively, under these (heterogeneous) priors, each agent i believes that $\mathbf{p}_t^i$, her estimates at period t, is based on all the available signals the network received up to period $t-2$, plus her own signal at $t-1$. Moreover, for each of her neighbor j, she thinks she knows all the neighbors of j, and thus her estimates of j's estimates include all the information j has inferred so far except for his most recent signal x_{t-1}^j. Similarly, because agent i believes she knows the connections among all her neighbors, she can form estimates the same way as they do, and thus her third (and higher)-order beliefs are all correct. In particular, each agent i believes that she can account for any mistakes a neighbor makes.[38] Therefore she can learn all the informative signals while the neighbor may form the wrong estimates. In fact, Proposition 5 can be generalized: If every agent believes that the network outside her local network consists of several unconnected components, each of which is a tree-like union of

[36]With more structures on the network, we may be able to show agents' learning outcomes are not Bayesian without fully characterizing the agent's learning outcomes. For instance, consider a network with two components, one of which contains simple circles only, and the other contains non-GCS subnetworks only. There is only one link between these two components, say between agent i in the first component and j in the second. Then if an agent in a simple circle receives an informative signal, the repeated signals will reach agents in the non-GCS component through the link. But information travels away from agent j by Proposition 4. Therefore the negatively correlated copies don't reach those in the simple circles. Clearly, the agents in the simple circle cannot form Bayesian estimates.

[37]For a given sequence of realized signals, these beliefs may not be the correct Bayesian posterior from the network's perspective.

[38]For instance, agent i in Example 2 knows that agent k does not know agent k', and thus agent k's belief that her own local network ijk is the entire network is wrong. Agent i accordingly expects agent k to make mistakes and accounts for it.

cliques with the root being one of her neighbors $j \in \mathrm{N_i}$, and this belief is common knowledge, then the agents also learn according to our learning procedure.[39]

One interesting implication of this result, however, is that if agent i is truly connected to everyone else such that her local network is the entire network, then her learning under our procedure must be strongly Bayesian regardless of the links among other agents. Furthermore, if all informative signals reach the network only through this agent, then everyone's learning is strongly Bayesian. To see this, note that because everyone is connected to i, they will infer agent i's signal x_t^i simultaneously at $t+2$. Next, if two of her neighbors j and k are connected, they expect each other to learn from i, and thus will not double count i's signal. This suggests that if one wants to promote a program such as fertilizer use or microfinancing in a local community, it is better to send all the relevant information through one agent connected to all others. This helps the community to learn correctly, without being unduly influenced by a few early failures.

5 Extensions

In this section, we first consider how varying the agents' local networks may affect their learning. Next, we allow agents to treat information differently depending on its sources. Then we turn to a natural rule-of-thumb agents may use to account for some correlation of information arriving from outside their local networks.

5.1 Expanding local networks

So far, every agent is assumed to know only her immediate neighbors and the links among them. In reality, they may know less or more about their neighbors. We now accommodate this possibility by shrinking or expanding an agent's local network accordingly. We also show simple circles and non-GCS networks remain the key impediments to agents' learning.

We first shrink the local network. Abusing notation slightly, we use (g_i^0, G_i^0) to denote the case when agents only know their neighbors, but not the links among them. This is the minimum information agents need about their local networks to learn from their neighbors.

[39]Consider, for instance, one such component, which is linked to agent $j \in \mathrm{N}_i$ via some agent $k \notin \mathrm{N}_i$. By Proposition 2, agent j can learn all the signals received by agents in this component via agent k without repetition or distortions. Also any signal k infers from j must travel to every agent of this component and never comes back to j. Therefore there is no loss to treat signals received by the component as exogenous signals agent j receives from nature. So Proposition 5 holds in this case as well. This types of beliefs are supported by Fainmesser and Goldberg (2016) who show that when the population is large and the number of each agent's neighbors is bounded, each agent believes asymptotically the network is a random tree where she is the root agent.

That is, $g_i^0 = \mathrm{N}_i \cup \{i\}$ as before, and $G_i^0 = \{ij : ij \in G\}$. Moreover, the shared local 0-network (g_{ij}^0, G_{ij}^0) is defined as $g_{ij}^0 = g_i^0 \cap g_j^0$, $G_{ij}^0 = G_i^0 \cap G_j^0 = \{ij\}$. To see the difference, recall the triangle in Example 1: $g = \{1,2,3\}$ and $G = \{12, 23, 13\}$. Here $g_1^0 = g_2^0 = g_{12}^0 = g$. But $G_1^0 = \{12, 13\}$ because agent 1 does not know agent 2 and 3 are linked. Similarly, $G_2^0 = \{12, 23\}$, and thus $G_{12}^0 = \{12\}$.

When the agents use our procedure to learn, the main difference is that in forming agent i's second-order estimates about agent j, she only uses reports from i and j since $G_{ij}^0 = \{ij\}$. It is straightforward to see the agents don't form any third-order (or higher-order) estimates since $G_{ijk}^0 = G_i^0 \cap G_j^0 \cap G_k^0 = \emptyset$ for any fully connected agents i, j and k. The counterpart of Proposition 2 and Corollary 2 is as follows.

Proposition 6. *Suppose that every agent i knows (g_i^0, G_i^0) only. Their learning outcomes are Bayesian for all sequences of realized signals if and only if the network is a tree.*

For sufficiency, note that in a tree, there is a unique path—not just a unique shortest path as in social quilts—between any pair of agents, and each signal reaches an agent through this path. Each agent can infer a signal correctly from her neighbor. Once an agent infers a signal, it cannot travel back and reach her again because a tree contains no circles of any kind. Thus the agents' learning outcomes are strongly Bayesian and they finish learning at $T + D$. For necessity, first observe that if a network is not a tree, then it must contain some simple circles or triangles, which make agents' learning outcomes not Bayesian. Consider a triangle ijk. Agent i receives the only initial signal x_0^i. Agent j and k both infer the signal at $t = 2$. But agent j does not know k and i are connected. He thinks that k has received an independent signal, and vice versa for k. Their estimates are not Bayesian from $t = 3$ onward. Moreover, when agents only know (g_i^0, G_i^0), all networks satisfy GCS because each agent i's shared local 0-network contains only one link.[40] This implies that there is no negative correlation among the inferred signals, and thus no agent forms oscillating estimates as in Example 4. Since the network is path-connected, as time goes on, all agents in this network believe there are infinitely many copies of x_0^i, which is clearly not Bayesian.

We now expand the agents' local networks to any $d \geq 2$. Because signals travel faster, we modify the definition of strongly Bayesian learning outcomes accordingly:

$$q_{t+1}^i(n, d) = \Pr\left(s_n \middle| X^1_{t-\lceil d(i1)/d \rceil}, \ldots, X^L_{t-\lceil d(iL)/d \rceil}\right).$$

That is, for all l with $d(il) \in (0, d]$, agent i learns their signals one period after these signals are incorporated into l's report. If $d(il) \in (d+1, 2d]$, agent i learns their signals two periods

[40]In particular, (g_i^0, G_i^0) satisfies LCS, because agents with links in G_{ij}^0 are fully connected for all $j \in \mathrm{N}_i$.

afterwards, and so on. We call the learning outcomes *d-strongly Bayesian* if $\mathbf{p}_t^i = \mathbf{q}_t^i(d)$ for all i and t, where $\mathbf{q}_t^i(d) = \{q_t^i(1,d), \ldots, q_t^i(n,d)\}$.

We assume as before that agent i can observe all the agents within distance d of her, and the links in the original network among them. We proceed to define each agent's *local d-network*: (g_i^d, G_i^d). Recall that $d(ii) = 0$, let $g_i^d = \{l : d(il) \leq d\}$. Then $g_i^d \setminus \{i\}$ are the agents with whom agent i exchanges reports and learns from. Agent i's observational d-network, G_i^d, consists of two parts. The first part includes all the links among agents in g_i^d: if $j,k \in g_i^d$ and $jk \in G$, the link $jk \in G_i^d$. Next, it also includes each pseudo link $\widehat{lh}$ for any two agents $l,h \in g_i^d$ such that $lh \notin G$, but there exists a path of distinct agents $\{l_0 \ldots l_z\} \in g_i^d$ such that $l_0 = l, l_z = h$ and $z \leq d$. That is, if such a path exists between agent l and h,

$$G_i^d = \left\{ jk \in G : j,k \in g_i^d \right\} \cup \left\{ \widehat{lh} : lh \notin G \text{ and } \{l, l_1, \ldots, l_{z-1}, h\} \in g_i^d \right\}.$$

Intuitively, $\widehat{lh} \in G_i^d$ means that agent i knows that agent l and h can observe and learn from each other, and she includes this information when she updates her estimates. We assume that every agent i knows (g_i^d, G_i^d). Next, the shared local d-network between i and $j \in g_i^d \setminus \{i\}$ is (g_{ij}^d, G_{ij}^d), where $g_{ij}^d = g_i^d \cap g_j^d$, and $G_{ij}^d = G_i^d \cap G_j^d$. Similarly for all the higher-order shared local d-network. To illustrate, consider a diamond: $g = \{1,2,3,4\}$, and $G = \{12, 23, 34, 14\}$. Agent 1's local 2-network consists of $g_1^2 = \{1,2,3,4\}$ and $G_1^2 = \{12, 23, 34, 14, \widehat{13}, \widehat{24}\}$, because agent 1 knows agent 1 and 3, and 2 and 4 can learn from each other. That is, agent i's observational 2-network is a four-agent clique. Similarly, agent 2's local 2-network consists of $g_2^2 = \{1,2,3,4\}$ and $G_2^2 = G_1^2$. Their shared local 2-network is: $g_{12}^2 = \{1,2,3,4\}$, and $G_{12}^2 = G_1^2 = G_2^2$. They both know each other can observe all four agents' reports.

As in the main model, each agent observes all the reports from others in her local d-network. Similarly, each agent forms estimates and higher-order estimates of her neighbors in the local d-network, and she updates using their reports as well. Our procedure is modified accordingly. Clearly, at $t = 0$ and $t = 1$, each agent i's learning remains unchanged. For all $t \geq 1$, she first identifies new information from each neighbor in g_i^d as before. But when she forms her second-order estimates $\mathbf{p}_t^{ij}$, she uses reports from i, j, and the agents who are common neighbors of i and j based on G_{ij}^d. That is, if both ik (or $\widehat{ik}$) and jk (or $\widehat{jk}$) are in G_{ij}^d, then k is a common neighbor of i and j. Similarly, her higher-order estimates $\mathbf{p}_t^{ii_1 \ldots i_l}$ are well-defined if the distinct agents $\{i, i_1, \ldots, i_l\} \subset g_i^d$ are fully connected based on $G_{ii_1 \ldots i_l}^d$.[41] She

[41]This is different from the $d = 1$ case, where G_i as defined before is agent i's observational 1-network. To see this, note that if i's two neighbors $i_k i_h$ are connected, agent i knows it when $d = 1$. But when $d \geq 2$, it is possible that i_k and i_h can see each other, but i does not know they can. For example, consider a simple circle $\{123456\}$, agent 1 does not know 3 and 5 know each other even though they do.

updates $\mathbf{p}_t^{ii_1\dots i_l}$ using reports from $\{i, i_1, \dots, i_l\}$ and the agents who are common neighbors of all $\{i, i_1, \dots, i_l\}$ based on $G^d_{ii_1\dots i_l}$. Clearly most properties of our learning procedure still hold. For example, as in Lemma 1, the order of agents in the higher-order estimates still does not matter. In particular, any pair of agents would agree, $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ji}$ for all $i \in g$, $j \in g_i^d \setminus \{i\}$ and $t \geq 1$. Also, signals can still be decomposed.

Because social quilts include trees as a special case, agents are able to learn correctly in more networks when their knowledge expands from (g_i^0, G_i^0) to (g_i, G_i). A natural hypothesis is that as each agent's local network continues to expand, for instance, to (g_i^2, G_i^2), the agents' learning becomes Bayesian in a larger set of networks. This turns out to be false.

Proposition 7. *If all agents know their local d-networks ($d \geq 2$), the agents' learning outcomes are d-strongly Bayesian for all sequences of realized signals if and only if $D \leq d$.*

If $D \leq d$, then from each agent's perspective (g_i^d, G_i^d) is a clique, and thus the learning outcomes are strongly Bayesian by Proposition 2. If $D > d$, the network does not satisfy GCS, and thus some agents must make mistakes. Consider a linear chain of four agents 1234 and $d = 2$. Then $g_1^2 = \{1, 2, 3\}$ and $G_1^2 = \{12, 23, \widehat{13}\}$; $g_2^2 = g_3^2 = g$ and $G_2^2 = G_3^2 = \{12, 23, 34, \widehat{13}, \widehat{24}\}$; $g_4^2 = \{2, 3, 4\}$ and $G_4^2 = \{23, 34, \widehat{24}\}$. For each agent, it is observationally equivalent to a diamond with a link between 2 and 3. Therefore, similar to Example 2, if agent 1 receives the only signal x_0^1, agent 4's estimates oscillate forever from $t = 3$ onwards. Proposition 7 thus shows there is a non-monotonicity when agents can observe more reports from their local d-network. If $D \leq d$, all agents' learning outcomes are strongly Bayesian, which cannot happen if $D > d$ at least for some sequence of realized signals.

In addition, characterizing learning outcomes when $D > d$ becomes more involved because a network may contain both simple circles and non-GCS subnetworks even though it only has one feature previously.[42] But in networks that have only one feature, we can show some agents' learning outcomes are not Bayesian for some sequence of realized signals similar to Proposition 4. This is the case in a linear chain with $D > d > 1$, which is a network of non-GCS subnetworks only. Observing more agents, however, may help locally: if $D > d_2 > d_1 > 1$, fewer agents make mistakes in learning when they observe the local d_2-network than the local d_1-network.

Intuitively, an agent's learning may deteriorate if she observes more reports from those who have received the signal from the same source because she cannot differentiate correlated

[42] For instance, consider a simple circle $\{123456\}$. With $d = 2$, the network features both simple circles and non-GCS subnetworks. First, this network contains a simple circle with $k = 3$, and thus if agent 1 has the only informative signal x_0^1, it is easy to see that at $t = 4$, agent 2 and 3 respectively learn two and one new copy of the signal from 4, which came indirectly from agent 1. Therefore agents still repeatedly infer the same signals due to the simple circle. In addition, agent 1 and 4 both can observe 2 and 3, but they do not know each other. Similar to the linear chain example above, their local d-networks do not satisfy LCS.

information. But we can view this problem through a different lens: The agents' local networks expand and become more complex if $D > d$ without any corresponding expansion in their communication possibilities. One possible solution is to allow agents to tag information locally, that is, sharing when a signal comes from the same source within their local networks. The following is one way to tag their signals. In addition to exchange reports about their first-order estimates, each agent may also report a message $\mathbf{m}_t^i = \{i, j, k, \ldots; x\}$, meaning that the reports from agent i, j, k and others at period t contain a signal x learned from the same source. This message is local: It can only be observed and understood by agents in the issuing agents' local d-network. That is, if agent l only know agent j and k, she can understand j and k learned the signal from one source, but not agent i also learned it because the part involving i does not mean anything to her. With locally tagged signals, we can show that, similar to Proposition 2, agents's learning outcomes in a d-social quilt, which is a tree with each node being a subnetwork with a diameter less than or equal to d, are strongly Bayesian.[43]

To illustrate, consider a simple network of five agents on a line: $\{12345\}$, where $d = 2$ and x_0^1 is the only signal. At $t = 1$, agent 1 reports as before. At $t = 2$, agent 2 and 3 respectively report $\mathbf{p}_2^2 = \{p_2(1), \ldots, p_2(n)\}$ where $p_2(n) = \Pr(s_n|x_0^1)$ for all s_n. Each agent also reports $\mathbf{m}_2^1 = \mathbf{m}_2^2 = \mathbf{m}_2^3 = \{1, 2, 3; x_0^1\}$. At $t = 3$, agent 4, who are within g_{23}^2, can observe both agent 2 and 3's reports and messages. Therefore he treats only one inferred signal as new and his estimates $\mathbf{p}_3^4$ are correct: $p_3^4(n) = \Pr(s_n|x_0^1)$ and $\mathbf{m}_3^4 = \{2, 3, 4, 5; x_0^1\}$ because 4 knows 5 learns from 3. Agent 5, however, does not know the existence of agent 2, and thus 3's message has no meaning to her. She observes agent 3's report and updates as before to $p_3^5(n) = \Pr(s_n|x_0^1)$ and $\mathbf{m}_3^5 = \{3, 4, 5; x_0^1\}$. At $t = 4$, all agents report the same estimates and their higher-order estimate also agree. Learning stops and all agents' estimates are Bayesian.

5.2 When all information is not equal

So far, the agents treat all information equally regardless of the sources and the arrival times. Yet for various reasons, agents may weigh their inferred signals differently. For instance, agents may trust some of their neighbors more than others; they may also discount new information as time goes on.[44] We now extend our model by letting the agents put

[43] Another possible solution is to allow agents to reduce the number of neighbors they communicate with. Though counterintuitive, this is consistent with the recent findings from Harel et al. (2014) and Alatas et al. (2016). Alatas et al. (2016) show that, controlling for other network characteristics, having a higher average number of connections has a negative effect on information aggregation.

[44] Doing so may improve their learning outcomes. Grimm and Mengel (2014) found, for instance, that subjects who earned above median payoffs reduce how much weight they put on their neighbors' information

different weights on their inferred signals.

Let $w_t^{ij} \geq 0$ be the weight agent i puts on her inferred signal $\boldsymbol{\alpha}_t^{ij}$ at time $t+1$. Previously, $w_t^{ij} = w_t^{ji} = 1$ for all the connected pairs ij and t. We continue to assume each agent i attaches a weight of $w_t^{ii} = 1$ to her own signals. In terms of her neighbors, if $w_t^{ij} > 1$ (resp. $w_t^{ij} < 1$), then agent i thinks inferred signals from agent j is more (resp. less) important than her own signals. If $w_t^{ij} = 0$, we assume that while agent i can still observe j's report, she thinks his information is useless.[45] More importantly, we assume that every agent not only knows her local network, she also knows the weights each neighbor j uses on their common neighbors in g_{ij}. Namely, for any agent $j, k \in \mathrm{N}_i$, i knows the weights her neighbors put on her information (w_t^{ji} and w_t^{ki}); and the weights they put on each other if they are connected (w_t^{jk} and w_t^{kj}). This makes it possible for agents in a shared local network to infer the same set of signals even when they attach different importance to these inferred signals.

Our learning procedure is accordingly modified. For every agent i, there is no change in each agent i's learning at $t = 0$ and $t = 1$. For all $t \geq 1$, agent i forms her (higher-order) estimates and exchanges reports with her neighbors as before. In Step 1, agent i continues to identify all her inferred signals by using expression (3) and (4).

In Step 2, agent i updates her own estimates using her private signal with a weight $w_t^{ii} = 1$, and the inferred signals from her neighbors with weights w_t^{ij} for each $j \in \mathrm{N}_i$. The counterpart of expression (5) becomes:

$$p_{t+1}^i(n) = \frac{p_t^i(n) \prod_{h \in g_i} (\alpha_t^{ih}(n))^{w_t^{ih}}}{\sum_{n'=1}^N p_t^i(n') \prod_{h \in g_i} (\alpha_t^{ih}(n'))^{w_t^{ih}}}. \tag{15}$$

Intuitively, agent i behaves as if she learns w_t^{ih} independent copies of $\boldsymbol{\alpha}_t^{ih}$. It is a natural formulation because all the inferred signals enter multiplicatively into agent i's estimates via Bayes' rule.[46]

In Step 3, when forming estimates of neighbor j's estimates, agent i starts with agent j's latest estimates $\mathbf{p}_t^j$ and incorporates the new information i thinks j has learned, $\alpha_t^{ijh}(n)$

over time.

[45]The alternative formulation is that a weight of 0 means agent i does not observe j's report. Doing so complicates our model without adding much insight because agent i cannot just ignore j. She still needs to form higher-order estimates involving j to keep track of the existing information, for instance information from one of their common neighbors inferred by j and then passed on to another common neighbor.

[46]Interestingly, this way of introducing weights is similar to that in the widely-used machine learning literature in computer science. The goal there is to classify test documents into categories, where the words are assumed to occur independently and the weights are the number of times a word occur. Although the independence assumption can be too strong when certain words often appear together in a given context, this method works very well empirically. For instance, see the survey "Naive (Bayes) at Forty: The Independence Assumption in Information Retrieval" in Lewis (1998) for more details.

with a weight w_t^{jh}. The counterpart of expression (6) thus becomes:

$$p_{t+1}^{ij}(n) = \frac{p_t^j(n)\prod_{h\in(g_{ij}\setminus\{j\})}(\alpha_t^{ijh}(n))^{w_t^{jh}}}{\sum_{n'=1}^N p_t^j(n')\prod_{h\in(g_{ij}\setminus\{j\})}(\alpha_t^{ijh}(n'))^{w_t^{jh}}}. \tag{16}$$

For each subset of g_i Agent i's higher-order estimates are formed in a similar way:

$$p_{t+1}^{ii_1\ldots i_l}(n) = \frac{p_t^{i_l}(n)\prod_{h\in(g_{ii_1\ldots i_l}\setminus\{i_l\})}(\alpha_t^{ii_1\ldots i_l h}(n))^{w_t^{i_l h}}}{\sum_{n'=1}^N p_t^{i_l}(n')\prod_{h\in(g_{ii_1\ldots i_l h}\setminus\{i_l\})}(\alpha_t^{ii_1\ldots i_l h}(n'))^{w_t^{i_l h}}}. \tag{17}$$

As before, for all sequences of fully connected and distinct agents $\{i, i_1, \ldots, i_l\}$, we allow the last agent to be repeated in agent i's estimates. Specifically, let $\mathbf{p}_{t+1}^{i\ldots i_l k} = \mathbf{p}_{t+1}^{i\ldots i_{k-1}i_{k+1}\ldots i_l k}$ for any $k \in \{i_1, \ldots, i_l\}$, where $i = i_0$ if $k = i_1$. Finally, for each agent i,

$$p_{t+1}^{ii_1\ldots i_l i}(n) = \frac{p_t^{i}(n)\prod_{h\in(g_{ii_1\ldots i_l}\setminus\{i\})}(\alpha_t^{ii_1\ldots i_l h}(n))^{w_t^{ih}}}{\sum_{n'=1}^N p_t^{i}(n')\prod_{h\in(g_{ii_1\ldots i_l h}\setminus\{i\})}(\alpha_t^{ii_1\ldots i_l h}(n'))^{w_t^{ih}}}. \tag{18}$$

Thus the highest order of estimates each agent i forms is $\widehat{L}_i + 1$.

Most of the properties in Section 3 hold. the next result is the counterpart of Lemma 1 and Lemma 2. Recall that $\{\beta(1), \ldots, \beta(l-1)\}$ is a permutation of $\{1, \ldots, l-1\}$.

Corollary **3.** *For every agent $j \in \mathrm{N}_i$, agent $k \in \mathrm{N}_i \cap \mathrm{N}_j$, and every $t \geq 1$,*

(1) Consider any fully connected subset $\{i_1, \ldots, i_l\}$ of g_i. Then $\mathbf{p}_t^{i_{\beta(1)}\ldots i_{\beta(l-1)}i_l}$ is the same for all $t \geq 1$.

(2) If (g_i, G_i) satisfies LCS, then $\mathbf{p}_t^{ik} = \mathbf{p}_t^{ijk} = \mathbf{p}_t^{ij\ldots k}$ for all distinct agents $i, j, \ldots, k \in g_{ij}$.

(3) If (g_l, G_l) satisfies LCS for every agent $l \in g_i$, then $\mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$.

First, notice that from each agent i's perspective, the order of agents does not matter except for the last agent. This is because the agents' higher-order estimates about agent i_l's estimates are all based on the weights agent i_l attaches to his inferred signals. Moreover, two neighbors still know what signal one infers from the other: $\boldsymbol{\alpha}_t^{iji} = \boldsymbol{\alpha}_t^{ji}$, even though $\mathbf{p}_t^{ij} \neq \mathbf{p}_t^{ji}$ because agent i and j attach different weights on each other's information. This is important for one agent to know what each neighbor infers from herself and from their common neighbors. Next, if (g_i, G_i) satisfies LCS, agent i's higher-order estimates on the same agent k must be the same. This is because from agent i's perspective, each shared local network is a clique, and agents within each clique must see the same set of reports.

Since agents in a clique also know exactly the weights each other attaches to new signals, if $k \in \mathrm{N}_i \cap \mathrm{N}_j$, i thinks j must agree with her estimates of k's estimates, $\mathbf{p}_t^{ik} = \mathbf{p}_t^{ijk}$. Also, if (g_l, G_l) satisfies LCS for every $l \in g_i$, each of agent i's shared local networks is indeed a clique, therefore i and j must agree: $\mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$. In addition, because the weights are part of the known local network information, signals can be decomposed similar to Lemma 4.

We first consider how social influence, measured by the weights agents attach to each other's opinion, affect the agents' learning outcomes when the network is a social quilt.[47] The following result extends Proposition 2.

Corollary 4. *If the network is a social quilt, learning stops at time $T + D$. Agent i's estimates from period $T + D$ onward are, for all $t \leq T - 1$:*

$$p^i_{T+D}(n) = \frac{\prod_t \alpha_t^{ii}(n) \prod_{l \neq i} \prod_t (\alpha_t^{ll}(n))^{w^{il}_{t+1}}}{\sum_{n'} \prod_t \alpha_t^{ii}(n') \prod_{l \neq i} \prod_t (\alpha_t^{ll}(n'))^{w^{il}_{t+1}}}, \tag{19}$$

where $w^{il}_{t+1} = w^{k'l}_{t+1} w^{kk'}_{t+2} \dots w^{ij}_{t+h-1}$, and $l, k', k, ...j, i$ is the unique shortest path from l to i.

Similar to the main model, the agents learn all the signals once and only once if all the weights are positive. But here the agents may not reach consensus. Observe that if all the weights are 1, expression (19) is exactly the Bayesian posterior given all the signals X_T. Using Corollary 4, we can see how an agent's social influence affects the learning outcomes in a social quilt. First, suppose agent l is a *local opinion leader* and agent j her *follower* if agent l's neighbors put a high weight on her information while agent l puts a low weight on theirs. That is, for any $j \in \mathrm{N}_l$, $w_t^{jl} = \frac{1}{\varepsilon}$ and $w_t^{lj} = \varepsilon$ for some small $\varepsilon > 0$. For simplicity, suppose agents use weights of 1 in all other links except for those between l and $j \in \mathrm{N}_l$. A local opinion leader can unduly influence the opinions of the entire network: all agents but l believes in the state(s) most likely given agent l's signals. This is because all agents overweight their inferred signals from l, directly or indirectly through her followers. Next, agent l is *stubborn* if she does not listen to anyone: $w_t^{lj} = \varepsilon$ for all $j \in \mathrm{N}_l$. A stubborn agent cannot bias the network's learning outcomes toward her opinion, but she may cause fragmentation by blocking the efficient aggregation of information.[48]

The weights may also capture the idea of imperfect information diffusion. Suppose $w < 1$ for everyone because information may not reach another agent with some probability. From expression (19), we can see that in a social quilt, agent i's weight on agent l's signal, which

[47] In the information partition model, clearly the weights do not matter.

[48] In comparison, under myopic learning, all agents are mislead by the stubborn agent. In our model, agents know their information is undervalued by the stubborn agent, so they do not infer much new information from the stubborn agent's unchanging estimates.

is $w^{d(il)}$, depends on their distance only. As a result, opinions may become polarized: Agents close to each other are similar in their opinions, but disagreement grows in their distance.

In networks that are not social quilts, Section 4.3 shows that agents may never stop learning because they keep inferring "new" signals due to the network structure. The next result shows that agents may want to discount later signals as a rudimentary way to reduce the correlations in their neighbors' estimates, which is supported by experimental evidence.[49] Recall that L_i is the number of agent i's neighbors, then we have:

PROPOSITION **8.** *Suppose that each agent i's weights w_t^{ij} are below a cutoff value $\overline{w}^i < \frac{1}{2L_i}$ for all agents $i \in g, j \in \mathrm{N}_i$ from some time τ onward. Then the agents' estimates must converge: $\lim_{t\to\infty} \mathbf{p}_t^i = \mathbf{p}_{t\to\infty}^i$. Moreover, $p_{t\to\infty}^i(n) \in (0,1)$ for all s_n.*

One important consequence of adding weights is that if agents attach a sufficiently small weight to their inferred signals, then convergence can be restored in any network. Proposition 8 is clearly true if all the weights are zero, in which case every agent is isolated and her estimates are simply the Bayesian posterior given her own signals. But this intuition is incomplete because we have shown that as long as all the weights are positive, signals may travel through a network repeatedly, possibly at an exponential rate. To rule out the possibility that signals inferred later are sufficiently influential and cause oscillation in the agents' estimates, we need to show their influence is bounded and decreasing regardless of the network structure. Recall that we consider the case of non-partitional signals, the agents' estimates and their higher-order estimates at time T are bounded because they only receive a finite number of signals. So are the agents' (higher-order) inferred signals at time T. In each ensuing period, each inferred signal exerts a smaller influence on agent i's estimates. We can thus find some period t sufficiently larger than T such that after t, the inferred signals have negligible influence. Consequently, every agent's estimates converge, and the limit estimates are strictly between 0 and 1.

Several remarks are in order. First, Proposition 8 shows that adding weights may moderate the opinions of the agents. Otherwise, they may believe in some wrong state almost surely as time goes on as in the case of Corollary 1. Second, agents with more connections may need to discount her neighbor's information more than those with fewer connections, because the former has more chances of being misled by repetition of the same information. Third, if agents attach a sufficiently small weight on information from their neighbors, then they can learn when they receive an arbitrarily large number of signals as $T \to \infty$.

[49] Celen and Kariv (2004) and Grimm and Mengel (2014), among others, found subjects weigh their own information more than their neighbors' information.

5.3 Correlations in inferred signals

While the inferred signals are indeed independent in social quilts, they are clearly correlated in other networks such as those with simple circles. One often suggested solution to information repetition problem is to consider an information environment where all signals are generic. That is, if no agents observe identical signals due to individual idiosyncrasies, and the agents know that, they would dismiss identical copies of a previously learned signal as old information, and thus avoid double-counting the same signal. We now examine this environment and show the potential pitfalls of using simple rule-of-thumbs like these.

We assume every agent learns her initial information structure before any signal is realized. In each period $t = 0, 1 \ldots, T-1$, her information structure evolves such that the probability of agent i observing signal x^i_m given state s_n in period t is $\phi^{it}_{mn} = \phi^i_{mn} + e^i_t$, where $e^i_t \in U[-\varepsilon^i_t, \varepsilon^i_t]$ for a sufficiently small $\varepsilon^i_t > 0$.[50] All the noises are independent. Consequently, even if all agents have identical initial information structures, the probability that two signals have identical conditional distributions is zero. Furthermore, every agent knows this.[51]

Each agent accounts for repeated information by keeping track of a set of signals that she has already learned, denoted as $\mathbf{A}^i_t$, which evolves as i learns new information in period t. Agents use the simplest rule to avoid repetition: Each time agent i infers a new signal $\boldsymbol{\alpha}^{ij}_{t-1} \notin \mathbf{A}^i_t$ in period t, she updates her estimates using $\boldsymbol{\alpha}^{ij}_{t-1}$ and then stores it inside $\mathbf{A}^i_t$. But if the inferred signal $\boldsymbol{\alpha}^{ij}_{t-1} \in \mathbf{A}^i_t$, she dismisses it as uninformative. She also puts her own signal $\boldsymbol{\alpha}^{ii}_t$ into $\mathbf{A}^i_t$ after incorporating it into her estimates. We assume that every agent uses this simple rule.

To implement this, our learning procedure is modified accordingly. At $t = 1$, all the reports reflect new and independent signals, so the initial estimates remain the same. Then at each period, agent i first infers all the new signals as in Step 1. Next, she reviews them one by one by the above rule, only keeping those inferred signals not already in her set $\mathbf{A}^i_t$. After reviewing all the inferred signals, she updates her estimates as in Step 2 of our learning procedure, using only the inferred signals she thinks are new. More importantly, she does the same for her higher-order estimates by keeping track of the new signals she thinks agent j infers, $\mathbf{A}^{ij}_t$. She also reviews all the signals she thinks j infers and dismisses those she thinks j would dismiss. She then updates $\mathbf{p}^{ij}_t$ using only the new signals she thinks j infers from their common neighbors, and so on.

This simple rule can reduce some errors in the agents' learning such as those in Example

[50]This is well defined since we focus on nonpartitional information structures in which all $\phi^i_{mn} \in (0, 1)$.

[51]The signal generating process of each agent is still the agent's private information. The additional common knowledge is that every agent knows the realized signals are subject to small amount of noises and thus cannot be identical. This rules out, for instance, the familiar symmetric binary signal case.

2 when agent k' receives an informative signal. Clearly, agent k infers the same signal from both agent i and j at $t = 2$, and treats only one of them as new. Therefore everyone in that network learns correctly. The following result generalizes this intuition.

PROPOSITION **9.** *If only one agent receives all informative signals, the agents' learning outcomes are strongly Bayesian in any network.*

If all signals reach the same agent, they diffuse sequentially throughout the network. Each agent infers (possibly multiple copies of) a signal at one time, therefore they can use the simple rule to identify correlated signals. More precisely, suppose agent i receives one initial signal x_0^i. Classify her neighbors according to $d(ij)$, the distance to her. Her immediate neighbors learn the signal correctly at $t = 2$. Next, for her indirect neighbors, when the signal first travels to a neighbor l at time $t = d(il) + 1$, either agent l learns it from one neighbor, in which case he infers the signal and passes it on to others; or he learns it from multiple neighbors. This can happen either because the signal travels through a circle and reaches him from different directions; or in the diamond-with-a-link case the signal reaches him from an unknown common source. Either way, he incorporates only one copy. Because each agent at any moment incorporates at most one copy of the signal which reaches him through the shortest path between him and the source of the information, everyone's estimates are strongly Bayesian. This logic holds when the same agent receives multiple signals over time because each informative signal reaches other agents sequentially and thus can be identified individually. For the same reason, we can show that in a simple circle of k agents with initial signals only, all agents' learning outcomes are strongly Bayesian.

In general, however, problems such as internal inconsistency arise when agents use simple rules like the above. This is because each agent's set of stored signals for a neighbor can differ from the actual set of stored signals of that neighbor. Consider the following example.

EXAMPLE **6.** *Consider a simple circle with four agents* $\{1, 2, 3, 4\}$*. There are two informative signals: Agent* 1 *receives signal* x_0^1 *at* $t = 0$ *and agent* 2 *receives* x_1^2 *at* $t = 1$.

To ease exposition, define three new signals x, y, z such that given the symmetric prior, $\Pr(s_n|x) = \Pr(s_n|x_0^1)$, $\Pr(s_n|y) = \Pr(s_n|x_1^2)$ and $\Pr(s_n|z) = \Pr(s_n|x_0^1, x_1^2)$ for all s_n. Further, let $-x$ be a signal such that $\Pr(s_n|x, -x) = 1/N$.[52] The correct Bayesian posterior given the signals is just $\Pr(s_n|x, y)$.

To begin with, at $t = 1$, only agent 1 has an informative signal, therefore $\mathbf{A}_1^1 = \{x\}$ and $\mathbf{A}_1^2 = \mathbf{A}_1^3 = \mathbf{A}_1^4 = \emptyset$. At $t = 2$, agent 2 learns from 1 and also learns her private signal. Agent 4 only learns from 1. Thus $\mathbf{A}_2^2 = \{x, y\}$. $\mathbf{A}_2^1 = \{x\}, \mathbf{A}_2^4 = \{x\}$ and $\mathbf{A}_2^3 = \emptyset$.

[52] Because even though the agent's inferred signals are not identical to the original signals as discussed in Section 2, the ensuing estimates are the same. So we use these three signals to avoid referring to each agent's inferred signals when they lead to the same estimates.

At $t = 3$, agent 3 infers z, the combination of (x_0^1, x_1^2), from 2 because she does not know the two separate signals. She also infers x from 4. As they are different, both are treated as new information. Therefore $\mathbf{A}_3^3 = \{z, x\}$. Observe, however, $\mathbf{A}_3^{32} = \{z\} \neq \mathbf{A}_3^2 = \{x, y\}$. In addition, $\mathbf{A}_3^1 = \{x, y\}$, $\mathbf{A}_3^2 = \{x, y\}$ and $\mathbf{A}_3^4 = \{x\}$. At $t = 4$, agent 4 infers y from 1 and z from 3. As they are both different from 4's old information x, $\mathbf{A}_4^4 = \{x, y, z\}$. Other agents' sets of stored signals do not change.

At $t = 5$, agent 1 infers z from 4, and agent 3 infers y from 4. Notice that agent 3 expects agent 2 to learn x from himself, $\mathbf{A}_5^{32} = \{z, x\}$, but agent 2 already knows x and thus does not change. This makes agent 3 believe agent 2 received an offsetting signal $-x$ from agent 2, so $\mathbf{A}_5^3 = \{z, x, y, -x\}$. Agent 2, however, expects 3 to learn $-x$ because she knows that 3 does not know she learned x before. That is, $\mathbf{A}_5^{23} = \{z, x, -x\}$. At the end of this period, $\mathbf{A}_5^1 = \{x, y, z\}$, $\mathbf{A}_5^2 = \{x, y\}$ and $\mathbf{A}_5^4 = \{x, y, z\}$.

At $t = 6$, agent 2 infers z from 1. Since she only knows x and y as two separate signals, her new list becomes $\mathbf{A}_6^2 = \{x, y, z\}$. Also, $\mathbf{A}_6^1 = \{x, y, z\}$, $\mathbf{A}_6^3 = \{z, x, -x, y\}$ and $\mathbf{A}_6^4 = \{x, y, z, -x\}$. At $t = 7$, agent 3 infers x from 2, because she expects 2 to learn y from her while agent 2 learned z instead. Since $x \in \mathbf{A}_6^3$, 3 treats it as old information, so she does not update. Also, $\mathbf{A}_7^1 = \{x, y, z, -x\}$, $\mathbf{A}_7^2 = \{x, y, z\}$ and $A_7^4 = \{x, y, z, -x\}$.

Finally, at $t = 8$, only agent 2 learns the signal $-x$ from agent 1. All agents' sets of inferred signals agree: $\mathbf{A}_8^i = \{x, y, z, -x\}$ and the learning stops. They have consensus, but they are wrong. $p_t^i(n) = \Pr(s_n | y, z)$ for all $i = 1, 2, 3, 4$, all s_n and all $t > 8$. $\diamond$

This example first illustrates that Lemma 4 no longer holds. Clearly, agents can use the simple rule to learn each signal correctly if x or y is the only informative signal. But when both are present, their consensus is wrong. But Lemma 4 is central to obtain results for any sequence of realized signals. Moreover, the pairwise agreement in Lemma 1 also fails because agents may no longer infer the same set of signals after observing the same reports. In the above example, $\mathbf{A}_5^{23} \neq \mathbf{A}_5^{32}$. Therefore even neighbors 2 and 3, whose only shared information are each other's reports, disagree. This will lead to internal inconsistency. Agent 3, for instance, notices that at $t = 7$, agent 2 seems to learn another copy of x, which agent 2 already learned at $t = 4$ from agent 3's perspective. Therefore agent 3 is confused and does not know if agent 2 truly follows the simple rule.

Despite the intuitive appeal of using simple rules to account for correlations in information, they may lead to learning problems in even very simple networks. Once an agent realizes an inconsistency as illustrated above, further learning becomes difficult. Although it is possible to ask agents to ignore such inconsistency as a further simple rule, or to develop more sophisticated rules, this type of internal consistency problems would still arise. How to construct an internally consistent model using simple rules is a topic for further research.

6 Conclusion

We propose a simple and tractable learning procedure when agents only know their local networks. Although the agents do not form sophisticated beliefs about the entire network, they can still form the Bayesian posterior in two benchmark cases. Our procedure cuts down the heavy computational burden the agents face when the network is unknown. It lies between the fully Bayesian learning model and the myopic learning model. In appendix B.1, we show that if the network is common knowledge, our agents can learn correctly; and in appendix B.2, we show that our agents make fewer mistakes than under myopic learning.

The tractability of our learning procedure makes it a suitable building block for more complicated network models. For instance, agents may believe that their local networks are the entire network with a high probability, and that her neighbors are connected indirectly with small probabilities. Initially, agents learn using our procedure. But they also update their beliefs about the network, especially if they accumulate enough evidence suggesting that their neighbors may be connected indirectly and there is repeated information. One challenge in updating both one's estimates about the state and one's beliefs about the network is that signals may no longer be decomposable. That is, even for a fixed sequence of realized signals, an agent's estimates given all the signals may be different from the combination of her estimates if she has received the signals individually. Stronger conditions are needed to obtain general results in more complex network models.

In terms of applications, our results suggest that to improve learning in networks, how and where to inject information is an important question. For example, when the network is tightly connected such as a rural village, some agent may be connected to everyone else in the network. Then sending all the informative signals to this central agent can guarantee the whole village learn correctly. For another example, suppose that the agents try to reduce information repetition by incorporating each new signal only once as in section 5.3, then Proposition 9 suggests that injecting all signals through one agent can guarantee strongly Bayesian learning.

Our model also suggests that to help the isolated communities in our introduction, policy responses should not be limited to information campaigns and awareness-raising. As shown in Example 5, if a neighborhood is connected through a sequence of interlinked simple circles, an early piece of wrong information is not offset by the many pieces of correct information later—wrong beliefs persist despite the new information. Rather, changes to the network structures such as building links with the outside world become necessary.

References

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar. 2014. "Dynamics of information exchange in endogenous social networks." *Theoretical Economics*, 9(1): 41–97.

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. 2011. "Bayesian Learning in Social Networks." *Review of Economic Studies*, 78(4): 1201–1236.

Adamic, Lada A. 1999. "The Small World Web." *Proceedings of the ECDL, Lecture Notes in Computer Science*, 1696: 443–452.

Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken. 2016. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *American Economic Review*, 106(7): 1663–1704.

Alevy, Jonathan E., Michael S. Haigh, and John A. List. 2007. "Information Cascades: Evidence from a Field Experiment with Financial Market Professionals." *Journal of Finance*, 62: 51–180.

Ali, S. Nageeb, and David A. Miller. 2013. "Enforcing Cooperation in Networked Societies." *working paper.*

Anderson, Lisa R., and Charles A. Holt. 1997. "nformation Cascades in the Laboratory." *American Economic Review*, 87: 847–862.

Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics*, 4(6): 1236–1239.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbours." *The Review of Economic Studies*, 65(3): 595–621.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107: 797–819.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy*, 10(1): 992–1026.

Cai, Hongbin, Yuyu Chen, and Hanming Fang. 2009. "Observational Learning: Evidence from a Randomized Natural Field Experiment." *American Economic Review*, 99(3): 864–882.

Casciaro, Tiziana. 1998. "Seeing things clearly: social structure, personality, and accuracy in social network perception." *Social Networks*, 20(4): 331–351.

Celen, Bogachan, and Shachar Kariv. 2004. "Distinguishing Informational Cascades from Herd Behavior in the Laborator." *American Economic Review*, 94: 484–498.

Chandrasekhar, Arun G., Emily Breza, and Alireza Tahbaz-Salehi. 2016. "Seeing the forest for the trees? an investigation of network knowledge." *working paper.*

Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri. 2012. "Testing Models of Social Learning on Networks: Evidence from a Framed Field Experiment." mimeo.

Conley, Timothy, and Christopher Udry. 2001. "Social Learning through Networks: The Adoption of New Agricultural Technologies in Ghana." *American Journal of Agricultural Economics*, 83(3): 668–673.

DeGroot, Morris H. 1974. "Reaching a Consensus." *Journal of the American Statistical Association*, 69(345): 118–121.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions." *Quarterly Journal of Economics*, 118(3): 909–968.

Duflo, Esther, and Emmanuel Saez. 2002. "Participation and investment decisions in a retirement plan: The influence of colleagues choices." *Journal of Public Economics*, 85: 121–148.

Ellison, Glenn, and Drew Fudenberg. 1993. "Rules of Thumb for Social Learning." *Journal of Political Economy*, 101(4): 612–643.

Ellison, Glenn, and Drew Fudenberg. 1995. "Word-of-Mouth Communication and Social learning." *Quarterly Journal of Economics*, 110(1): 93–125.

Enke, Benjamin, and Florian Zimmermann. 2015. "Correlation Neglect in Belief Formation." mimeo.

Eyster, Erik, and Matthew Rabin. 2014. "Extensive imitation is irrational and harmful." *The Quarterly Journal of Economics*, 129(4): 1861–1898.

Fainmesser, Itay P, and David A Goldberg. 2016. "Cooperation in partly observable networked markets." *working paper.*

Gale, Douglas, and Shachar Kariv. 2003. "Bayesian learning in social networks." *Games and Economic Behavior*, 45(2): 329–346.

Geanakoplos, John D., and Heraklis M. Polemarchakis. 1982. "We can't disagree forever." *Journal of Economic Theory*, 28(1): 192–200.

Golosov, Mikhail, Nancy Qian, and Yan Kai. 2015. "Understanding the Influence of Government Controlled Media: Evidence from Air Pollution in China." *working paper.*

Golub, Benjamin, and Matthew O. Jackson. 2010. "Naive Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics*, 2(1): 112–49.

Goyal, Sanjeev, Marco J. van der Leij, and Jose Luis Moraga-Gonzlez. 2006. "Economics: An Emerging Small World." *Journal of Political Economy*, 114(2): 403–412.

Grimm, Veronika, and Friederike Mengel. 2014. "Experiments on Belief Formation in Networks." mimeo.

Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz. 2014. "When more information reduces the speed of learning." *working paper.*

Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein. 2005. "Thy neighbors portfolio: Word-of-mouth effects in the holdings and trades of money managers." *Journal of Finance*, 60: 2801–2824.

Ioannides, Yannis M., and Linda Datcher Loury. 2004. "Job Information Networks, Neighborhood Effects, and Inequality." *The Journal of Economic Literature*, 42(2): 1056–1093.

Jackson, Matthew O., Tomas Rodriguez-Barraquer, and Xu Tan. 2012. "Social Capital and Social Quilts: Network Patterns of Favor Exchange." *American Economic Review*, 102(5): 1857–1897.

Krackhardt, David. 1990. “Assessing the Political Landscape: Structure, Cognition, and Power in Organizations.” *Administrative Science Quarterly*, 35(2): 342–369.

Lee, In-Ho. 1993. “On the Convergence of Informational Cascades.” *Journal of Economic Theory*, 61: 395–411.

Lewis, David D. 1998. “Machine Learning: ECML-98.” , ed. Claire Nedellec and Celine Rouveirol, Chapter 2, 4–15. Springer Berlin Heidelberg.

Lones, Smith, and Peter Sorensen. 2000. “Pathological Outcomes of Observational Learning.” *Econometrica*, 68(2): 371–398.

MacRae, Duncan. 1960. “Direct Factor Analysis of Sociometric Data.” *Sociometry*, 23(4): 360–371.

Mobius, Markus, Tuan Phan, and Adam Szeidl. 2015. “Treasure Hunt: Social Learning in the Field.” *NBER Working Paper No. 21014*.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie. 2016. “Foundations of non-bayesian social learning.” *working paper*.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2015. “Strategic Learning and the Topology of Social Networks.” *Econometrica*, 83(5): 1755–1794.

Mueller-Frank, Manuel. 2013. “A general framework for rational learning in social networks.” *Theoretical Economics*, 8(1): 1–40.

Mobius, Markus M., and Tanya S. Rosenblat. 2001. “The Process of Ghetto Formation: Evidence from Chicago.” mimeo.

Parikh, Rohit, and Paul Krasucki. 1990. “Communication, consensus, and knowledge.” *Journal of Economic Theory*, 52(1): 178–189.

Wilson, William J., James M. Quane, and Bruce H. Rankin. 1998. “Locked in the Poorhouse: Cities, Race, and Poverty in the United States.” , ed. Fred R. Harris and Lynn Curtis, Chapter 4: The New Urban Poverty: Consequences of the Economic and Social Decline of Inner-City Neighborhoods. Rowman and Littlefield.

A Appendix: An extension and proofs

A.1 A model with agents forming more higher-order estimates

In Section 2, agents form their estimates up to order $\widehat{L}_i$. Then the estimates involving a repeated agent (as the last agent) are set to be equal to the estimates without her. We now show that doing so is without loss. Suppose each agent i forms estimates up to order $\overline{L}$, $\overline{L} > \widehat{L}_i$. We require $\overline{L}$, which can be arbitrarily large, to be finite so that each agent can finish her updating in each period.

In the following procedure, we only describe how agents form their estimates with at least one repeated agent. All their estimates involving only distinct agents in any fully connected subset of g_i remain as before. Let $distinct\{ij\ldots k\} = \{i_0, i_1, \ldots, i_l\}$ be the set of distinct agents in this sequence, such that $i_h \neq i_{h'}$ for any $i_h, i_{h'} \in \{i_0, i_1, \ldots, i_l\}$. By definition, the length of $distinct\{ij\ldots k\}$ is at most $\widehat{L}_i$.

At $t = 0$ and $t = 1$, agent i learns as before. All her higher-order estimates $\mathbf{p}_1^{ij\ldots k} = \{1/N, \ldots, 1/N\}$. For all $t \geq 1$, each agent follows steps similar to those in Section 2 to update their estimates at period $t+1$, with the following additions.

Step 1: Identify new information. For each sequence of fully connected agents $\{ij\ldots k\}$ with $distinct\{ij\ldots k\} \subseteq g_i$, agent i believes that agent j believes...that agent k infers $\boldsymbol{\alpha}_t^{ij\ldots kh} \equiv \{\alpha_t^{ij\ldots kh}(1), \ldots, \alpha_t^{ij\ldots kh}(N)\}$ from agent $h \in g_{ij\ldots k}$ such that

$$\alpha_t^{ij\ldots kh}(n) = \frac{p_t^h(n)}{p_t^{ij\ldots kh}(n)} \Big/ \sum_{n'} \frac{p_t^h(n')}{p_t^{ij\ldots kh}(n')}. \tag{20}$$

Step 2: Update own estimates. This is exactly the same as before.

Step 3: Update estimates of neighbors' estimates. Agent i's higher-order estimates up to order $\overline{L}$ are formed similarly:

$$p_{t+1}^{ij\ldots k}(n) = \frac{p_t^k(n) \prod_{h \in \left(g_{ij\ldots k} \setminus \{k\}\right)} \alpha_t^{ij\ldots kh}(n)}{\sum_{n'=1}^N p_t^k(n') \prod_{h \in \left(g_{ij\ldots k} \setminus \{k\}\right)} \alpha_t^{ij\ldots kh}(n')}. \tag{21}$$

Finally, let $\mathbf{p}_{t+1}^{ij\ldots kh} = \mathbf{p}_{t+1}^{ij\ldots k}$ for all $h \in distinct\{ij\ldots k\}$. The agent's $\overline{L}+1$-order estimates are degenerate as before. Each agent only uses her $\overline{L}+1$-order estimate to infer her $\overline{L}+1$-order new information in period $t+2$. $\|$

We now show that it is without loss of generality to let agents form estimates only for sequences of distinct agents.

Lemma **6.** *For any two sequences of fully connected agents $\{i\ldots j\}$ and $\{k'\ldots k\}$, if $distinct\{i\ldots j\} = distinct\{k'\ldots k\} \subseteq g_i$, then $\mathbf{p}_t^{i\ldots j} = \mathbf{p}_t^{k'\ldots k}$ for all $t \geq 1$.*

Proof of Lemma 6: We prove this lemma by induction on time t. At $t = 1$, by assumption, $\mathbf{p}_1^{i\ldots j} = \mathbf{p}_1^{k'\ldots k} = \{1/N, \ldots, 1/N\}$. Next, suppose this is true at period t. Then at $t+1$, for some agent l connected to all the distinct agents, agent i and k''s higher-order inferred signals from agent l are respectively:

$$\alpha_t^{i\ldots jl}(n) = \frac{p_t^l(n)}{p_t^{i\ldots jl}(n)} \Big/ \sum_{n'} \frac{p_t^l(n')}{p_t^{i\ldots jl}(n')}, \text{ and } \alpha_t^{k'\ldots kl}(n) = \frac{p_t^l(n)}{p_t^{k'\ldots kl}(n)} \Big/ \sum_{n'} \frac{p_t^l(n')}{p_t^{k'\ldots kl}(n')}.$$

These inferred signals are identical because $\{i\ldots jl\}$ and $\{k'\ldots kl\}$ contain the same set of distinct agents, and thus $\mathbf{p}_t^{i\ldots jl} = \mathbf{p}_t^{k'\ldots kl}$ by our induction hypothesis. Similarly, since

$\{i \dots jk\}$ and $\{k' \dots kj\}$ contain the same set of distinct agents, we have

$$p_t^j(n) \frac{p_t^k(n)}{p_t^{i \dots jk}(n)} = p_t^k(n) \frac{p_t^j(n)}{p_t^{k' \dots kj}(n)}.$$

By Step 3 of our learning procedure, the agents' estimates about others become:

$$\begin{aligned} p_{t+1}^{i \dots j}(n) &= \frac{p_t^j(n) \alpha_t^{i \dots jk}(n) \prod_{l \in g_{i \dots j}/\{j,k\}} \alpha_t^{i \dots jl}(n)}{\sum_{n'} p_t^j(n') \alpha_t^{i \dots jk}(n') \prod_{l \in g_{i \dots j}/\{j,k\}} \alpha_t^{i \dots jl}(n')} \\ &= \frac{p_t^k(n) \alpha_t^{k' \dots kj}(n) \prod_{l \in g_{i \dots j}/\{j,k\}} \alpha_t^{k' \dots kl}(n)}{\sum_{n'} p_t^k(n') \alpha_t^{k' \dots kj}(n') \prod_{l \in g_{i \dots j}/\{j,k\}} \alpha_t^{k' \dots kl}(n')} \\ &= p_{t+1}^{k' \dots k}(n). \end{aligned} \tag{22}$$

The second equality holds because the denominators of $\alpha_t^{i \dots jk}(n)$ and $\alpha_t^{k' \dots kj}(n)$ cancel out in the updating formula. Thus the induction hypothesis is true at $t+1$. ∎

Lemma 6 shows that forming any order of estimates higher than $\widehat{L}_i$ does not change agent's learning outcomes. Therefore even though our results apply for any finite-order estimates, we restrict attention to estimates involving distinct agents in fully connected subsets of g_i for each agent i as in the text.

A.2 Proofs

Proof of Lemma 1: A direct corollary of Lemma 6. ∎

Proof of Lemma 2: For part (1), if (g_i, G_i) satisfies LCS, g_{ij} is fully connected for any $j \in \mathrm{N_i}$. If $g_{ij} = \{i, j\}$, then the claim does not apply. Otherwise, for any $k \in g_{ij} \setminus \{i\}$, $g_{ij} = g_{ik} = g_{ijk}$. We now show by induction that $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik} = \mathbf{p}_t^{ijk} = \mathbf{p}_t^{ij \dots k}$ for all $t \geq 1$.

First, by assumption at $t = 1$, $\mathbf{p}_1^{ij} = \mathbf{p}_1^{ik} = \mathbf{p}_1^{ijk} = \mathbf{p}_1^{ij \dots k} = \{1/N, \dots, 1/N\}$. Next, suppose this is true at period t. At period $t+1$, agent i observe the reports from all her neighbors in g_{ij}. Then by the updating rules given in (6) and (7), we can see that the numerator of agent i's estimates $p_{t+1}^{ij}(n)$ is the same as that of her estimates $p_{t+1}^{ik}(n)$: For $l \in g_{ij} \setminus \{j, k\}$,

$$\alpha_t^{ijl}(n) = \frac{p_t^l(n)}{p_t^{ijl}(n)} \Big/ \sum_{n'} \frac{p_t^l(n')}{p_t^{ijl}(n')}, \text{ which is equal to } \alpha_t^{ikl}(n) = \frac{p_t^l(n)}{p_t^{ikl}(n)} \Big/ \sum_{n'} \frac{p_t^l(n')}{p_t^{ikl}(n')},$$

because $\mathbf{p}_t^{ijl} = \mathbf{p}_t^{ikl}$ by the induction hypothesis. Similar to the proof of Lemma 6, since agent i thinks j and k infer the same signals, the counterpart of equation (22) shows that

$\mathbf{p}^{ij}_{t+1} = \mathbf{p}^{ik}_{t+1}$. Moreover, $\mathbf{p}^{ijk}_{t+1}$ and all the higher-order estimates $\mathbf{p}^{ij\dots k}_{t+1}$ are calculated using the same set of reports from g_{ij}, and thus they must all be the same.

(2) From part (1), we can see that if (g_l, G_l) satisfies LCS for every $l \in g_i$, then $g_{ij} = g_{ik}$ and $g_{ji} = g_{jk}$. By definition, $g_{ij} = g_{ji}$. Therefore $\mathbf{p}^{ij}_t = \mathbf{p}^{ik}_t = \mathbf{p}^{jk}_t$. Moreover, all the higher-order estimates are also equal to $\mathbf{p}^{ij}_t$. ∎

Proof of Lemma 3: By definition,

$$\begin{aligned}
\alpha^{ij}_{t+1}(n) &= \frac{p^j_{t+1}(n)}{p^{ij}_{t+1}(n)} \Big/ \sum_{n'} \frac{p^j_{t+1}(n')}{p^{ij}_{t+1}(n')} \\
&= \frac{\frac{p^j_t(n)\prod_{h\in g_j}\alpha^{jh}_t(n)}{p^j_t(n)\prod_{h\in g_{ij}\setminus\{j\}}\alpha^{ijh}_t(n)} \cdot \frac{\sum_{n'} p^j_t(n')\prod_{h\in g_{ij}\setminus\{j\}}\alpha^{ijh}_t(n')}{\sum_{n'} p^j_t(n')\prod_{h\in g_j}\alpha^{jh}_t(n')}}{\sum_{n'}\left(\frac{p^j_t(n')\prod_{h\in g_j}\alpha^{jh}_t(n')}{p^j_t(n')\prod_{h\in g_{ij}\setminus\{j\}}\alpha^{ijh}_t(n')} \cdot \frac{\sum_{n'} p^j_t(n')\prod_{h\in g_{ij}\setminus\{j\}}\alpha^{ijh}_t(n')}{\sum_{n'} p^j_t(n')\prod_{h\in g_j}\alpha^{jh}_t(n')}\right)} \\
&= \prod_{l\in((g_j\setminus g_i)\cup j)} \alpha^{jl}_t(n) \prod_{h\in g_{ij}\setminus\{j\}} \frac{\alpha^{jh}_t(n)}{\alpha^{ijh}_t(n)} \Big/ \left(\sum_{n'} \prod_{l\in((g_j\setminus g_i)\cup j)} \alpha^{jl}_t(n') \prod_{h\in g_{ij}\setminus\{j\}} \frac{\alpha^{jh}_t(n')}{\alpha^{ijh}_t(n')}\right). \qquad (23)
\end{aligned}$$

The third equality holds because the second term of the numerator and that of the denominator cancel out. ∎

Proof of Lemma 4: We prove the lemma by induction. Recall that $x^i_\emptyset$ is agent i's uninformative signal. The initial signals $\{x^{a,i}_0, x^{b,i}_0\}$ are simply $\{x^i_0, x^i_\emptyset\}$. That is, agent i is uninformed in one of $\{X^a_0, X^b_0\}$, and learns x^i_0 in the other. Therefore, $\{\mathbf{p}^{a,i}_1, \mathbf{p}^{b,i}_1\} = \{\mathbf{p}^i_1, (\frac{1}{N}, \dots, \frac{1}{N})\}$. Moreover, $\mathbf{p}^{ij}_1 = \mathbf{p}^{a,ij}_1 = \mathbf{p}^{b,ij}_1 = \{\frac{1}{N}, \dots, \frac{1}{N}\}$, and the same for all the higher-order estimates. Thus expression (11), (12), and (13) all hold at $t = 1$.

Suppose the result holds at time t. We now show it also holds at time $t+1$. In Step 1, recall that the inferred signals under X^a_t and X^b_t are respectively

$$\alpha^{a,ij}_t(n) = \frac{p^{a,j}_t(n)}{p^{a,ij}_t(n)} \Big/ \sum_{n'} \frac{p^{a,j}_t(n')}{p^{a,ij}_t(n')}, \text{ and } \alpha^{b,ij}_t(n) = \frac{p^{b,j}_t(n)}{p^{b,ij}_t(n)} \Big/ \sum_{n'} \frac{p^{b,j}_t(n')}{p^{b,ij}_t(n')}.$$

Further, using (11) and (12), we have:

$$\begin{aligned}
\alpha^{ij}_t(n) &= \frac{p^j_t(n)}{p^{ij}_t(n)} \Big/ \sum_{n'} \frac{p^j_t(n')}{p^{ij}_t(n')} \\
&= \frac{p^{a,j}_t(n)p^{b,j}_t(n)}{\sum_{n'} p^{a,j}_t(n')p^{b,j}_t(n')} \cdot \frac{\sum_{n'} p^{a,ij}_t(n')p^{b,ij}_t(n')}{p^{a,ij}_t(n)p^{b,ij}_t(n)} \Big/ \sum_{n'} \frac{p^j_t(n')}{p^{ij}_t(n')} \\
&= \alpha^{a,ij}_t(n)\alpha^{b,ij}_t(n) \sum_{n'} \frac{p^{a,j}_t(n')}{p^{a,ij}_t(n')} \sum_{n'} \frac{p^{b,j}_t(n')}{p^{b,ij}_t(n')} \frac{\sum_{n'} p^{a,ij}_t(n')p^{b,ij}_t(n')}{\sum_{n'} p^{a,j}_t(n')p^{b,j}_t(n')} \Big/ \sum_{n'} \frac{p^j_t(n')}{p^{ij}_t(n')}. \qquad (24)
\end{aligned}$$

In Step 2, since $\{x_t^{a,i}, x_t^{b,i}\} = \{x_t^i, x_\emptyset^i\}$, then $\{\alpha_t^{a,ii}(n), \alpha_t^{b,ii}(n)\} = \{\alpha_t^{ii}(n), \frac{1}{N}\}$.

$$\begin{aligned}
p_{t+1}^i(n) &= \frac{p_t^i(n)\prod_{h\in g_i}\alpha_t^{ih}(n)}{\sum_{n'=1}^N p_t^i(n')\prod_{h\in g_i}\alpha_t^{ih}(n')} \\
&= \frac{p_t^{a,i}(n)p_t^{b,i}(n)\prod_{h\in g_i}\alpha_t^{a,ih}(n)\alpha_t^{b,ih}(n)}{\sum_{n'=1}^N p_t^{a,i}(n')p_t^{b,i}(n')\prod_{h\in g_i}\alpha_t^{a,ih}(n')\alpha_t^{b,ih}(n')} \\
&= \frac{p_{t+1}^{a,i}(n)p_{t+1}^{b,i}(n)}{\sum_{n'=1}^N p_{t+1}^{a,i}(n')p_{t+1}^{b,i}(n')}.
\end{aligned}$$

The second equality holds by (11) and (24), and the last equality holds because it is the Step 2 of the learning procedure under X_t^a and X_t^b respectively. Thus (11) holds at time $t+1$. Similarly, the higher-order inferred signals follow the same pattern as that in equation (24), and in Step 3,

$$\begin{aligned}
p_{t+1}^{ij}(n) &= \frac{p_t^j(n)\prod_{h\in g_{ij}\setminus\{j\}}\alpha_t^{ijh}(n)}{\sum_{n'=1}^N p_t^j(n')\prod_{h\in g_{ij}\setminus\{j\}}\alpha_t^{ijh}(n')} \\
&= \frac{p_t^{a,j}(n)p_t^{b,j}(n)\prod_{h\in g_{ij}\setminus\{j\}}\alpha_t^{a,ijh}(n)\alpha_t^{b,ijh}(n)}{\sum_{n'=1}^N p_t^{a,j}(n')p_t^{b,j}(n')\prod_{h\in g_{ij}\setminus\{j\}}\alpha_t^{a,ijh}(n')\alpha_t^{b,ijh}(n')} \\
&= \frac{p_{t+1}^{a,ij}(n)p_{t+1}^{b,ij}(n)}{\sum_{n'=1}^N p_{t+1}^{a,ij}(n')p_{t+1}^{b,ij}(n')}.
\end{aligned}$$

Thus (12) also holds at time $t+1$. Lastly, for any sequence of all fully connected agents $\{ij\ldots k\}$,

$$\begin{aligned}
p_{t+1}^{ij\ldots k}(n) &= \frac{p_t^k(n)\prod_{h\in g_{ij\ldots k}\setminus k}\alpha_t^{ij\ldots kh}(n)}{\sum_{n'=1}^N p_t^k(n')\prod_{h\in g_{ij\ldots k}\setminus k}\alpha_t^{ij\ldots kh}(n')} \\
&= \frac{p_t^{a,k}(n)p_t^{b,k}(n)\prod_{h\in g_{ij\ldots k}\setminus k}\alpha_t^{a,ij\ldots kh}(n)\alpha_t^{b,ij\ldots kh}(n)}{\sum_{n'=1}^N p_t^{a,k}(n')p_t^{b,k}(n')\prod_{h\in g_{ij\ldots k}\setminus k}\alpha_t^{a,ij\ldots kh}(n')\alpha_t^{b,ij\ldots kh}(n')} \\
&= \frac{p_{t+1}^{a,ij\ldots k}(n)p_{t+1}^{b,ij\ldots k}(n)}{\sum_{n'=1}^N p_{t+1}^{a,ij\ldots k}(n')p_{t+1}^{b,ij\ldots k}(n')}.
\end{aligned}$$

Thus (13) also holds at time $t+1$. ∎

Proof of Proposition 1: Consider agent i and her local network (g_i, G_i). Let $|\cdot|$ represent the number of states in any subset of S, for instance $|S| = N$ and $|\emptyset| = 0$. At $t = 1$, $p_1^i(n) = 1/|P^i(s_1)|$ if $s_n \in P^i(s_1)$, and 0 otherwise. By definition (14), this is the correct Bayesian posterior for i at $t = 1$.

At $t = 2$, by Step 1 of our learning procedure,

$$\alpha_1^{ij}(n) = \frac{p_1^j(n)}{p_1^{ij}(n)} \Bigg/ \sum_{n'} \frac{p_1^j(n')}{p_1^{ij}(n')}.$$

Clearly, for all $s_n \in P^j(s_1)$, $\alpha_1^{ij}(n) = 1/|P^j(s_1)|$ and 0 otherwise. Thus the inferred signal has the same distribution as j's estimates: $\boldsymbol{\alpha}_1^{ij} = \mathbf{p}_1^j$. Similarly, $\boldsymbol{\alpha}_1^{kj} = \mathbf{p}_1^j$. Let be intersection of the partitional elements containing true state s_1 of all agents in g_i be $P_1^{g_i}(s_1) \equiv \cap\{P^h(s_1)\}_{h\in g_i}$. Then, by Step 2 of our learning procedure, agent i's estimates are

$$p_2^i(n) = \frac{p_1^i(n)\prod_{h\in g_i}\alpha_1^{ih}(n)}{\sum_{n'} p_1^i(n')\prod_{h\in g_i}\alpha_1^{ih}(n')} = \frac{1}{|P_1^{g_i}(s_1)|}$$

if $s_n \in P_1^{g_i}(s_1)$, and $p_2^i(n) = 0$ otherwise. Similarly, let $P_1^{g_{ij}}(s_1) \equiv \cap\{P^h(s_1)\}_{h\in g_{ij}}$, then agent i's second-order estimates are $p_2^{ij}(n) = 1/|P_1^{g_{ij}}(s_1)|$, for $s_n \in P_1^{g_{ij}}(s_1)$, and $p_2^{ij}(n) = 0$ otherwise. Clearly, $P_1^{g_i}(s_1) \subseteq P_1^{g_{ij}}(s_1)$ since $g_{ij} \subseteq g_i$. And so on for all higher-order estimates.

At $t = 3$, if $\mathbf{p}_2^j \neq \mathbf{p}_2^{ij}$, there must be some states in $P_1^{g_{ij}}(s_1)$ that have zero probability under $\mathbf{p}_2^j$. As at $t = 2$, the inferred signal has the same distribution as j's estimates, $\boldsymbol{\alpha}_2^{ij} = \mathbf{p}_2^j$. Let $P_2^{g_i}(s_1)$ be the set of states agent i thinks are still possible given her inferred signals, then $P_2^{g_i}(s_1) \subset P_1^{g_i}(s_1)$. It is important to notice that, because $P_1^{g_j}(s_1) = \cap\{P^h(s_1)\}_{h\in g_j}$,

$$P_2^{g_i}(s_1) \equiv \cap\{P_1^{g_h}(s_1)\}_{h\in g_i} = \cap\{P^l(s_1)\}_{l\in g_i^2},$$

where $g_i^d \equiv \{l \in g : d(il) \leq d\}$. That is, $P_2^{g_i}(s_1)$ is the intersection of the partitional elements containing s_1 of all the $d \leq 2$ neighbors of agent i. Therefore $p_3^i(n) = 1/|P_2^{g_i}(s_1)|$ if $s_n \in P_2^{g_i}(s_1)$, and 0 otherwise.

Iteratively, we can show that $p_t^i(n) = 1/|P_{t-1}^{g_i}(s_1)|$ if $s_n \in P_{t-1}^{g_i}(s_1) \equiv \cap\{P^h(s_1)\}_{h\in g_i^{t-1}}$, and 0 otherwise. By at most $t = D + 1$, all the initial signals have reached agent i through her neighbors according to the travel path of the signals. Let $P^g(s_1) \equiv \cap\{P^l(s_1)\}_{l\in g}$ be the intersection of all agents' partitional elements containing state s_1. Then agent i's estimates are simply $1/|P^g(s_1)|$ if $s_n \in P^g(s_1)$, and 0 otherwise. This is the case for all the agents in the network, therefore their learning outcomes are strongly Bayesian. ∎

Proof of Lemma 5: For necessity, if a network is a social quilt, it does not contain a simple circle by definition. Moreover, it satisfies GCS because for any agent i and 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. By the definition of social quilts, $kk' \in G$. Thus LCS holds for any agent i, and the network satisfies GCS.

For sufficiency, we prove the claim by induction that any circle of more than three agents that satisfies GCS and does not contain a simple circle must be a clique. First, any four agent circle must be part of a clique. No simple circle means that there must be at least one link between two nonadjacent agents. Then by GCS, these four agents must be connected in a clique. Next, suppose any circle of l (or fewer) agents, $l \geq 5$ 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 fewer 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, and by GCS, they are connected. Therefore this circle of $l+1$ agent must be a clique, which is the definition of a social quilt. ∎

Proof of Proposition 2: We begin with two properties of social quilts. First, if $d(ij) = d$, then there must be a unique path of length d from j to i. Suppose instead, there are two distinct paths with length d between i and j. 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 two numbers k and h, $0 \leq k < h \leq d$ and $h - k \geq 2$ such that:

$$\begin{cases} i_k = j_k; & \\ i_l \neq j_l, & \text{if } k < l < h; \\ i_h = j_h. & \end{cases}$$

Clearly, $\{i_k i_{k+1} \dots i_h j_{h-1} \dots j_{k+1}\}$ must be a circle, going from i_k to herself through distinct agents. They are distinct because $i_l \neq j_l$ 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 connected. Thus agent i_k and i_h must be connected and $i_k i_h$ is a unique shortest path between them, which is a contradiction to $\{ii_1i_2 \dots i_{d-1}j\}$ being a shortest path.

The second property of social quilts is that if agent i's signal travels from agent l to k, and then inferred by k's neighbor j who is not connected to l, then j must be further away from i, Specifically, if l is the agent next to k on the shortest path from i to k, such that $d(ik) = d(il) + 1$ and $kl \in G$, then for any j with $jk \in G$ and $jl \notin G$, the shortest path from i to j must go through l and k: $d(ij) = d(ik) + 1$. To see this, note that since $jk \in G$, the maximum possible distance between i and j is $d(ij) = d(ik) + 1$. Next, if $d(ij) \leq d(ik) - 1$, then the path through l cannot be the unique shortest path between i and k. If $d(ij) = d(ik)$, then the shortest path between i and j must not involve k, or agent l since $jl \notin G$. Thus we have a circle involving $\{j, k, l\}$ and i's shortest path to agent j and l, which is a contradiction to the definition of social quilts. Therefore $d(ij) = d(ik) + 1$. This implies that once a signal reaches an agent, it cannot travel back and reach her again.

We now proceed to prove the claim. By Proposition 4, if we can show that agents' learning outcomes are strongly Bayesian for each signal, then it is also true for multiple signals. Without loss of generality, let agent i receive an initial signal $x_0^i = x_m^i$. Let $\mu_{nm}^i = \Pr(s_n | x_m^i)$ for agent i. We want to show that each agent j infers this signal at $t = d(ij) + 1$ from some neighbor k (who can be agent i), and this is the only signal j infers from her neighbors at any time. Specifically, $\alpha_t^{jk}(n) = \mu_{nm}^i$ if and only if $t = d(ik) + 1 = d(ij)$, and $\alpha_t^{jk}(n) = 1/N$ otherwise (as well as all the higher-order estimates). Notice that this implies agent j learns the signal and changes his estimates once at $t = d(ij) + 1$.

We prove this by induction on time t. First, this holds at $t = 1$. If $d(ij) = 1$, or $j \in \mathrm{N}_i$, then agent j infers the signal from agent i's report $\mathbf{p}_1^i$ such that $\alpha_1^{ji}(n) = \mu_{nm}^i$. No other agents (including agent i) infer any new signal from their neighbors. Next, if $\alpha_1^{jk}(n) = \mu_{nm}^i$, then by Lemma 3, $\alpha_0^{kh}(n) = \mu_{nm}^i$ for some h not connected to j. Clearly, $h = k = i$ and $d(ik) = 0$, $d(ij) = 1$.

Suppose this holds at period t, we want to show it also holds at $t + 1$. If $\alpha_t^{jk}(n) = \mu_{mn}^i$, using expression (3) in Lemma 3, we have

$$\alpha_t^{jk}(n) = \frac{\prod_{h \in ((g_k \setminus g_j) \cup k)} \alpha_{t-1}^{kh}(n)}{\sum_{n'} \prod_{h \in ((g_k \setminus g_j) \cup k)} \alpha_{t-1}^{kh}(n')}.$$

That is, agent k must infer the signal from someone (say l) outside g_j in the previous period, so $jl \notin G$. By induction, since $\alpha_{t-1}^{kl}(n) = \mu_{nm}^i$, we have $d(ik) = t-1$ and $d(il) = t-2$. By the second property above, it must be true that $d(ij) = t$. On the other hand, if $d(ik) = t-1$, by induction $\alpha_{t-1}^{kl}(n) = \mu_{nm}^i$ for some neighbor l. As $d(il) = t-2$ and $d(ij) = t$, it is true that $l \in g_k \setminus g_j$. Because agent j has not learned any new information so far, $\alpha_t^{jk}(n) = \mu_{mn}^i$. Thus $\alpha_t^{jk}(n) = \mu_{nm}^i$ if and only if $d(ij) = t$ and $d(ik) = t - 1$.

Since signal x_0^i arrives at each agent $j \in g$ exactly once at period $d(ij) + 1$, $p_t^j(n) = \mu_{nm}^i$ if $t > d(ij)$ and $p_t^j(n) = 1/N$ otherwise. Everyone learns x_0^i at period $D + 1$ since D is the diameter of the network. Thus the learning is strongly Bayesian with signal x_0^i. When there are multiple signals, Lemma 4 ensures that the learning remains strongly Bayesian. ∎

Proof of Proposition 3: Suppose that agent i receives the only informative signal $x_0^i = x_m^i$. We first describe the repetition of this signal within a simple circle. For any k-agent simple circle $c = \{i_1 \dots i_k\}$, there are two separate cases: agent $i \in c$ or $i \notin c$.

First, suppose $i \in c$. Without loss, assume $i = i_k$. Then at $t = 2$, agent i_1 and i_{k-1}'s inferred signals are $\alpha_1^{i_1 i_k}(n) = \alpha_1^{i_{k-1} i_k}(n) = \mu_{nm}^i$. Let $\mu_{nm}^i(\eta)$ be the Bayesian posterior of the true state being s_n if an agent infers η copies of identical x_0^i: $\mu_{nm}^i(\eta) = \frac{(\mu_{nm}^i)^\eta}{\sum_{n'} (\mu_{n'm}^i)^\eta}$. By Lemma

3, when GCS holds, all the inferred signals must have the distribution equal to $\mu^i_{nm}(\eta)$ for some non-negative η. Let $\alpha^{jk}_t(n) = \mu^i_{nm}(\eta^{jk}_t)$. Lemma 3 can be rewritten as

$$\eta^{jk}_{t+1} = \sum_{l \in g_k \setminus g_j} \eta^{kl}_t. \tag{25}$$

At period $t = k+1$, the signal finishes traveling around the simple circle in both directions, and thus $\eta^{ii_{k-1}}_k \geq 1$ and $\eta^{ii_1}_k \geq 1$. Both may be higher than 1 because there may be more than one simple circle. Thus agent i and all other agents in the simple circle infer at least two copies of x^i_0 in every k periods.

Next, if $i \notin c$, then without loss of generality, assume i_k is the first one of the simple circle (or one of the first ones) to learn the signal, such that $\eta^{i_k j}_t \geq 1$ for some $j \in \mathrm{N}_{i_0}$. Because i_1 and i_{k-1} are not connected by definition of a simple circle, j cannot be connected with i_1 and i_{k-1} at the same time because the network is assumed to satisfy GCS. Suppose $ji_1 \notin G$, then $\eta^{i_1 i_0}_{t+1} \geq 1$, and it is passed on to i_2, i_3 and so on. Also, the signal travels through i_{k-1} to i_{k-2}, because i_{k-1} learns from either j or i_k. Similar to the first case, we can show agent i_k and all other agents in the simple circle infer at least two copies of x^i_0 in every k periods.

We now turn to how a signal travels in a network that satisfies GCS but contains multiple simple circles. Recall that the largest simple circle in the network (g, G) has k agents. If k is even, it takes $k/2$ periods for a signal to reach every agent in the simple circle and another $k/2$ periods for everyone to infer the second copy of the signal. If k is odd, it takes $(k-1)/2$ periods for everyone to learn the first copy and $(k+1)/2$ periods for the second. For simplicity, we consider the case when k is even.

Since x^i_0 is the only informative signal, at $t = 1 + D$, everyone in the network learns at least one copy of x^i_0, which travels along each simple circle in both directions. From now on, we focus on agents belonging to at least one simple circle. At period $t = 1 + D + k/2$, everyone in a simple circle infers at least two copies of x^i_0 and at least one agent in each simple circle infers at least three copies. Treat this two "new" copies of the signal as new information, which travels to all the other agents in D periods. Thus at $t = 1 + 2D + k/2$, each agent learns at least $3 + 2(k_{sc} - 1)$ copies. The first part comes from each agent infers at least one new copy from her own simple circle. The second part comes from learning two copies of the signal from each of the remaining $(k_{sc} - 1)$ simple circles. Similarly, at $t = 1 + 2(D + k/2)$, every agent in the network learns at least $4 + 4(k_{sc} - 1)$ copies and generates at least $4(k_{sc} - 1)$ "new" copies of the signal. This is clearly true if there is only one agent (say agent h) connecting to the rest of the network from one simple circle. Because by period $t = 1 + 2D + k/2$, the $2(k_{sc} - 1)$ copies have already traveled half of the simple circle, and they come back to agent h in another $k/2$ periods in both directions. If there are two or more agents connecting to the outside, then each sends at least $2(k_{sc}-1)$ copies to the outside as these copies travel around the simple circle in $k/2$ periods, so the total number of new signals to the outside is at least $4(k_{sc}-1)$. Then in another $D+k/2$ periods, each agent infers at least two more copies of the signal because the initial information finishes another trip in her own simple circle. In addition, she learns at least $2(k_{sc}-1)\cdot 4(k_{sc}-1)$ copies from other simple circles.

Counting the number of copies iteratively, for any $t \in \left[\tau(D+\lceil k/2\rceil)+1, (\tau+1)(D+\lceil k/2\rceil)\right]$, every agent believes there are at least $2\tau + 2\sum_{\tau'=1}^{\tau-1}(2(k_{sc}-1))^{\tau'}$ copies of signal x_0^i if τ is an integer larger than 1. ∎

Proof of Corollary 1: Suppose the only simple circle is $c=\{i_1\dots i_k\}$ with distinct agents. By assumption, the network outside of c must satisfy GCS and contain no simple circles. First, consider the case with only one informative signal x_0^i. Then the first time this signal arrives at the circle, it must reach either only one agent (say i_k), or two connected 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 from i_k. 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. Further, once the signal reaches the circle, agents in the simple circle do not learn any other new information from outside, because there is no other simple circle through which information can travel back to c.

Now consider the case with multiple informative signals. For ease of notation, suppose each agent receives at most one signal (the other case is similar). Let signal x_t^l reach the circle at agent $i_{h(l)}\in c$ ($h(l)\in\{1,\dots,k\}$) and at time $\tau(l)$ for the first time. Then for agent i_k, if $h(l)=k$, i_k receives two more copies at each time $\tau(l)+ko$ for any integer o, otherwise she receives one more copy at each time $\tau(l)+h(l)+ok$ and $\tau(l)+(o+1)k-h(l)$. At time $t=T+D$, all signals must have reached the circle,

$$p_{T+D}^{i_k}(n)=\frac{\prod_{l\in g}(\alpha_t^{ll}(n))^{\eta_{T+D}^{i_k}(x_t^l)}}{\sum_{n'}\prod_{l\in g}(\alpha_t^{ll}(n'))^{\eta_{T+D}^{i_k}(x_t^l)}}, \tag{26}$$

where $\eta_{T+D}^{i_k}(x_t^l)$ is the number of copies of x_t^l agent i_k learned at time $T+D$. As we argued above, in every k periods, agent i_k must receive two more copies of each signal, such that

$$p_{T+D+ok}^{i_k}(n)=\frac{\prod_{l\in g}(\alpha_t^{ll}(n))^{\eta_{T+D}^{i_k}(x_t^l)+2o}}{\sum_{n'}\prod_{l\in g}(\alpha_t^{ll}(n'))^{\eta_{T+D}^{i_k}(x_t^l)+2o}}. \tag{27}$$

By assumption, there is a unique state $\tilde{s}$ such that $\Pr\left(\tilde{s}|\{\alpha_t^{ll}\}_{l\in g,0\leq t<T}\right) > \Pr\left(s'|\{\alpha_t^{ll}\}_{l\in g,0\leq t<T}\right)$ for all $s' \neq \tilde{s}$. As $o \to \infty$, $p_{T+D+ok}^{i_k}(\tilde{s}) \to 1$. For any other large time t, which is between $T + D + ok$ and $T + D + (o+1)k$, let $\Delta t = t - (T + D + ok)$, so $\Delta t \in \{1, \ldots, k-1\}$,

$$p_t^{i_k}(n) = \frac{\prod_{l\in g}(\alpha_t^{ll}(n))^{\eta_{T+D+\Delta t}^{i_k}(x_t^l)+2o}}{\sum_{n'}\prod_{l\in g}(\alpha_t^{ll}(n'))^{\eta_{T+D+\Delta t}^{i_k}(x_t^l)+2o}}. \tag{28}$$

Similarly, as $t \to \infty$, $o \to \infty$, and so $p_t^{i_k}(\tilde{s}) \to 1$. The same argument applies to all agents in the circle, and information in the circle is learned by all others in the whole network. ∎

Proof of Proposition 4: Suppose $x_0^l = x_m^l$ is the only signal. We can classify all agents based on their distance to l, that is, $\mathrm{N}_l^d = \{h \in g : d(lh) = d\}$, and $\mathrm{N}_l^1 = \mathrm{N}_l$. To begin with, if agent a and $b \in \mathrm{N}_l^d$ are both connected to some agent h in N_l^{d+1}, then $ab \in G$. Find a's connection to some agent f in N_l^{d-1}, then agent f and h must not be connected, because their distance must be 2. Similarly the agent who is connected to b in N_l^{d-1}, say f', cannot be connected to h. If agent a and b are not connected, then there exists a simple circle consisting of agent f, a, h and b (with possibly other agents like f' and l), which is a contradiction.

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

The argument above shows that agent l never learns any new information and thus her estimates remain at μ^l_{nm}. Moreover, estimates of agents in N_l must remain at μ^l_{nm}. This is because first, they cannot infer new information from their successors. Second, for any connected agents in N_l, they learn from agent l simultaneously and expect each other to learn the signal. Therefore they cannot infer new information from each other.

Lastly, we claim that there must exist some agent $l' \in \mathrm{N}^2_l$, who is connected to at least two agents in N_l but does not infer new signals from his peers (those with the same distance as him). Therefore the estimates of agent l' oscillate and never agree with the Bayesian posterior. Recall that $l \in \tilde{g}$ implies that there exist $i, j \in \mathrm{N}_l$ and $k \in \mathrm{N}^2_l$ such that $k \in g_{ij}$. We start with this agent k who is connected to i and j, and possibly more agents in N_l. If k does not infer new signals from his peers in N^2_l, then we can show k must keep oscillating. This is because by the claim above agents in N_l who are connected to k must be connected with each other. So k must keep inferring multiple copies of x^l_m in even periods, and multiple copies of the signal that offsets x^l_m is odd periods for $t \geq 3$.

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

In the second case, agent k learns new information indirectly from some peer $h' \in \mathrm{N}^2_l$. 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 connected to agent h' while kj' and hj' cannot be connected. Moreover, ih must also be connected here to avoid a simple circle, so in this case $ijhk$ is a clique. In fact, $ijhh'$ is also a clique. Therefore h' is connected 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 $ijkh$ is a clique, then $ijhh'$ has to be a clique, $ijh'h^{''}$ 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 connected to multiple agents (more than i, j) in N_l. ∎

Proof of Corollary 2: If the network is not a social quilt, it must contain either at least one simple circle, or some agent i such that LCS does not hold, or both. We prove the result by showing that in either case, there exists some sequence of signals such that an agent's learning outcome is not strongly Bayesian.

First, suppose there is a simple circle $\{i_1 \dots i_k\}$ and GCS holds. We focus on the case when k is even, and the other case is analogous. Let $j = i_{k/2}$ be the agent opposite of $i = i_k$ in the circle. Then at $t = j + 1$, agent j's estimates are

$$p_t^j(n) = \Pr\left(s_n \middle| \eta_t^j(x_\tau^h) x_\tau^h, \forall h \in g, \tau \in \{0, \dots, t - d(jh) - 1\}\right).$$

Here $\eta_t^j(x_\tau^h)$ is the number of copies of signal x_τ^h that agent j has learned by period t. Clearly, $\eta_t^j(x_0^i) \geq 2$ due to the simple circle, and $\eta_t^j(x_\tau^h) \geq 1$ for $\forall h \in g, \tau \in \{0, \dots, t - d(jh) - 1\}$. If agent j's learning outcome is strongly Bayesian, his estimates should equal to the Bayesian posterior belief conditional on exactly one copy of each signal that has arrived. This implies that for all $\eta_t^j(x_\tau^h) > 1$, the extra copies of signals must cancel each other in terms of the Bayesian posterior belief. That is,

$$\Pr\left(s_n \middle| (\eta_t^j(x_\tau^h) - 1) x_\tau^h, \forall h \in g, \tau \in \{0, ..., t - d(jh) - 1\}\right) = 1/N.$$

Since $\eta_t^j(x_0^i) - 1 \geq 1$, this cannot hold if x_0^i is the only informative signal.

Second, suppose there is some agent i such that LCS does not hold. Then there must exist $j \in \mathrm{N}_i$ and $k, k' \in \mathrm{N}_i \cap \mathrm{N}_j$, such that $kk' \notin G$. Without loss of generality, assume $x_0^{k'}$ is the only informative signal. Then at $t = 3$, agent k's estimates much include at least two copies of $x_0^{k'}$, which is not strongly Bayesian. ∎

Proof of Proposition 5: We show that if all agents hold this belief, then $\mathbf{p}_t^i$, $\mathbf{p}_t^{ij}$, $\mathbf{p}_t^{ij\dots k}$ as defined in Section 2 and Section A.1 are the correct Bayesian posterior beliefs from every agent i's perspective, and thus we call them beliefs instead of estimates in this proof. More precisely, let $B^i(z)$ denote what agent i believes z is with probability 1, which may or may not be equal to z. Then we have, for all $t \geq 2$, all s_n and all $i \in g$:

$$\text{(i)} \quad p_t^i(n) = B^i\left(\Pr\left(s_n | \{x_\tau^h\}_{\tau \leq t-2,\ h \in g}, x_{t-1}^i\right)\right); \text{ and}$$

$$\text{(ii)} \quad p_t^{ij}(n) = B^i\left(\frac{p_t^j(n)/\alpha_{t-1}^{jj}(n)}{\sum_{n'} p_t^j(n')/\alpha_{t-1}^{jj}(n')}\right); \text{ and}$$

$$\text{(iii)} \quad \mathbf{p}_t^{ij\dots k} = B^i\left(\mathbf{p}_t^{j\dots k}\right).$$

That is, agent i believes that (i) her belief at period t is based on all the available signals from the entire network up to period $t-2$ plus her own signal at $t-1$; and (ii) her belief of j's belief includes all the information j has inferred except for his most recent signal x_{t-1}^j; and (iii) her third (and higher)-order beliefs are all correct.

We prove this result by induction. First, agent i gets x_0^i and her Bayesian posterior belief is $p_1^i(n) = \Pr(s_n|x_0^i)$ by definition. All the higher-order beliefs $\mathbf{p}_1^{ij\ldots k} = \{1/N, \ldots, 1/N\}$, which are their prior, are also correct because they have not learned anything from their neighbors. At $t = 2$, the inferred signals $\boldsymbol{\alpha}_2^{ij}$ for all $j \in \mathrm{N_i}$ are the same as x_0^j. Note that agent i believes $g_j \subseteq g_i = g$, and thus $g_{ij} = B^i(g_j)$. Similarly, agent i believes $g_{j\ldots k} \subseteq g_i$, and thus $g_{ij\ldots k} = B^i(g_{j\ldots k})$. Using expression (5), we have

$$p_2^i(n) = \Pr\left(s_n|\{x_0^h\}_{h\in g_i}, x_1^i\right) = B^i\left(\Pr\left(s_n|\{x_0^h\}_{h\in g}, x_1^i\right)\right).$$

Using expression (6), we have:

$$p_2^{ij}(n) = \Pr\left(s_n|\{x_0^h\}_{h\in g_{ij}}\right) = B^i\left(\Pr\left(s_n|\{x_0^h\}_{h\in g_j}\right)\right),$$

and thus any difference from $p_t^j(n)$ can only be attributed to x_1^j. Moreover,

$$p_2^{ij\ldots k}(n) = B^i\left(\Pr\left(s_n|\{x_0^h\}_{h\in g_{j\ldots k}}\right)\right) = B^i(p_2^{j\ldots k}(n)).$$

Next, suppose the induction hypothesis is true at t. Then from agent i's point of view, if $\mathbf{p}_t^j \neq \mathbf{p}_t^{ij}$, according to (ii) above, the only difference is caused by j's private signal x_{t-1}^j. That is, the inferred signal y_{t-1}^{ij} is the same as x_{t-1}^j. Recall from the discussion in Section 2.2, the inferred signal $\boldsymbol{\alpha}_t^{ij}$ contains all the information in y_{t-1}^{ij} (and from i's perspective, x_{t-1}^j) she needs to update her beliefs. Clearly, $\boldsymbol{\alpha}_t^{ij}$ must be new information to i, and thus agent i will follow Step 2 to update her own beliefs. Therefore (i) holds: $p_{t+1}^i(n) = B^i\left(\Pr(s_n|\{x_\tau^h\}_{\tau\leq t-1, h\in g}, x_t^i)\right)$, which is the correct posterior belief from agent i's perspective. Moreover, since agent i believes she observes all j's neighbors, she must update her beliefs $\mathbf{p}_{t+1}^{ij}$ according to Step 3 of our procedure. By the induction hypothesis, $\boldsymbol{\alpha}_t^{ijk} = B^i\left(\boldsymbol{\alpha}_t^{jk}\right)$, and thus i thinks she can infer all the signals j inferred from his neighbors. The only thing missing is j's own signal, which is exactly point (ii). To see this, note that

$$p_{t+1}^{ij}(n) = \frac{p_t^j(n)\prod_{h\in g_{ij}\setminus\{j\}}\alpha_t^{ijh}(n)}{\sum_{n'=1}^N p_t^j(n')\prod_{h\in g_{ij}\setminus\{j\}}\alpha_t^{ijh}(n')} = B^i\left(\frac{p_t^j(n)\prod_{h\in g_j\setminus\{j\}}\alpha_t^{jh}(n)}{\sum_{n'=1}^N p_t^j(n')\prod_{h\in g_j\setminus\{j\}}\alpha_t^{jh}(n')}\right)$$

while

$$B^i\left(p^j_{t+1}(n)\right) = B^i\left(\frac{p^j_t(n)\prod_{h\in g_j}\alpha^{jh}_t(n)}{\sum_{n'=1}^N p^j_t(n')\prod_{h\in g_j}\alpha^{jh}_t(n')}\right).$$

Similarly, because $\boldsymbol{\alpha}^{ij\ldots kl}_t = B^i\left(\boldsymbol{\alpha}^{j\ldots kl}_t\right)$, $\mathbf{p}^{ij\ldots k}_{t+1} = B^i\left(\mathbf{p}^{j\ldots k}_{t+1}\right)$ and (iii) holds. ▮

Proof of Proposition 6: The sufficiency closely follows the proof of Proposition 2. In particular, if the network is a tree and x^i_0 is the only signal, then agent $h \in g\setminus\{i\}$ infers the signal at and only at period $t = d(ih)+1$. So the learning is strongly Bayesian.

The necessity is due to the fact that if the network is not a tree, it must contain a triangle or a simple circle. In both cases, agents infer positively correlated signals. Because every agent's local 0-network satisfies LCS, any network satisfies GCS, and thus agents do not infer negatively correlated signals. Therefore even with one informative signal, the agents infer infinite copies of the signal as $t\to\infty$, and their estimates are not Bayesian. ▮

Proof of Proposition 7: First, if the diameter $D \le d$, $g^d_i = g$ and $G^d_i = \{jk : jk \in G\}\cup\{\widehat{lh} : lh \notin G\}$. From each agent's perspective, her local d-network is a clique. Their learning outcomes are d-strongly Bayesian because once a signal reaches any agent, all other agents in the network infer it two periods later, and they all know they have learned it from the same agent. Thus each signal is learned once and only once. The learning stops at $T+1$, and the agents' estimates are the Bayesian posterior given the signals in X_T.

Second, if $D > d$, there exists a pair of agents l and l' such that $d(ll') = d+1$. Let $\{ll_1\ldots l_d l'\}$ be one shortest path between agent l and l'. While agent l and l' do not observe each other's report, they both can observe agents $l_1,\ldots,l_d$. Suppose the only signal is x^l_0. At $t=1$, $p^l_1(n) = \Pr(s_n|x^l_0)$ for all s_n. At $t=2$, the signal is inferred by agents $l_1,\ldots,l_d$, and agent l_k reports $p^{l_k}_2(n) = \Pr(s_n|x^l_0)$ for all $k\in\{1,\ldots,d\}$. At $t=3$, agent l' infers one copy of the signal from each agent l_k, so his estimates must include at least d copies of x^l_0, possibly more if there are multiple paths between l and l'. Clearly, agent l''s estimates are not d-strongly Bayesian. ▮

Proof of Proposition 8: By assumption, no new signals arrive at or after time T. Without loss, suppose that $\tau \le T$. The other case can be proved by replacing T with τ from the following expressions. For all $s_n, s_{n'}$ and any sequence of fully connected agents $\{ik\ldots k'\}$, let

$$\overline{\alpha}_T = \max_{n,n'}\frac{\alpha^{ik\ldots k'}_T(n)}{\alpha^{ik\ldots k'}_T(n')}$$

be the highest ratio among all the elements of the distributions of inferred signals at time $T+1$, including all the higher-order inferred signals ($\overline{\alpha}_t$ for any $t\ge 1$ is defined similarly). Because $\phi^i_{mn}\in(0,1)$ for all $i\in g$, all the agents' estimates and higher-order estimates

$p^i_T, p^{ij}_T, \ldots$ are bounded away from 0 and 1. Since the inferred signals are the ratio of the agents' (higher-order) estimates, and there are a finite number of higher-order estimates, the maximum $\overline{\alpha}_T$ exists, and $\overline{\alpha}_T \in [1, \infty)$.

Let $V \equiv \max_{i \in g}(\overline{w}^i L_i)$. Clearly, $V < \frac{1}{2}$. By the definition of higher-order estimates, we can rewrite the ratio of any higher-order inferred signals at $T+2$ as:

$$\begin{aligned}
\frac{\alpha^{ik\ldots k'}_{T+1}(n)}{\alpha^{ik\ldots k'}_{T+1}(n')} &= \frac{p^{k'}_{T+1}(n)}{p^{ik\ldots k'}_{T+1}(n)} \cdot \frac{p^{ik\ldots k'}_{T+1}(n')}{p^{k'}_{T+1}(n')} \\
&= \frac{p^{k'}_T(n) \prod_{h \in g_{k'} \setminus \{k'\}} (\alpha^{k'h}_T(n))^{w^{k'h}_T}}{p^{k'}_T(n) \prod_{l \in g_{ik\ldots k'} \setminus \{k'\}} (\alpha^{ik\ldots k'l}_T(n))^{w^{k'l}_T}} \cdot \frac{p^{k'}_T(n') \prod_{l \in g_{ik\ldots k'} \setminus \{k'\}} (\alpha^{ik\ldots k'l}_T(n'))^{w^{k'l}_T}}{p^{k'}_T(n') \prod_{h \in g_{k'} \setminus \{k'\}} (\alpha^{k'h}_T(n'))^{w^{k'h}_T}} \\
&= \prod_{h \in g_{k'} \setminus \{k'\}} \left(\frac{\alpha^{k'h}_T(n)}{\alpha^{k'h}_T(n')} \right)^{w^{k'h}_T} \cdot \prod_{l \in g_{ik\ldots k'} \setminus \{k'\}} \left(\frac{\alpha^{ik\ldots k'l}_T(n')}{\alpha^{ik\ldots k'l}_T(n)} \right)^{w^{k'l}_T} \\
&\leq (\overline{\alpha}_T)^{2V}. \qquad (29)
\end{aligned}$$

Therefore the maximum of the above ratios $\overline{\alpha}_{T+1} \leq (\overline{\alpha}_T)^{2V}$. Similarly, we can show that $\overline{\alpha}_{t+1} \leq (\overline{\alpha}_t)^{2V}$ for all $t > T$, and thus $\overline{\alpha}_t \leq (\overline{\alpha}_T)^{(2V)^{t-T}}$.

Denote the aggregate new information i learns from all her neighbors at time $t+1$ as,

$$\alpha^i_t(n) = \prod_{j \in g_i \setminus \{i\}} \frac{(\alpha^{ij}_t(n))^{w^{ij}_t}}{\sum_{n'} (\alpha^{ij}_t(n'))^{w^{ij}_t}}.$$

Then the ratio

$$\frac{\alpha^i_t(n)}{\alpha^i_t(n')} = \prod_{j \in g_i \setminus \{i\}} \left(\frac{\alpha^{ij}_t(n)}{\alpha^{ij}_t(n')} \right)^{w^{ij}_t} \leq (\overline{\alpha}_t)^V.$$

This is because each new signal agent i learns from her neighbor is weighted by at most $\overline{w}^i$ and agent i has L_i neighbors. Using the fact that $\overline{\alpha}_t \leq (\overline{\alpha}_T)^{(2V)^{t-T}}$, we can show when $t = T + z$, $\alpha^i_t(n)/\alpha^i_t(n') \leq (\overline{\alpha}_T)^{2^z V^{z+1}}$.

We can then express agent i's estimates at time $t+h$ as a function of her estimates and inferred signals from period t onward:

$$\begin{aligned}
p^i_{t+h}(n) &= \frac{p^i_{t+h-1}(n)\alpha^i_{t+h-1}(n)}{\sum_{n'} p^i_{t+h-1}(n')\alpha^i_{t+h-1}(n')} \\
&= \frac{p^i_{t+h-2}(n)\alpha^i_{t+h-2}(n)\alpha^i_{t+h-1}(n)}{\sum_{n'} p^i_{t+h-2}(n')\alpha^i_{t+h-2}(n')\alpha^i_{t+h-1}(n')} \\
&= \frac{p^i_t(n)\alpha^i_t(n)\ldots\alpha^i_{t+h-2}(n)\alpha^i_{t+h-1}(n)}{\sum_{n'} p^i_t(n')\alpha^i_t(n')\ldots\alpha^i_{t+h-2}(n')\alpha^i_{t+h-1}(n')} \\
&= \frac{p^i_t(n)}{\sum_{n'} p^i_t(n')\frac{\alpha^i_t(n')}{\alpha^i_t(n)}\ldots\frac{\alpha^i_{t+h-2}(n')}{\alpha^i_{t+h-2}(n)}\frac{\alpha^i_{t+h-1}(n')}{\alpha^i_{t+h-1}(n')}}. \qquad (30)
\end{aligned}$$

Using the results derived above, we can see that for any $\varepsilon > 0$, there exists some z such that when $t > T + z$ and $2V < 1$,

$$\begin{aligned}
p^i_{t+h}(n) &\geq p^i_t(n)\cdot\left((\overline{\alpha}_T)^{-2^z V^{z+1}}\right)^{\frac{1-(2V)^h}{1-2V}} \geq p^i_t(n)(1-\varepsilon), \text{ and} \\
p^i_{t+h}(n) &\leq p^i_t(n)\cdot\left((\overline{\alpha}_T)^{2^z V^{z+1}}\right)^{\frac{1-(2V)^h}{1-2V}} \leq p^i_t(n)(1+\varepsilon),
\end{aligned}$$

for any $h > 0$ and state s_n. This shows that if $V < \frac{1}{2}$, then $\left|p^i_{t+h}(n) - p^i_t(n)\right| < \varepsilon$ for t sufficiently large. Therefore the agents' estimates converge as $t \to \infty$.

Finally, fix t and z and let h go to infinity, we have

$$p^i_t(n)\cdot\left((\overline{\alpha}_T)^{-2^z V^{z+1}}\right)^{\frac{1}{1-2V}} \leq p^i_{t\to\infty}(n) \leq p^i_t(n)\cdot\left((\overline{\alpha}_T)^{2^z V^{z+1}}\right)^{\frac{1}{1-2V}}.$$

Therefore the limit estimates are strictly between 0 and 1. ∎

Proof of Proposition 9: Let agent i be the only agent who receives informative signals from nature. We claim that for any agent $h \in g$, at period t, $t \geq d(ih)+1$, she infers signal $x^i_{t-d(ih)-1}$ for the first time from all her neighbors in $\mathrm{N}^{d(ih)-1}_i$ and puts it in her set A^t_h. At period t, her estimates become $p^h_t(n) = \Pr\left(s_n \middle| x^i_0, \ldots, x^i_{t-d(ih)-1}\right)$, which are strongly Bayesian. Moreover, agent h's estimates of any neighbor h' who either is in $\mathrm{N}^{d(ih)-1}_i$ or shares a common neighbor with h in $\mathrm{N}^{d(ih)-1}_i$ are the same as $\mathbf{p}^h_t$. Her estimates of any other neighbor h'' must be $\mathbf{p}^{hh''}_t = \mathbf{p}^h_{t-1}$ because h'' should learn from h, and similarly for her higher-order estimates. Agent h does not infer any (higher-order) new signal from her neighbors who are not strictly closer to i. Let $l \in \mathrm{N}_h$ and $d(il) \geq d(ih)$, then either $\alpha^{hl}_t(n) = 1/N$ or $\boldsymbol{\alpha}^{hl}_t \in A^h_{t-1}$ for all $t \geq 1$.

First, the claim is clearly true at $t = 2$, that is, i's immediate neighbors infer x^i_0 from agent i at $t = 2$. For anyone who is not connected to i, her estimates (and higher-order estimates) remain unchanged. Now consider agent j who is connected to h and is strictly closer to i, that is, $j \in \mathrm{N}_i^{d(ih)-1}$. By the claim, $p_t^j(n) = \Pr\left(s_n \middle| x_0^i, \ldots, x_{t-d(ih)}^i\right)$ and agent h has inferred $x_0^i, \ldots, x_{t-d(ih)-1}^i$ over time. Then at period $t+1$, agent h must infer $x_{t-d(ih)}^i$ from agent j, and by the same argument she must infer the same signal from all her neighbors in $\mathrm{N}_i^{d(ih)-1}$. Agent h uses the simple rule and treats only one copy of these inferred signals as new information. Moreover, because agent h has already learned $x_0^i, \ldots, x_{t-d(ih)-1}^i$, she does not infer any new signal from her neighbors in $\mathrm{N}_i^{d(il)}$ such that $d(il) \geq d(ih)$. Therefore she updates her estimates to $p_{t+1}^h(n) = \Pr\left(s_n \middle| x_0^i, \ldots, x_{t-d(ih)}^i\right)$. ∎

B Appendix: Further discussions

B.1 Common knowledge of the network

In our model, there is a limit to the agents' learning because they don't know the network. We now show that if the network (g, G) and T, the period from which no new signals arrive, are common knowledge, then agents reach the correct Bayesian posterior in finite time. For simplicity, we consider the case with initial signals only.[53] We follow our learning procedure in which each agent continues to report her posterior distribution of the states and receives reports from her local network. The main difference is that each agent may temporarily treat some inferred signals as independent if she cannot tell some of the signals apart. But she gradually revises her report until it is based only on the true signals. And she can do that because she knows the network and how everyone reports.

More specifically, at time t, agent i's information set consists of all the reports she observes (her initial signal is included as her own report at period 1): $I_t^i = \{\mathbf{p}_\tau^h : h \in g_i, \tau \leq t\}$. Because there are only initial signals, to ease notations, we use x^h for x_0^h for all $h \in g$. At $t = 1$, each report is simply the posterior distribution of the states based on each agent's initial signal as before. For $t > 1$, each agent begins with finding the inferred signals $\boldsymbol{\alpha}_t^{ij}, \boldsymbol{\alpha}_t^{ik}, \ldots,$ as in Step 1 of our learning procedure. Because every agent knows the network and how all agents form their estimates, she knows when and how each signal reaches each neighbor (possibly in combination with other signals). Therefore she knows whether any two

[53] It takes longer if $T > 1$ because each agent needs to learn up to TL signals. But their learning outcomes are still Bayesian as long as the agents follow the learning procedure. Alternatively, we can imagine that the agents report and learn each signal sequentially. Namely, they learn all signals arrived at $t = 0$ in the first L^2 periods (see Observation 1 for details), and then the signals arrived at $t = 1$ in the next L^2 periods, and so on. Yet another possibility is for agents to wait until period T to communicate their best estimates based on all their individual signals.

of her inferred signals are independent. If they are, then her estimates $\mathbf{p}_t^i$ are based on the set of true signals available to her: $p_t^i(n) = \Pr(s_n|x^j, \ldots, x^l)$ for all j, l such that $d(ij) < t$ and $d(il) < t$. If some inferred signals are not independent, agent i forms her estimates based on a particular combination of signals from her information set I_t^i. For example,

$$\Pr\left(s_n \middle| u_t^{i1}x^1, \left(u_t^{i2,1}x^2, u_t^{i3}x^3\right), \left(u_t^{i2,2}x^2, u_t^{i4}x^4\right), \ldots, u_t^{iL}x^L\right),$$

where the parentheses denote signals agent i cannot tell apart. For example, the first parentheses means agent i only knows an inferred signal which contains $u_t^{i2,1}$ copies of x^2 and u_t^{i3} copies of x^3. If x^j is contained in multiple inferred signals of agent i, let $u_t^{ij} = \sum_k u_t^{ij,k}$ be the total number of copies of x^j in agent i's estimates at t.

We require that first, the estimates are complete such that agent i's estimates contain all the signals that she has learned. That is, each u_t^{ij} must be a positive integer if $t > d(ij)$, and $u_t^{ij} = 0$ otherwise. Second, the estimates must feature minimal repetition. We order vectors $\mathbf{u}_t^i = (u_t^{i1}, \ldots, u_t^{iL})$ and $\mathbf{v}_t^i = (v_t^{i1}, \ldots, v_t^{iL})$ lexicographically such that $\mathbf{u}_t^i < \mathbf{v}_t^i$ if $u_t^{i1} < v_t^{i1}$; or if $u_t^{i1} = v_t^{i1}$ and $u_t^{i2} < v_t^{i2}$, and so on. Then $\mathbf{u}_t^i$ must be the smallest vector such that $\Pr(s_n|u_t^{i1}x^1, ..., u_t^{iL}x^L)$ is known to agent i.[54] It is easy to see that if agent i knows x^j at time t, u_t^{ij} must be 1, but the reverse is not true. Also, if an agent learns all the individual signals at t, then $\mathbf{u}_t^i = (1, \ldots, 1)$ and her estimates are the correct Bayesian posterior belief. We say agent i *changes* her information set when she infers a new signal that she does not know before. The above learning procedure has the following properties.

LEMMA **7.** *(1) If no one changes their information set at time t, the agents stop learning.*

(2) An agent's estimates agree with the correct Bayesian posterior after at most L changes of her information set.

(3) If $\mathbf{u}_t^i \neq \mathbf{u}_t^j$ and $ij \in G$, at least one of them changes her information set at t.

Proof of Lemma 7: We start with property (1). If no one changes their information set at some time t, it means no one learns new information in the previous period. This also means that their estimates remain the same, and thus no one learns new information in this period. Therefore the agents stop learning. For property (2), each time an agent changes her information set, she must have inferred a new signal containing a different combination

[54] Agents form estimates this way to facilitate the travel of signals to all agents. For instance, suppose agent i's estimates are based on $(x^1, 2x^2, x^3)$, while agent j's estimates are based on $(2x^1, x^2, 2x^3)$ at t. Then at $t+1$, both report $\Pr(s_n|(x^1, 2x^2, x^3))$ by the lexicographical order. This does not mean, however, they discard information contained in agent j's report. They both store it in their information set and use it later to tell signals apart. Clearly, the agents can learn faster if the message space allows them to report both. But we want the learning procedure to be as close to that in the main model as possible.

of signals. Because there are L initial signals, an agent needs to know at most L different combinations of signals to learn all the individual signals. Afterwards, $\mathbf{u}_t^i = (1, \ldots, 1)$ for all t, and her estimates remain constant. Lastly for property (3), without loss of generality, suppose $\mathbf{u}_t^i < \mathbf{u}_t^j$ and $ij \in G$. Then agent j must not know $\Pr(s_n | u_t^{i1} x^1, \ldots, u_t^{iL} x^L)$. If he knew, then $\mathbf{u}_t^j$ is not the smallest vector he can use, which contradicts the learning procedure. Therefore agent j must learn something new and change his information set. ∎

Part (1) is true because if no one infers new signals at period t, they have no new information to pass on to their neighbors, and thus no one learns anything new in the next period. Next, because there are L initial signals, if an agent makes L changes to her information set, she has a system of L non-degenerate equations involving these signals.[55] Clearly, she can solve for each of these signals individually. Lastly, if two neighbors disagree, one of their reports must be unknown to the other agent, as they report the estimates with the minimal repetition given their information set. Therefore at least one of them must learn new information. Using these properties, we can show next that agents always form the correct consensus.[56]

Observation **1.** *There exists some period $t \leq L^2$ such that all agents' estimates after t agree with the Bayesian posterior belief given the signals.*

Intuitively, agents are able to learn eventually because they know the network and they know how information travels. By part (1) of Lemma 7, if learning has not stopped, then there must be at least one agent changing her information set in each period. By part (2), no agent would change more than L times. Therefore their learning must stop within L^2 periods. And part (3) suggests that first, once the learning stops, agents must have consensus. Moreover, since each agent i knows x^i, in her estimates u_t^{ii} must be 1, and thus in the consensus, all agents' estimates contain exactly one copy of each agent's signal.

The following example illustrates the above procedure and shows how agents may be temporarily wrong, but they gradually revise their estimates until they are correct.

Example **7.** *An 8-agent network in Figure 3, which is common knowledge.*

We focus on the learning of agent 1 using the above learning procedure. At $t = 1$, agent 1's information set includes agent 2, 4 and 5's reports $\mathbf{p}_1^2$, $\mathbf{p}_1^4$ and $\mathbf{p}_1^5$. At $t = 2$, agent 1 infers the signals from her neighbors which reflect their signals. We use signals directly instead of

[55] By the lexicographical order, if she makes L changes, the dimension of this system of L equations must be L. It is possible, however, for some agents to change their information sets fewer than L times and thus they cannot distinguish all signals individually. Nevertheless, their estimates still contain one copy of each signal, as shown in the next result.

[56] This result shares the same intuition with Theorem 3 in DeMarzo, Vayanos and Zwiebel (2003).

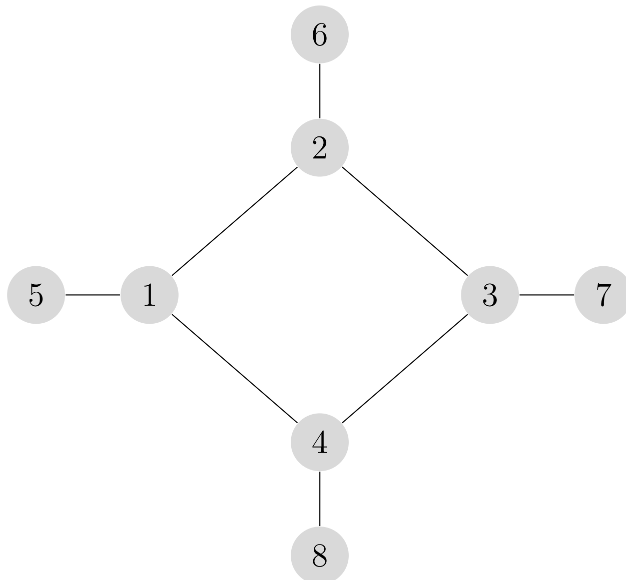

Figure 3: A network of 8 agents.

reports in the following discussion if there is little room for confusion. At $t = 2$, agent 1 reports $p_2^1(n) = \Pr(s_n|x^1, x^2, x^4, x^5)$ and agent 2 reports $p_2^2(n) = \Pr(s_n|x^1, x^2, x^3, x^6)$.

At $t = 3$, agent 1 infers $\boldsymbol{\alpha}_2^{12}$ from agent 2 which contains x^3 and x^6. She also infers $\boldsymbol{\alpha}_2^{14}$ from agent 4 which contains x^3 and x^8. Agent 1 knows the two inferred signals are correlated through x^3, but she cannot tell them apart for now. Similarly, agent 2 knows his inferred signals are correlated through x^4. Agent 1's estimates become $p_3^1(n) = \Pr(s_n|x^1, x^2, x^4, x^5, (x^3, x^6), (x^3, x^8))$ using the above learning procedure. Similarly, agent 2 report $p_3^2(n) = \Pr(s_n|x^1, x^2, x^3, x^6, (x^4, x^5), (x^4, x^7))$. For agent 1, $\mathbf{u}_3^1 = (1, 1, 2, 1, 1, 1, 0, 1)$. She also knows that for agent 2, $\mathbf{u}_3^2 = (1, 1, 1, 2, 1, 1, 1, 0)$.

At $t = 4$, since agent 1 knows x^4, she can easily compute $\Pr(s_n|x^1, x^2, x^3, 2x^4, x^5, x^6)$ and learn x^7 from agent 2. Her estimates become $p_4^1(n) = \Pr(s_n|x^1, x^2, x^4, x^5, (x^3, x^6), (x^3, x^8), x^7)$. Similarly, agent 2 can infer x^8 and report $p_4^2(n) = \Pr(s_n|x^1, x^2, x^3, x^6, (x^4, x^5), (x^4, x^7), x^8)$.

At $t = 5$, agent 1 infers x^8 from agent 2, then she can infer x^3 because she knows $\Pr(s_n|(x^3, x^8))$. Similarly she can infer x^6 from $\Pr(s_n|(x^3, x^6))$. She now knows all the signals individually. So her updated estimates are correct: $p_5^1(n) = \Pr(s_n|x^1, ..., x^8)$. Similarly, agents 2-4 learn all the individual signals.

At $t = 6$, agents 5-8 form the correct Bayesian posterior beliefs, but they do not learn all the individual signals. For instance, agent 5 cannot tell x^2 and x^4 apart. $\diamond$

B.2 Improvement over myopic learning

In this subsection, we compare our learning procedure with the myopic learning model in the spirit of DeGroot (1974) and DeMarzo, Vayanos and Zwiebel (2003), among many others. The essence of myopic learning is that agents update their beliefs by repeatedly treating their neighbors' reports as new information, without accounting for any possible correlations in these reports. In particular, we consider a version of the myopic learning model which shares the same information environment and message space with our main model. That is, the state and signal distributions are finite and an agent reports her estimates of the state distribution. But as in myopic learning models, each agent treats a neighbor's report as a new signal in each period.[57]

More specifically, in Step 1 of this myopic learning procedure, the inferred signal from j is j's estimates, $\boldsymbol{\alpha}_t^{ij} = \mathbf{p}_t^j$. In Step 2, agents update their estimates using the inferred signals as before.[58] Their estimates are formed according to the counterpart of expression (5), where we add m to the superscript to denote estimates in our version of the myopic learning model:

$$p_{t+1}^{i,m}(n) = \frac{\alpha_t^{ii}(n) \prod_{h \in g_i} p_t^{h,m}(n)}{\sum_{n'} \alpha_t^{ii}(n') \prod_{h \in g_i} p_t^{h,m}(n')}. \tag{31}$$

Because myopic learning features local information repetition between any two connected agents, it is not surprising that agents' learning becomes worse. Suppose that in the aggregate, the realized signals are informative: $\Pr(s_n|X_T) \neq 1/N$ for some s_n. Moreover to compare these two learning procedure, let $\eta_t^h(x_\tau^i)$ be the number of copies of x_τ^i ($\tau < T$) in $\mathbf{p}_t^h$, and let $\eta_t^{h,m}(x_\tau^i)$ be the number of copies of x_τ^i in $\mathbf{p}_t^{h,m}$.

Observation 2. *Under myopic learning, the agents' estimates never agree with the Bayesian posterior beliefs in any network if the realized signals are informative. Furthermore, information repetition becomes strictly worse: If a network satisfies GCS,* $\lim_{t\to\infty}(\eta_t^{h,m}(x_\tau^i) - \eta_t^h(x_\tau^i)) = \infty$ *for all* $h \in g$ *and* $x_\tau^i \in X_T$.

Proof of Observation 2: For the first part, suppose that the signals are informative, but the agents agree with the correct Bayesian posterior beliefs at some time $t > T$: $p_t^i(n) = \Pr(s_n|X_T) \neq 1/N$ for some s_n and for all $i \in g$. At $t+1$, agents infer new signals from all neighbors. Each of these inferred signals has a distribution of $\alpha_t^{ij} = \Pr(s_n|X_T)$. Therefore by our myopic learning procedure, $p_{t+1}^i(n) = \Pr(s_n|(L_i+1)X_T)$ where L_i is the number of

[57] This implies that the agents don't remember how they made use of their own signals or their neighbors' reports in the past.

[58] Step 3 is unnecessary because the agents no longer differentiate the new information of each neighbor from the old, which is why agents in our main model need to compute the higher-order estimates.

agent i's neighbors. They continue to do so in the following periods until they believe in the state(s) most likely under X_T with probability one.

For the second part, we start with the case where x_0^i is the only signal. Under myopic learning, agents behave as if their existing estimates are the prior, and treating every report from every neighbor as a new signal. By induction, we can easily show $\eta_t^{h,m}(x_0^i) \geq \eta_t^h(x_0^i)$. Let $\triangle\eta_t^h(x_0^i) = \eta_{t+1}^h(x_0^i) - \eta_t^h(x_0^i)$ be the number of copies of x_0^i that agent h infers at time $t+1$, and similarly $\triangle\eta_t^{h,m}(x_0^i) = \eta_{t+1}^{h,m}(x_0^i) - \eta_t^{h,m}(x_0^i)$. Then we claim that when $t > D$, $\triangle\eta_t^{h,m}(x_0^i) \geq \triangle\eta_t^h(x_0^i) + 1$ for all $h \in g$. When $t > D$, everyone must learn at least one copy of x_0^i. In our learning procedure, everyone must also know that at least one of her neighbors has learned the signal before (because she learns the signal from that neighbor). Therefore the new copies of x_0^i one infers must be strictly fewer than the total number of copies reported by all neighbors, $\triangle\eta_t^h(x_0^i) < \sum_{j\in \mathrm{N}_h} \eta_t^j(x_0^i)$. But under myopic learning, $\sum_{j\in \mathrm{N}_h} \eta_t^{j,m}(x_0^i) = \triangle\eta_t^{h,m}(x_0^i)$. Thus we have:

$$\triangle\eta_t^h(x_0^i) < \sum_{j\in \mathrm{N}_h} \eta_t^j(x_0^i) \leq \sum_{j\in \mathrm{N}_h} \eta_t^{j,m}(x_0^i) = \triangle\eta_t^{h,m}(x_0^i),$$

which implies $\triangle\eta_t^{h,m}(x_0^i) \geq \triangle\eta_t^h(x_0^i) + 1$. It follows that

$$\eta_{t'+1}^{h,m}(x_0^i) - \eta_{t'+1}^h(x_0^i) \geq \sum_{t=D+1}^{t'} \left(\triangle\eta_t^{h,m}(x_0^i) - \triangle\eta_t^h(x_0^i)\right) \geq t' - D.$$

Hence the difference between the number of copies of x_0^i included in the agents' estimates goes to infinity as $t' \to \infty$. ▮

To see why myopic learning cannot lead to the correct learning outcomes if the signals are informative, consider the simplest network with just two connected agents, i and j. At $t = 1$, agent i uses her initial signal and reports $p_1^{i,m}(n) = \Pr(s_n|x_0^i)$, while agent j does not see a signal and reports the symmetric prior. Then at $t = 2$, agent j learns the signal from agent i while agent i learns nothing from agent j, so $p_2^{i,m}(n) = p_2^{j,m}(n) = \Pr(s_n|x_0^i)$. So far, the estimates are the same as in our main model, which are correct given the signals. But learning stops in the main model at $t = 2$, because agent i remembers that j learns the signal from herself and agent j remembers that i has already learned the signal in the last period, and thus there is no new information. In the myopic learning model, however, each agent treats all reports from her neighbor as new and continues to combine her estimates with her neighbor's estimates. It is easy to see that at $t \geq 2$, $p_t^{i,m}(n) = p_t^{j,m}(n) = \Pr(s_n|2^{t-2}x_0^i)$. As t goes to infinity, agents' estimates converge. They believe the true state is one of the state(s) mostly likely given signal x_0^i. If two agents cannot learn one initial signal correctly, no network can guarantee Bayesian learning outcomes.[59]

Recall that when the network is a social quilt, learning is strongly Bayesian in our main model, while it is generally not Bayesian under myopic learning. Furthermore, even when our learning procedure cannot avoid information repetition, for instance in networks that satisfy GCS but contain simple circles, the information repetition is far worse under myopic learning. This shows that the ability to remove information repetition within one's local network can improve the learning outcomes significantly. Lastly, results in Observation 2 continue to hold even when agents have the minimal knowledge of the local network, that is, when eacj agent i only knows (g_i^0, G_i^0). This is because the repetition between any pair of connected agents is the fastest and most severe repetition of information.

[59]This contrasts with the conclusions in DeMarzo, Vayanos and Zwiebel (2003) and Golub and Jackson (2010) that when the network satisfies certain properties, agents can form the correct posterior beliefs asymptotically. Roughly speaking, these models consider large networks where $L \to \infty$. They show that if no single agent's information exerts a disproportionately large influence on others, then by the weak LLN, there are enough agents who receive signals with independent noises such that all of them learn correctly in the limit. Instead, our model focuses on finite networks and a finite number of realized signals.