Learning in Local Networks*

PRELIMINARY AND INCOMPLETE

Wei Li
University of British Columbia

Xu Tan
University of Washington at Seattle

April 17, 2016

Abstract

Agents in a social network learn about the true state of the world over time from their own signals and reports from immediate neighbors. Each agent only knows her local network, consisting of her neighbors and any connections among them. In each period, every agent updates her own estimates about the state distribution based on her 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 the agent's estimates of his estimates, the agent attributes the difference to new information. The agents form the correct Bayesian posterior beliefs 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 completely connected subgroups. Under this procedure, the agents make fewer mistakes than under myopic learning; and they learn correctly if the network is common knowledge.

JEL: D03, D83, D85

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

*We thank Vitor Farinha Luz, Drew Fudenberg, Li, Hao, Stephen Morris, Wing Suen, Jacques Lawarrée, Fahad Khalil, Quan Wen, and especially Matt Jackson and Mike Peters for helpful discussions and comments. Email: wei.li@ubc.ca; tanxu@uw.edu.

1 Introduction

It is often observed that in social networks, all learning is *local learning*: We learn from the people we interact with, who in turn talk to and learn from their neighbors, and so on. How one reacts to such information, often coarsely communicated, depends on one's best assessment of its accuracy. Incorrect inferences may lead to entrenched poverty, political polarization, and other financial and personal failures. A relatively less emphasized feature of social networks is that people often have very little idea of who their indirect neighbors are, and the connections among them. That is, they often know their *local networks* only. This paper investigates what one can learn given such limited knowledge of the network, and any possible inefficiencies that may arise as a consequence.

Consider, for instance, the plight of those who live in a poor neighborhood. Lack of good information of the wider network means that a poor kid may have no incentive to even try because the perceived underlying state is so unfavorable ("the system is against us") that the chance of success is negligible.[1] 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 to get out of the ghetto. One singular failure, however, may reach the poor kid through multiple neighbors. But because he only knows his immediate neighbors, he fails to account for the repetition and duplication of such information. Consequently he believes erroneously—and increasingly so if the information travels back to him again—that the underlying state is far less favorable than it actually is.

The crux of the problem lies in finding a systematic and practical way to process information when agents only know their local networks—their immediate neighbors and any connections among them. One strand of existing literature studies learning when the network structure is common knowledge *and* the agents are capable of sophisticated Bayesian updating. Agents can form correct Bayesian beliefs if everyone knows the structure of their social network—who one's neighbor is

[1] 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. Even more ominously, "only 61 percent of the youth in ghetto neighborhoods reported the most of their friends attended school regularly, compared to 89 percent in low-poverty neighborhoods." See also Mobius and Rosenblat (2001) and Ioannides and Loury (2004) for more discussions.

talking to and their neighbors—as well as the source and time received of all information, as in the tagged information system of Acemoglu, Bimpikis and Ozdaglar (2014). In reality, however, most network knowledge is local.[2] One can conceivably assume that instead, agents have common prior beliefs of the network, and they update both their beliefs of the network and the underlying true state as they learn. But such updating is infeasible for all but the simplest networks.[3]

The other strand of the literature does not assume common knowledge of the network. These papers eschew the complexity of Bayesian learning by assuming that agents follow some "rule of the thumb" to process information (see DeGroot (1974), Ellison and Fudenberg (1993), Ellison and Fudenberg (1995), DeMarzo, Vayanos and Zwiebel (2003), Golub and Jackson (2010), among others.) The classic model of DeGroot (1974) assumes that agents treat their neighbors' information in each period as new and incorporate it into their beliefs. All their neighbors behave in the same way. Thus the same information is repeated both in each agent's local network and through common sources farther away in the network. Often called myopic learning in the literature, this influential model may feature high levels of repetition and no memory of information communication history.[4]

This paper introduces an iterative learning procedure to model learning in social networks. We depart from the existing literature in two ways. First, we do not assume the agents have common knowledge of the network. Instead, each agent behaves as if her local network is the whole network. Thus they do not have to form complicated inferences about their indirect neighbors as they would in a model of Bayesian learning without common knowledge of the network. Second, it is more sophisticated than myopic learning. Agents receive information from their neighbors and try to avoid repetition *within their local network* by separating old, existing information from new information. They then process the perceived new information by Bayes' rule. Doing so reduces information repetition and improves the agents' learning outcomes in the network significantly.

[2]For instance, Krackhardt (1990) show that the accuracy of knowing the connections of other people is between 15%-48% in a small startup with 36 people, and Casciaro (1998) show the accuracy is 42-45% in a research center of 25 people.

[3]The number of possible networks for a given number of agents is often astronomical. 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 $2^{L(L-1)/2}$. With only thirty agents, the number of possible networks is 2^{435}. This is comparable to the number of atoms in the universe, estimated to be between 2^{158} and 2^{246}.

[4]Several recent experiments, such as Grimm and Mengel (2014), find that people learn in a more sophisticated way in that they do try to avoid repeatedly learning the same information.

More specifically, in every period, each agent reports her current estimates of the state distribution to all her neighbors. Each agent then infers the new information contained in each neighbor's report. Next, each agent receives a (possibly uninformative) signal from nature, and forms new estimates of the state distribution based on the new information from her neighbors (if any) and any new signal of her own. The innovation of our model lies in how an agent identifies the new information. To do so, after hearing from all her neighbors, she forms estimates of each of her neighbor's estimates of the state distribution based on her *own* information. Namely, the reports she thinks each neighbor has also heard, which are in general different from the actual reports the neighbor has heard.[5] If any neighbor's report differs 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 agent learns any new information *and* no agent thinks her neighbors learn any new information. The agents form a consensus if their estimates agree at the end of the learning procedure.

We first characterize the properties of this learning procedure. The fact that each agent only knows her own local network means that she does not know her neighbor's local network except for their common neighbors and connections among them. Therefore in forming estimates about each other's estimates, they have to form estimates of what new information they think their neighbors have learned from their common neighbors. In other words, they have to form higher-order estimates of one neighbor's estimates of another neighbor's estimates, and so on. We show, however, each agent only needs to form finitely many orders of estimates, which can be far lower than the actual number of her neighbors. In simple networks with the property that no two agents share two common neighbors who don't know each other, second-order estimates suffice. Thus our procedure is far simpler than it first appears. Furthermore, the signals can be decomposed and analyzed separately. For a fixed signal sequence, if agents' estimates agree with the correct Bayesian posterior beliefs under each signal, they also agree under multiple signals which may arrive at the

[5] Notice that she forms such estimates based on the reports she thinks that a neighbor has also heard, rather than taking the expectation over her best estimate of each neighbor's local network. Therefore we use estimates instead of beliefs to reflect the fact that the agents are not fully Bayesian.

network at different agents and different times.

Next, we show that if the agents' information structures are all partitional as in Aumann (1976), then their estimates are Bayesian regardless of the network structure. In this case, each agent's initial signal informs her about the element of the partition the true state is in. Due to the special feature of the information partition model, each agent's signal can be imprecise, but never wrong. Therefore as agents exchange reports locally, each piece of new information helps eliminate some possible states from an agent's original estimates. Furthermore, repetition does not matter. Suppose an agent hears two identical new reports from her neighbors who have learned the new information from a common neighbor. The agent's estimates are unaffected in that she removes the same set of states from what she thinks is possible as she would with one report. Eventually, the states in the intersections of every agent's element of partition containing the true state are considered equally likely. Agents assign zero probability to all the other states.

For more general information structures, we show that the agents' estimates are Bayesian if the network is a social quilt—a tree-like union of completely connected subgroups. In such networks, there is a unique shortest path from one agent to another; and once a piece of information reaches an agent, it will not reach her again. Suppose there is only one signal. Because each subgroup of agents is fully connected, when this piece of information arrives at one member of the subgroup, all members can identify this as new information in the next period. More importantly, all members correctly infer that all others in the subgroup have learned this information at the same time (instead of each member has received a new, identical piece of information). Thereby they avoid learning the same information repeatedly. In this way, each agent learns every signal once and exactly once, and they form the correct estimates D periods after the last informative signal, where D is the diameter—the longest shortest distance between any pair of agents.

If a network is not a social quilt, it must contain a simple circle of more than three agents in which every agent is only connected to her two adjacent neighbors; and/or involve agents with asymmetric information of their local network.[6] In the first case, suppose that there is only one simple circle in the network which receives one signal. Once the signal reaches the circle, it travels

[6] In a simple circle, the first agent is the last, and each agent has exactly two links to others in the circle.

in both directions. The last agent of the circle thinks that there are two copies of the same signal, and passes her estimates onto her neighbors. Each of her two neighbors knows that one copy comes from himself, but treats the other copy as new. Each piece of information thus gets double counted every time it travels around the circle. This process continues until everyone in the circle becomes extremely confident, but wrong. They believe, with probability one, the state is the one most likely given the initial signal. The problem is exacerbated in the presence of multiple circles. Not only the agents may become overconfident; their estimates may never agree. With one simple circle, the number of repeated signals grows linearly in time, and thus the signal with the highest precision dominates in the long run. But with multiple circles, the number of repeated signals grows exponentially, causing some agents' estimates to oscillate and never converge.

In the second case, when a network is not a social quilt, agents may have asymmetric information even within a local network. Suppose that agents i and j share two common neighbors, k and k' who are not connected to each other. Then a signal from agent k', which is the only signal, has different effect on agents i, j and agent k. Agent i and j learn correctly and do not change their estimates afterwards. Agent k, however, initially believes that there are two copies of the same signal, one from i and one from j. More importantly, agent k attributes agent i and j's unchanging reports to them each receiving an opposite signal in the following period. Agent k's estimates oscillate forever, which is clearly not Bayesian, because he thinks that he is learning new information from i and j every period. These two impediments to Bayesian learning, simple circles and local asymmetric information, imply that for any network that is not a social quilt, there exists a signal sequence such that the agents form incorrect estimates.

To relate our model to the existing literature, we show that if our agents have common knowledge of the network, then they form the correct Bayesian beliefs within a finite number of periods. In particular, agents' estimates may still temporarily include the same signal multiple times, but as the network is common knowledge, agents can account for such double counting and eventually learn correctly. If instead, our agents learn myopically, then they still learn from their neighbors' reports and update their own estimates, but they treat the reports in each period as new information. Due to this information repetition even within their local networks, agents fail to learn correctly even in social quilts. In general, our agents do better than those under myopic learning because in each period, their estimates do not feature information repetition in each local network.

In one extensions of our model, we allow agents to put different weights on the reports, depending on where and when they are received. Doing so allow us to consider the possibility of opinion leaders, stubborn agents, and imperfect information diffusion. The main properties of our learning procedure remain. When agents discount information received later significantly, however, they may disagree with each other forever at the end of the learning procedure. In addition, the influence one has over another decreases in their social distance (see Mobius, Phan and Szeidl (2015)); and polarization of opinion can appear between agents farther away from each other. We also consider the extension in which agents build their communication network endogenously. In this case, a social quilt can be an equilibrium outcome when agents are sufficiently patient.

Our paper contributes to the large and growing literature on social learning in networks. In addition to the papers cited above, our paper is also related to Bayesian models of observational learning (for example, Banerjee (1993), Bala and Goyal (1998); Acemoglu et al. (2011); Mossel, Sly and Tamuz (2015), among others). Their focus is on agents' mistakes in inferring useful information from the observed actions, leading to errors in information aggregation. Instead, we use a richer message space of the agents' estimates in order to make it possible for the agents to infer the new information from their local networks.

Our paper is also related to the recent literature on providing theoretical foundation for fully connected subnetworks, or cliques, and high clustering often found in real networks.[7] Jackson, Rodriguez-Barraquer and Tan (2012) show that a favor exchange relationship can be sustained when a pair of agents shares common friends, who can punish a deviator by removing their links with him. As a result, cliques are robust because any removal of one link only leads to further link deletion in the local clique. Relatedly, Ali and Miller (2013) show that cliques can support more cooperation by reducing the travel time of bad behavior and speeding up the punishment for deviations. Our model provides a complementary reason: Cliques and high clustering allow agents

[7]Clustering is a measure of the likelihood that one agent's neighbors are neighbors with one another. Empirical analysis found clustering to be much higher than predicted in a random network. See for example MacRae (1960), Goyal, van der Leij and Moraga-Gonzlez (2006) and Adamic (1999).

to distinguish new information from old ones, and thus form correct Bayesian posterior beliefs even when they only know their immediate neighbors.

This paper is also related to the literature on knowledge and consensus (see Aumann (1976), Geanakoplos and Polemarchakis (1982), Parikh and Krasucki (1990), Mueller-Frank (2013), among many others). This literature analyzes under which conditions—and under what reporting protocols—repeated communication among a finite set of individuals leads to consensus. In this context, one can think of our procedure as a more general reporting protocol which allows more than pairwise communication, and expands the message space to the posterior beliefs of all the states. As a result, not only the agents in our model agree, they form the correct consensus sooner.

Moreover, a large existing literature studies networks empirically, such as Conley and Udry (2001), Munshi (2003) and Munshi (2004). Several recent papers use experiments to study social learning in networks. Chandrasekhar, Larreguy and Xandri (2012) compared Bayesian learning with myopic learning in the lab. They found that while the myopic learning model performs better than Bayesian learning, it can only explain 76% of the actions taken. More recently, Grimm and Mengel (2014) found that subjects do account for correlations in their neighbors' estimates of the true state, but they do so in a more rudimentary way than Bayesian learners. Similarly, Mobius, Phan and Szeidl (2015) also show that their data is more consistent with a sophisticated learning model in which people try to avoid double-counting signals.

We introduce our learning procedure in Section 2.1 and characterize its properties in Section 3. Section 4 presents our results on when Bayesian learning outcomes are achievable, and impediments to Bayesian learning. We then adapt our learning procedure to the case of normally distributed signals before comparing it with the classic DeGroot model in Section 5. Section 6 extends our procedure to allow agents to weigh information differently. Section 7 considers several extensions of the main model and concludes.

2 An iterative learning procedure

We begin with a formal description of our learning procedure, and discuss several important aspects of our model afterwards.

2.1 Model setup

Consider a network (g, G): $g = \{1, 2, ..., L\}$ represents a finite set of agents, and G represents the links among the agents. $G(ij) = 1$ if i and j are connected; $G(ij) = 0$ otherwise. The network is undirected, so information flows both ways: $G(ij) = G(ji)$. The network is also path-connected, so it is possible for information to diffuse throughout the network. Formally, for any $i, j \in g$, there is a path $\{i_0, i_1, .., i_k\}$ such that $i_0 = i$, $i_k = j$ and $G(i_l, i_{l+1}) = 1$ for all $l < k$. Denote $\mathrm{N}_i = \{j : G(ij) = 1\}$ as the set of agent i's neighbors, and $L_i = |\mathrm{N}_i|$ as the number of agent i's neighbors. Let $g_i = \mathrm{N}_i \cup i$. We refer to (g_i, G_i) as agent i's *local network*, consisting of agent i, her neighbors, and all the links among them in the original network.

We do not assume the network (g, G) is common knowledge; nor do we assume agents have common priors over the network structure.[8] To give our model enough structure to make it tractable, we maintain the following two assumptions, one on knowledge and one on behavior, throughout the main model.

A1: *know thy local network.* Every agent i knows (g_i, G_i).

A2: *out of sight, out of mind.* Every agent i behaves as if (g_i, G_i) is the whole network (g, G).

Under Assumption A1, each agent has local knowledge. She knows her neighbors and the connections among them. Under Assumption A2, agents use only the information from their local network to learn. They are completely agnostic about the part of the network they don't know.

We first give an intuitive description of our iterative learning procedure. The agents in the network aim to learn the true state of the world s given all the information available. Time is discrete with an infinite horizon: $t = 1, 2, ...$ In each period prior to T, each agent receives a (possibly uninformative) signal. The signals are independent conditional on the state across agents and time. At the beginning of each period, each agent forms and reports her most up-to-date estimates of the state distribution given her information so far. Agent i learns all the reports in her local network and later, her new signal. She then proceeds to update her estimates given these

[8] The only exception is when we discuss the property of our learning procedure when the network is common knowledge in Section 7.2.

new information. The innovation of our model is that she also forms estimates of each neighbor's estimates in her local network, based on the information she knows a neighbor has observed so far. Then in the next period, if her neighbor j's reported estimates differ from i's estimates of j's estimates, agent i knows that agent j must have received information somewhere in period t unobservable to her. She then incorporates this information in forming new estimates, both her own and those of her neighbors. This procedure continues until every agent's estimates stop changing.

We now describe our learning procedure more formally. There are finitely many states of the world and signals. The state s is distributed as $s \in S = \{s_1, s_2, s_3, ..., s_n, ..., s_N\}$. Let the agents have symmetric prior beliefs: $\Pr(s = s_n) = 1/N$ for all s_n and all $i \in g$.[9] The true state s^* is realized at the beginning of $t = 0$. The set of agent i's signals is finite and includes an uninformative signal, $x_\emptyset$. Each signal occurs with a probability strictly between 0 and 1. The above is common knowledge to all agents in the network. The agents, however, do not know T, the period since which agents receive no informative signals.

Agent i's signal x^i is distributed as $x^i \in X^i = \{x_\emptyset, x_1, x_2, ...x_m, ..., x_{M_i}\}$. Let the unconditional probability of agent i receiving the uninformative signal be $\psi_0^i \in (0,1)$ and that of receiving an informative signal x_m be $\psi_m^i \in (0,1)$. Also, let the distribution of the signals conditional on the state be Φ^i such that $\phi_{mn}^i = \Pr(x_m|s_n)$ for agent i. By Bayes' rule,

$$\psi_m^i = \sum_{n'} \Pr(x_m|s_{n'})\Pr(s = s_{n'}) = \sum_{n=1}^{N} \phi_{mn}^i/N.$$

[9] If the agents have asymmetric prior beliefs at $t = 0$, we can model a previous period $t = -1$ in which agents with symmetric priors receive different signals that lead to different posterior beliefs.

The set X^i, the ex ante distribution of signals, as well as the following matrix,

$$\begin{array}{c} \\ s_1 \\ \vdots \\ s_n \\ \vdots \\ s_N \end{array} \begin{array}{c} \begin{array}{cccccc} x_\emptyset & x_1 & \dots & x_m & \dots & x_{M_i} \end{array} \\ \begin{pmatrix} \psi_0^i & \phi_{11}^i & \dots & \phi_{m1}^i & \dots & \phi_{M_i1}^i \\ \psi_0^i & \dots & \dots & \dots & \dots & \dots \\ \psi_0^i & \phi_{1n}^i & \dots & \phi_{mn}^i & \dots & \phi_{M_in}^i \\ \psi_0^i & \dots & \dots & \dots & \dots & \dots \\ \psi_0^i & \phi_{1N}^i & \dots & \phi_{mN}^i & \dots & \phi_{M_iN}^i \end{pmatrix} \end{array}$$

is agent i's private information. In each period, agent i observes a private signal x_t^i which is generated according to the matrix above.[10] All the entries of the matrix $\phi_{mn}^i \in [0,1]$, with each row summing up to 1 and each column summing up to a positive value. The distribution of state conditional on signals can then be derived from Bayes' rule:

$$\mu_{nm}^i \equiv \Pr(s_n|x_m) = \frac{\phi_{mn}^i}{\sum_{n'} \phi_{mn'}^i} = \frac{1}{N} \cdot \frac{\phi_{mn}^i}{\psi_i^m}.$$

Clearly, $\mu_{nm}^i \in [0,1]$ for all i. Every signal except for $x_\emptyset$ is informative. That is, for every signal x_m, there must exist a state s_n such that $\phi_{mn}^i \neq \psi_m^i$.

We now turn to how agents learn from their local network. Since $t = 0$ is the first period, each agent simultaneously reports her prior beliefs of the state,

$$\mathbf{p}_0^i = \{p_0^i(1), ..., p_0^i(N)\} = \{1/N, 1/N, ..., 1/N\}.$$

Agent i observes the reports of all agents who are connected to her, and then agent i receives signal $x_0^i \in X^i$. Denote all the information available to agent i *at the end of period* t as I_t^i, where the superscript denotes agent and the subscript denotes time. Then $I_0^i = \{\mathbf{p}_0^h : h \in g_i\} \cup x_0^i$ for all i. At the beginning of $t = 1$, she proceeds to update her estimates of the state given her

[10] It is straightforward to allow the signal generating process to vary over time, but it won't change our results qualitatively as long as the signal generating process is independent across time and agents.

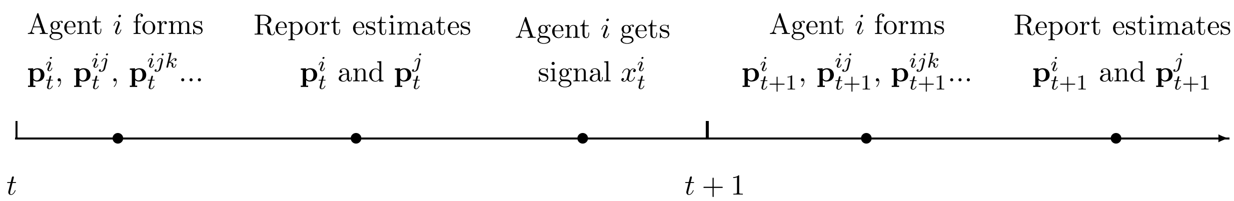

Figure 1: Time line

information I_0^i. That is, $\mathbf{p}_1^i = \{p_1^i(1), ..., p_1^i(N)\}$ where $p_1^i(n) = \Pr(s_n | I_0^i)$. She then simultaneously reports $\mathbf{p}_1^i$ to all the other agents in her network. Her information set at the end of period 1 is $I_1^i = \{\mathbf{p}_1^h : h \in g_i\} \cup \{I_0^i, x_1^i\}$.

The innovation in our model is that each agent tries to incorporate only the new information from her neighbors in updating her own estimates. To do so, agent i needs to form initial estimates of each of her neighbors' estimates of the state distribution. At the beginning of period 1, agent i forms an estimate about agent j's estimates: $\mathbf{p}_1^{ij} = \{p_1^{ij}(1), ..., p_1^{ij}(N)\}$. To do so, agent i uses the information that i knows j observes at the end of period 0. Let $g_{ij} = g_i \cap g_j$ be the *shared local network* of agent i and j. Then

$$p_1^{ij}(n) = \Pr(s_n | I_0^{ij}), \text{ where } I_0^{ij} = \{\mathbf{p}_0^k : k \in g_{ij}\}.$$

Clearly, $\mathbf{p}_1^{ij} = \{1/N, 1/N, ..., 1/N\}$ for all i, j. Note that i's estimate of j's estimates is based only on j's information observable to agent i. Similarly, agent i forms estimates of agent j's estimates of agent k's estimates when both $j, k \in g_i$. This is based on information i knows j knows k observes, that is, $I_0^{ijk} = \{\mathbf{p}_0^h : h \in g_{ijk}\}$ where $g_{ijk} = g_i \cap g_j \cap g_k$. We can define estimates $\mathbf{p}_1^{ijkh}, \mathbf{p}_1^{ijkjh}, ...$, in a similar fashion. Initially, all these estimates are the symmetric priors.

Given these initial set of estimates, agents update their estimates in a similar way in each period. Let $\mathbf{p}_t^i$ denote agent i's estimates of the state distribution at the beginning of time t, before hearing her neighbors' reports and receiving any new signal x_t^i. Also, $\mathbf{p}_t^{ij}$ denotes i's estimates of j's estimates of the state at the beginning of time t; $\mathbf{p}_t^{ijk}$, $\mathbf{p}^{ijkh}$, ..., are formed similarly. Agents then report their estimates $\mathbf{p}_t^i$ and $\mathbf{p}_t^j$. Then each agent receives a new private signal. Figure 1 summarizes the timing of events in period t for agent i's local network.

Step 1: Identify new information. From agent i's perspective, for every $j \in \mathrm{N}_i$, j has no new information if $\mathbf{p}_t^j = \mathbf{p}_t^{ij}$ (the equality holds component wise). Otherwise j must have had new information since his previous report $\mathbf{p}_{t-1}^j$ that i does not know. Note from the above time line that in our updating procedure, i has already incorporated information that i thinks j has learned from her and their common neighbors in period $t-1$ in forming estimate $\mathbf{p}_t^{ij}$ (to be specified in Step 3). Similarly, if $\mathbf{p}_t^j = \mathbf{p}_t^{ikj}$, agent i thinks that k does not think j has new information since j's previous report. Agent i proceeds to compare all her higher-order estimates with $\mathbf{p}_t^j$.

If $\mathbf{p}_t^j \neq \mathbf{p}_t^{ij}$, agent i thinks that j had new information since j's previous report. Let y_{t-1}^{ij} be the *inferred signal* agent i learns from j at period t. That is, for each state s_n,

$$p_t^j(n) = \Pr(s_n \big| I_{t-1}^{ij}, y_{t-1}^{ij}).$$

Then from agent i's perspective, by Bayes' rule, for each s_n:

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

The numerator of the right hand side of equation (1) is, for agent i, the joint probability of observing this inferred signal and the state is s_n; and the denominator is the total probability that agent i believes agent j receives signal y_{t-1}^{ij}. Then the vector $\mathbf{\Delta p}_t^{ij} \equiv \{\alpha_t^{ij}(1), ..., \alpha_t^{ij}(N)\}$, where

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

is the conditional distribution of the inferred signal i learns from j, up to a positive multiplier. By equation (1), if $p_t^{ij}(n) = 0$ for some state s_n, then $p_t^j(n) = 0$, in which case we define $\frac{0}{0}$ as 0. Clearly, $\sum_{n'} \alpha_t^{ij}(n') = 1$, and $\alpha_t^{ij}(n) = 1/N$ for all s_n if $\mathbf{p}_t^j = \mathbf{p}_t^{ij}$.[11]

In a similar way, i identifies the new information she thinks j learns from herself and from their common friend k in period t. If $\mathbf{p}_t^{iji} = \mathbf{p}_t^i$, agent i thinks j does not learn new information from herself. Also, if $\mathbf{p}_t^{ijk} = \mathbf{p}_t^k$, agent i thinks j does not learn new information from k. Otherwise, the

[11]Instead of cluttering the exposition of our learning procedure here, we discuss in the next subsection on why $\mathbf{\Delta p}_t^{ij}$ is the part of the inferred signal that can be learned by all j's neighbors.

new information i thinks j learns from herself and from k, is respectively:

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

That is, agent i treats the difference between k's reported, most up-to-date estimates, and what she believes to be j's estimates of k's estimates, as new information agent j learned from k. Let $\mathbf{\Delta p}_t^{iji} \equiv \{\alpha_t^{iji}(1), ..., \alpha_t^{iji}(N)\}$ and $\mathbf{\Delta p}_t^{ijk} \equiv \{\alpha_t^{ijk}(1), ..., \alpha_t^{ijk}(N)\}$. Agent i identifies new information on all higher-order estimates in a similar way.

Step 2: Update own estimates. Agent i may observe some new signal $x_t^i = x_m$ with the aforementioned signal distribution. For consistency and ease of exposition, let $\alpha_t^{ii}(n) = \mu_{nm}^i$ be the new information agent i learns from nature. Then $\mathbf{\Delta p}_t^{ii} = \{\alpha_t^{ii}(1), ..., \alpha_t^{ii}(N)\}$.[12] Given the new information $\mathbf{\Delta p}_t^{ij}$ for all $j \in \mathrm{N}_i$, she updates her own estimates using her new signal and the new information.

$$p_{t+1}^i(n) = \frac{p_t^i(n) \prod_{h \in g_i} \alpha_t^{ih}(n)}{\sum_{n'=1}^N p_t^i(n') \prod_{h \in g_i} \alpha_t^{ih}(n')}. \tag{4}$$

Step 3: Update estimates of neighbors' estimates. When forming new 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 herself, $\mathbf{\Delta p}_t^{iji}$; and their common neighbors $k \in \mathrm{N}_i \cap \mathrm{N}_j$, $\mathbf{\Delta p}_t^{ijk}$. Recall that $g_i = \mathrm{N}_i \cup i$, thus $g_{ij} = g_i \cap g_j = \{\mathrm{N}_i \cap \mathrm{N}_j\} \cup \{i, j\}$. Similarly for g_{ik}. Agent i's second-order estimates are formed by Bayes' rule such that for each s_n:

$$p_{t+1}^{ij}(n) = \frac{p_t^j(n) \prod_{h \in (g_{ij}/j)} \alpha_t^{ijh}(n)}{\sum_{n'=1}^N p_t^j(n') \prod_{h \in (g_{ij}/j)} \alpha_t^{ijh}(n')}. \tag{5}$$

Observe from equation (5) that agent i's estimates of j's estimates are simply j's own estimates one period earlier if i thinks that there is no new information available to j. Agent i's third-order

[12] If the signal from nature is uninformative, when $t \geq T$, then $\mathbf{\Delta p}_t^{ii} = \{1/N, ..., 1/N\}$.

estimates are formed similarly:

$$p_{t+1}^{ijk}(n) = \frac{p_t^k(n) \prod_{h \in ((g_{ij} \cap g_{ik})/k)} \alpha_t^{ijkh}(n)}{\sum_{n'=1}^{N} p_t^k(n') \prod_{h \in ((g_{ij} \cap g_{ik})/k)} \alpha_t^{ijkh}(n')}. \tag{6}$$

Each of agent i's lower-order estimates are thus formed from her next higher-order estimates. And this proceeds iteratively. ‖

2.2 Remarks and an illustrating example

Having defined our learning procedure in Section 2.1, we first comment on several features of our model. Then we use a simple but important example to illustrate how this procedure works.

Local knowledge of the network. We aim to design the simplest learning procedure that still allows agents to account locally for old information already learned. Assumption A1 helps the agents differentiate one piece of information circulating within her local network from new information arriving from outside the local network. Assumption A2 requires agents to use only the information from their local network to learn. They are completely agnostic about the part of the network they don't know. These assumptions are crucial in how agents identify new information in Step 1, because each agent compares new reports from her neighbors with her estimates of the same neighbors and treats the difference as the new inferred signals. Moreover, in order to estimate a neighbor j's estimates, agent i needs to infer the new information j learns from their common friend k, To do so, agent i uses only the neighbors who she knows are connected with j and k to estimate j's estimates of k's. She does not try to form prior beliefs of j and k's local networks and update them based on new information.

Knowledge of local network only is also the reason that agents may need to form higher-order estimates, such as i's estimates of j's estimates of k's. Observe from Step 1 of the procedure that $\mathbf{\Delta p}_t^{jk}$, the new information agent j learns from k, is based on j's information set I_{t-1}^{jk}. It is generally different from $\mathbf{\Delta p}_t^{ik}$, which is based on i's information set I_{t-1}^{ik}. Intuitively, this is the case if agent i and k share some common friends not connected to j, and thus they may have information unknown to j. Therefore, agent i needs to estimate $\mathbf{\Delta p}_t^{jk}$ using information available to all three of them, and so on for higher-order estimates.

Communication protocols. Our agents report their most up-to-date estimates of the distribution of the state. Although these reports are richer than those in the observational learning literature, which typically involve agents' actions or payoffs (Bala and Goyal 1998; 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 l," and so on. We want to model a simplified communication process since the agents neither know the network beyond their immediate neighbors, nor do they pass on information so precisely in real life. We discuss the message space needed for our agents to achieve full Bayesian learning in Section 7.

Inferred signals. Intuitively, agent i uses her own estimates of j's estimates of how likely the state is s_n as her prior, and agent j's actual report as her posterior. She then applies the Bayes' rule as if agent j's new information comes from one new signal, whether it is actually 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 (except for their common neighbors), she cannot differentiate the sources of j's information. Therefore treating all j's information as coming from one signal results in the same updating given agent i's local knowledge.

More specifically, the vector $\mathbf{\Delta p}_t^{ij}$ is the part of the inferred signal relevant to agent j's neighbors's learning. Recall that no one knows agent j's ex ante signal distribution (except that it is finitely distributed). Suppose that agent i believes, *e*x ante, j receives the inferred signal y_{t-1}^{ij} with probability ψ_{t-1}^{ij}. To find out all the $\Pr(y_{t-1}^{ij}|s_n)$, agent i has $N-1$ equations from (1) because $\sum_n p_t^j(n)=1$. Moreover, by Bayes' rule: $\sum_{n'}\Pr(y_{t-1}^{ij}|s_n)/N=\psi_{t-1}^{ij}$. Thus for any ψ_{t-1}^{ij},

$$\Pr(y_{t-1}^{ij}|s_n)=N\psi_{t-1}^{ij}\cdot\left(\frac{p_t^j(n)}{p_t^{ij}(n)}\Big/\sum_{n'}\frac{p_t^j(n')}{p_t^{ij}(n')}\right)=N\psi_{t-1}^{ij}\alpha_t^{ij}(n).$$

Because the inferred signals always enter an agent's updating multiplicatively, $\alpha_t^{ij}(n)$ is sufficient for the updating purpose for any ex ante belief of ψ_{t-1}^{ij}. Thus we treat $\mathbf{\Delta p}_t^{ij}$ as the inferred signal.

Distribution of the state and signals. For concreteness, we illustrate our learning procedure with a model of finitely many states and finitely many signals. But the idea behind our learning procedure is applicable to many 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 correspondence $\mathcal{P}^i : S \to 2^S/\emptyset$. $\mathcal{P}^i$ associates each state s_n with a non-empty element $P^i(s_n)$ such that at s_n, the agent considers $P^i(s_n)$ to be the set of possible states. Moreover, $\mathcal{P}^i$ induces an information partition over the state space if (1) for any $s_n \in S$, $s_n \in P^i(s_n)$; and (2) for any $s_n, s_{n'} \in S$, $s_n \in P^i(s_{n'})$ implies $P^i(s_n) = P^i(s_{n'})$. In our model, each signal x_m informs agent i of an element $P^i(s_n)$ of her partition: for all $s_n \in P^i(s_n)$, $\phi^i_{mn} = 1$; and $\phi^i_{mn} = 0$ otherwise. The number of signals each agent has corresponds to the number of elements in her partition. Intuitively, the information partition model is equivalent to all the entries in each agent's matrix above being 0 or 1.

Second, we can easily accommodate the model with uniformly distributed states and normally distributed signals even though both the state and the signals are continuous. This extension is presented in Section 5, which we then use to compare our model with the myopic learning models.

Finally, the binary state and binary signal model often used in experimental studies of social networks, is a special case of our model. Specifically, suppose that the true state is $s \in S = \{0, 1\}$ and that the signal is $x \in X = \{0, 1\}$. Then our procedure is particularly simple because we only need to track each agent's estimate of the true state being 1. We use this binary model repeatedly to illustrate different features of our learning procedure.

Example **1.** *Consider a binary state and binary signal setting, with* $\Pr(x^i_t = 1|s = 1) = \Pr(x^i_t = 0|s = 0) = \phi^i$. *$(g, G)$ is a fully connected network (clique), that is, $G(ij) = 1$ for all $i, j \in g$. All L agents hold symmetric prior beliefs, $p^i_0(1) = \frac{1}{2}$.*

To begin with, recall that $\mathbf{p}^t_i = (p^i_t(0), p^i_t(1))$. Because of the binary state, we only need to keep track of $p^i_t(1)$ for simplicity. At $t = 0$, $p^i_t(1) = \frac{1}{2}$ for all $i \in g$. Agent 1 observes a signal $x^1_0 = 1$. At $t = 1$, $p^1_t(1) = \phi^1$ by Bayes' rule, and $p^i_t(1) = \frac{1}{2}$ for all $i \geq 2$. Also, because the second-order estimates are based on everyone's report at $t = 0$, $p^{jk}_1(1) = p^{kj}_1(1) = \frac{1}{2}$ for all $j, k \in g$. Similarly, all the higher-order estimates remain $\frac{1}{2}$. We keep the notation that $i \geq 2$, $h \neq i$ and $i' \neq 1, i$.

Step 1: After hearing the reports at $t = 1$, agent i notices that $p^1_1(1) \neq p^{i1}_1(1)$. Agent i's inferred signal is $\mathbf{\Delta p}_1^{i1} = (\alpha_1^{i1}(1), \alpha_1^{i1}(0))$. By expression (2),

$$\alpha_1^{i1}(1) = \frac{p_1^1(1)}{p_1^{i1}(1)} \Bigg/ \left(\frac{p_1^1(0)}{p_1^{i1}(0)} + \frac{p_1^1(1)}{p_1^{i1}(1)} \right) = 2\phi^1/2 = \phi^1,$$

and thus $\alpha_1^{i1}(0)) = 1 - \phi^1$. From now on, we only track the probability that the state is 1 given the inferred signal. Similarly, for any sequence of agents $jk...i$, $p_1^1(1) \neq p_1^{jk...i1}(1)$, and thus $\alpha_1^{jk...i1}(1) = \phi^1$. Note that each agent (except agent 1) learns the same inferred signal from agent 1, and they also think other agents learn the same inferred signal from 1. Moreover, since $p_1^i(1) = p_1^{hi}(1)$ for any $h \neq i$, h does not learn any new information from i, $\alpha_1^{hi}(1) = \frac{1}{2}$. Similarly $\alpha_1^{jk...hi}(1) = \frac{1}{2}$.

Step 2: Agent 1 does not learn any new information from her neighbors, so $p_2^1(1) = p_1^1(1) = \phi^1$. Agent $i \geq 2$ learns $\alpha_1^{i1}(1) = \phi^1$ from agent 1 and no new information from other agents, so $p_2^i(1) = \phi^1$. These estimates agree with the correct Bayesian posterior beliefs.

Step 3: When agent i forms her estimates of agent 1, she starts with agent 1's estimates $p_1^1(1) = \phi^1$ and incorporates the new information i thinks 1 learns from i and from their common friends: $\alpha_1^{i1i}(1) = \frac{1}{2}$ and $\alpha_1^{i1i'}(1) = \frac{1}{2}$ ($i' \neq 1, i$) from step 1. So $p_2^{i1}(1) = \phi^1$. When agent $h \neq i$ forms her estimates of agent i, she follows the same procedure, using $p_1^i(1) = \frac{1}{2}$, $\alpha_1^{hi1}(1) = \phi^1$ and $\alpha_1^{hii'}(1) = \frac{1}{2}$. By expressions (5), $p_2^{hi}(1) = \phi^1$. Similarly, by expression (6), all high-order estimates that $p_2^{j...k}(1) = \phi^1$. So agents' estimates and all their high-order estimates agree.

At $t = 2$, $p_2^j(1) = p_2^{j...k}(1) = \phi^1$. So no one learns anything new. The learning stops with all agents's estimates agree with the Bayesian posterior beliefs given signal x_0^1. $\diamond$

3 Main properties

Having defined our learning procedure above, we proceed to characterize its main properties. To begin with, our learning procedure may seem to impose a heavy computational burden on the agents because it requires them to form higher and higher order of estimates. We now show that given assumptions A1 and A2, each agent only needs to form a finite—and possibly far lower than the number of her neighbors—order of estimates about neighbors in her local network.

Notice that $\mathbf{p}_t^{i...j}$ is a valid high-order estimate if and only if any pair of agents in $\{i...j\}$ is connected.[13] So we group g_i into subsets within which all agents are connected; these subsets may be overlapping. Without of loss generality, suppose the largest subset is $\{i, j, k_1, ...k_z\}$. Let the number of agents in the largest subset be $\overline{L}^i$, which is $2+z$. For example, in the network depicted in Figure 2 below, $\mathrm{N}_i \cap \mathrm{N}_j = \{k, k'\}$. Agent i has three neighbors, and the largest connected subsets

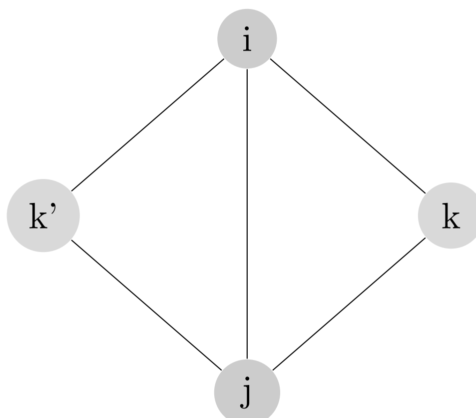

Figure 2: A four agent example

are $\{(ijk), (ijk')\}$ so $\overline{L}^i = 3$. Agent i, j and k have asymmetric information even in the same local network. When agent i form estimates about j's estimates, she knows that they both have access to reports from k and k' (in addition to their own reports). But when she estimates j's estimates of k's estimates, she only uses the reports from i, j and k. Thus it is necessary to explicitly calculate agent i and j's higher-order estimates when k and k' have informative signals.

Despite the possibility of asymmetric information within agents' local networks, we first show that each agent only needs to form a finite order of estimates. This is because agents can only base their higher-order estimates on the agents *who are connected to all of them*, and there are only finitely many of such neighbors. All the proofs are in the Appendix.

Proposition 1. *Consider the local network (g_i, G_i). The highest order of estimates agent i needs to form about her neighbors is $\overline{L}^i$.*

[13] If $G(kk') = 0$, agent i's higher-order estimate cannot feature agent k and k' together because they do not know the existence of each other.

Intuitively, let $distinct(i...j)$ be the set of all distinct agents in the sequence of (possibly repeated) agents $\{i...j\}$. We first show that for any high-order estimates, only distinct agents matter, that is if $distinct(i...j) = distinct(k'...k)$ and they are connected, then $\mathbf{p}_t^{i...j} = \mathbf{p}_t^{k'...k}$ for all t. This is because the shared local network they belong to is the same one:

$$g_{i...j} = g_i \cap ... \cap g_j = g_{k'} \cap ... \cap g_k = g_{k'...k}.$$

Consequently, the information sets they rely on are the same: $I_t^{i...j} = I_t^{k'...k} = \{\mathbf{p}_\tau^l : l \in g_{i...j}, \tau \leq t\}$, so the high-order estimates must be the same. With this property, we can focus on high-order estimates of distinct agents below.

Next, the more distinct agents we have in a sequence, the smaller their shared network becomes. For example, agent i's estimates of j's estimates, $\mathbf{p}_t^{ij}$, can easily differ from agent i's estimates of k_1's estimates of j's, $\mathbf{p}_t^{ik_1j}$. Observe from equation (5) and (6) in Step 3 of our learning procedure, agent i uses information from more agents to form her second-order estimates $\mathbf{p}_t^{ij}$ than her third-order estimates $\mathbf{p}_t^{ik_1j}$ ($h \in g_{ij}/j$ versus $h \in (g_{ij} \cap g_{ik_1})/j$ respectively). To see this, note that

$$g_{ij}/j = \{\mathrm{N}_i \cap \mathrm{N}_j\} \cup i = \{\mathrm{N}_i \cap \mathrm{N}_j \cap \mathrm{N}_{k_1}\} \cup \{(\mathrm{N}_i \cap \mathrm{N}_j)/\mathrm{N}_{k_1}\} \cup i \supseteq \{\mathrm{N}_i \cap \mathrm{N}_j \cap \mathrm{N}_{k_1}\} \cup \{i, k_1\} = (g_{ij} \cap g_{ik_1})/j;$$

and they are equal if $\{(\mathrm{N}_i \cap \mathrm{N}_j)/\mathrm{N}_{k_1}\} = k_1$. This is because higher-order estimates require agent i to use reports from common neighbors of all the involved agents. As the order of estimates increases, more distinct agents are involved, and they share fewer common neighbors. Since each of the local network is finite, the order of estimates reach a maximum at some point, with no more distinct common neighbors outside the sequence. This maximum is clearly $\overline{L}^i$, from the subset with the most distinct agents in agent i's shared local networks. From then on, all the higher-order estimates are based on the reports (and consequently the inferred signals) from the same set of agents. Therefore these estimates are identical.

The upper bound of Lemma 1 is tight in that it is easy to construct examples when agent i needs to form exactly $(\overline{L}^i)$th-order estimates. Consider the case $\mathrm{N}_i \cap \mathrm{N}_j = \{k_1, k_2, k_3, k_4\}$, with $G(k_1k_2) = G(k_2k_3) = 1$ and no other connections among these agents. Then $p_t^{ijk_2}$ is based on reports from i, j, k_1, k_2, k_3, while $p_t^{ijk_1k_2}$ is based on reports from i, j, k_1, k_2. Thus there exists a signal sequence such that agent i's forth-order estimates differ from her third-order estimates.

In special cases, agents need at most third-order estimates even though $\overline{L}^i$ is large. Suppose that all the subsets of fully connected agents in $\mathrm{N}_i \cap \mathrm{N}_j$ are mutually disjoint, and there are no other connections among these agents. The agents in each of the subsets think that they (and agent i, j) form a clique. Therefore agent i, j's estimates above or equal to the third-order are all the same within each subset. That is, $\mathbf{p}_t^{ijk} = \mathbf{p}_t^{jik} = \mathbf{p}_t^{kij} = \mathbf{p}_t^{ijk'}$ if $G(kk') = 1$; and $\mathbf{p}_t^{ijk}$ may differ from $\mathbf{p}_t^{ijk''}$ if $G(kk'') = 0$. But third-order estimates suffice for agent i and j. If the local network satisfies certain properties, the agents' estimates can be further simplified, involving only second-order estimates. Define the following property on agent i's local network (g_i, G_i).

Ni: (g_i, G_i) satisfies Ni if for every agent $j \in \mathrm{N}_i$, either (1) $\mathrm{N}_i \cap \mathrm{N}_j = \emptyset$, or (2) if there exists $k \in \mathrm{N}_i \cap \mathrm{N}_j$, there does not exist another agent k' such that $k' \in \mathrm{N}_i \cap \mathrm{N}_j$, but $G(kk') = 0$.

Notice that if i's local network satisfies Ni, g_{ij} is fully connected: Any agents connected to i and j must be connected themselves. In the example of Figure 2, because $G(kk') = 0$, agent i and j's local networks fail this property. But agent k and k''s local network satisfy Nk and Nk' respectively even though agent i and j have a common neighbor k who is not known to agent k'. This is because by assumption A2, agent k' *thinks* that every agent in $g_{ijk'}$ observes the same set of reports in forming estimates of another agent's estimates of the true state. Agents' estimates are particularly simple when their local networks satisfy Ni.

Corollary **1.** *For every agent $j \in \mathrm{N}_i$ and every period t,*

(1) Pairwise agreement: $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ji}$.

(2) Individual agreement: if (g_i, G_i) satisfies Ni, then $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik}$ and $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ijk} = \mathbf{p}_t^{ik...j...k'}$ for agents $k...j...k' \in g_{ij}/i$.

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

To begin with, part (1) above shows that even if agent i and j have different estimates in each 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 result does not rely on the properties of agent i or j's local networks. Rather, it depends on the fact that agent i and j's higher-order estimates involving each other all agree, up to the maximum $\overline{L}^{ij} + 2$ as given by Proposition 1. In the largest subset of $\mathrm{N}_i \cap \mathrm{N}_j$, the agents are fully connected with each other and i, j. Consequently i and j's highest-order estimates rely on the same reports and agree: $\mathbf{p}_t^{ijk...k'} = \mathbf{p}_t^{jik...k'}$ for any complete sequence of agents $k...k'$ in this subset. Similarly, there are a finite number of lower-order estimates, and we can show that they rely on the same reports for agent i and j. Since each of the lower-order estimates are derived from the next higher-order estimates, and they all agree, agent i and j's second-order estimates about each other also agree.

Next, if (g_i, G_i) satisfies Ni, from agent i's perspective, her estimates of two connected neighbors must be the same, that is, $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik}$. Recall that if Ni holds, g_{ij} is a clique and from agent i's perspective, everyone in g_{ij} has access to the same set of information. In fact, it is easy to show that $g_{ij} = g_{ik} = g_{ijk} = g_{ik...j...k'}$. Therefore agent i only needs to form second-order estimates of her neighbors' estimates.

Third, if (g_l, G_l) satisfies Nl for every agent $l \in g_i$, then the agents in each g_{ij} must agree on their second-order estimates, that is, $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$. Everyone's estimates of everyone else in g_{ij} are the same because they think that they have access to the same information. In particular, since $\mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$, agent i can directly use her inferred signal from k to replace j's inferred signal from k. Step 3 of our learning procedure then becomes much simpler, because there is no need to evaluate equation (6) and its higher-order counterparts. In particular, we have the following:

$$\alpha_t^{ji}(n) = \frac{p_t^i(n)}{p_t^{ij}(n)} \Bigg/ \sum_{n'} \frac{p_t^i(n')}{p_t^{ij}(n')}, \text{ and } p_{t+1}^{ij}(n) = \frac{p_t^j(n)\alpha_t^{ji}(n) \prod_{k \in \mathrm{N}_i \cap \mathrm{N}_j} \alpha_t^{ik}(n)}{\sum_{n'=1}^N p_t^j(n')\alpha_t^{ji}(n') \prod_{k \in \mathrm{N}_i \cap \mathrm{N}_j} \alpha_t^{ik}(n')}. \tag{7}$$

So far we described our learning procedure from a local network's perspective. We can extend this property to the whole network.

NG: Ni holds for every agent in the network (g, G).

Networks that satisfy NG are particularly well behaved because we can track the agents' learning in each local network easily. As the agents have the same estimates of their neighbors' estimates, each agent treats a new piece of information from outside the local network in an identical way. In particular, when NG holds, the inferred signals have the following property.

Corollary **2.** *Suppose the network (g, G) satisfies NG, then for every agent i,*

$$\alpha_{t+1}^{ij}(n) = \frac{\prod_{l\in((g_j/g_i)\cup j)} \alpha_t^{jl}(n)}{\sum_{n'} \prod_{l\in((g_j/g_i)\cup j)} \alpha_t^{jl}(n')}. \tag{8}$$

Corollary 2 shows that when NG holds, the inferred signal agent i learns from agent j in period t must be the combination of inferred signals j learns from nature and her neighbors who are not connected to i in the previous period by Bayes' rule. Two remarks are in order. First, Corollary 2 does not mean that agent i is able to learn all the other agents' signals correctly. If, say, a signal reaches multiple neighbors of an agent in a particular local network, she will infer there are multiple copies of that very signal according to equation (8). But if agent i receives each inferred signal only once, she will be able to form the correct Bayesian posteriors despite her lack of knowledge of (g, G). Second, her inferred signals will be wrong if there is asymmetric information, that is, when NG does not hold. For example, in the network in Figure 2, the signal agent k inferred from i can easily differ from the product of the inferred signals of i and j in the previous round.[14]

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 fixed signal sequence, information travels independently. Therefore we can analyze learning under each signal separately. More precisely, suppose the full sequence of signals is $X_T = \{x_t^i : \forall i, t < T\}$. Divide it into any two disjoint sets of signals, $X_T^a = \{x_t^{a,i} : \forall i, t < T\}$ and $X_T^b = \{x_t^{b,i} : \forall i, t < T\}$ such that $X_T = X_T^a \cup X_T^b$. Let $\mathbf{p}_t^i$ be agent i's estimates of the true state under information X_T, and $(\mathbf{p}_t^{a,i}, \mathbf{p}_t^{b,i})$ are her estimates under the corresponding information X_T^a and X_T^b. We say signals can be *decomposed* in a network if the agent's estimates under X_T is equal to the combination of her estimates under X_T^a and X_T^b

[14]See more details at Example 2 at the end of this section.

using Bayes' rule. That is, for all t, all i and all s_n.

$$p_t^i(n) = \frac{p_t^{a,i}(n)p_t^{b,i}(n)}{\sum_{n'=1}^N p_t^{a,i}(n')p_t^{b,i}(n')}; \tag{9}$$

$$p_t^{ij}(n) = \frac{p_t^{a,ij}(n)p_t^{b,ij}(n)}{\sum_{n'=1}^N p_t^{a,ij}(n')p_t^{b,ij}(n')}. \tag{10}$$

Moreover, for agent $k, ...j, ..., k' \in g_{ij}/i$,

$$p_t^{ik...j...k'}(n) = \frac{p_t^{a,ik...j...k'}(n)p_t^{b,ik...j...k'}(n)}{\sum_{n'=1}^N p_t^{a,ik...j...k'}(n')p_t^{b,ik...j...k'}(n')}. \tag{11}$$

PROPOSITION **2.** *Signals can be decomposed. If $X_T = X_T^a \cup X_T^b$, then equations (9), (10) and (11) hold for all i, t and s_n.*

Proposition 2 does not mean that agents can learn each signal correctly. Rather, it means that given a signal sequence, one signal travels, with its possible repetitions or distortions, independently from another. Consider two signals x_0^j and x_τ^k, which reaches N_i at time t in the form of y_t^j and y_t^k respectively. Then agent i's estimates are just the aggregation of y_t^j and y_t^k by Bayes' rule, the same as if we combine agent i's estimates at time t after receiving each signal individually. It implies that if the agents can form the correct Bayesian posterior beliefs under one signal, they can form the correct posterior beliefs under multiple signals.

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

EXAMPLE **2.** *In the network depicted in Figure 2, 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 signal and state distribution described in Example 1. The only informative signal is $x_0^{k'} = 1$.*

We only track $p_t^h(1)$, $h \in \{i, j, k', k\}$, agents' estimates that the true state is 1 in each period. At $t = 0$, the agents' first-order estimates, and thus their reports are $p_0^h(1) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. Agent k' gets the only signal $x_0^{k'} = 1$. At $t = 1$, the agents' first-order estimates are $p_1^{k'}(1) = \phi^1$ and $p_1^l(1) = \frac{1}{2}$, $l \in \{i, j, k\}$. All the higher-order estimates are $\frac{1}{2}$ because they are formed before agents exchange reports in period 1. After hearing k''s report, agent i and j observe that $p_1^{k'}(1) \neq \frac{1}{2}$.

At $t=2$, the agents' first-order estimates are $p_2^{k'}(1) = p_2^i(1) = p_2^j(1) = \phi^1, p_2^k(1) = \frac{1}{2}$. Agent i and j infer a signal $\alpha_2^{ik'}(1) = \alpha_2^{jk'}(1) = \phi^1$. Moreover, $\alpha_2^{ik'j}(1) = \alpha_2^{ik'...j}(1) = \frac{1}{2}$, and $\alpha_2^{ijk'}(1) = \alpha_2^{ij...k'}(1) = \phi^1$ by equation (3). Agent k' thinks that $p_2^{k'i}(1) = p_2^{k'j}(1) = p_2^{k'ij}(1) = ... = \phi^1$. Thus all the agents in the shared local network $g_{ik'}$ agree. But agent k observes that $p_2^i(1) = p_2^j(1) \neq \frac{1}{2}$.

At $t=3$, the agents' first-order estimates are $p_3^{k'}(1) = p_3^i(1) = p_3^j(1) = \phi^i$, and $p_3^k(1) = \frac{(\phi^1)^2}{(\phi^1)^2+(1-\phi^1)^2}$. All agents' estimates of every order in the shared local network of $g_{ik'}$ remain at ϕ^1, so they learn no new information. From now on, we focus on the shared local network g_{ik}.

Agent k infers two signals, $\alpha_3^{ki}(1) = \alpha_3^{kj}(1) = \phi^1$. Moreover, he thinks that agent i and j infers one new signal, $\alpha_3^{kij}(1) = \alpha_3^{kji}(1) = \phi^1$. Using Bayes' rule, all his estimates become $p_k^3(1) = p_3^{ki}(1) = p_3^{kj}(1) = p_3^{kij}(1) = ... = \frac{(\phi^1)^2}{(\phi^1)^2+(1-\phi^1)^2}$. Agent k also expects i and j to agree because they should incorporate each other's new signal. Agent i and j think that agent k will think there are two copies of such signals: $\alpha_3^{ikj}(1) = \alpha_3^{iki}(1) = \phi^1$, and thus $p_3^{ik}(1) = p_3^{jk}(1) = p_3^k(1)$. After observing their reports, however, agent k observes that $p_3^i(1) = p_3^j(1) \neq p_3^k(1)$.

At $t=4$, agent k', i and j's first-order estimates remain at ϕ^1, but $p_4^k(1) = \frac{1}{2}$. Agent k infers two opposite signals from i and j: $\alpha_4^{ki}(1) = \alpha_4^{kj}(1) = 1-\phi^1$. Agent i and j expect k to think they each received such an opposite signal. Incorporating these inferred signals, agent k thinks that $p_4^k(1) = p_4^{ki}(1) = p_4^{kj}(1) = ... = \frac{1}{2}$, but he notices after the reports that $p_4^i(1) = p_4^j(1) = \phi^1$ instead.

From then on, k', i, j's first-order estimates remain constant at ϕ^1, but agent k's estimates keep oscillating between $\frac{(\phi^1)^2}{(\phi^1)^2+(1-\phi^1)^2}$ in odd periods and $\frac{1}{2}$ in even periods. $\diamond$

4 Bayesian learning outcomes with local network information

How well can our agents learn given that they only know their local networks? This section compares the agents' estimates with the correct Bayesian posterior beliefs, which are based on all the signals available in network (g, G). We provide sufficient conditions for this to occur in two benchmark cases. We also show a necessary condition to illustrate why agents may not be able to learn correctly in general.

The outcomes from our learning procedure above are *Bayesian* if these exists some period t since which all agents' estimates agree with the Bayesian posterior beliefs and remain constant afterwards. This does not mean, however, agents make no mistakes during their learning process as shown in Example 3 below. Therefore we are also interested in a stronger notion of Bayesian learning, in which the agents' estimates agree with the Bayesian learning outcomes in every period *given* the travel paths of signals.[15] More precisely, let $d(ij)$ be the social distance between i and j, that is, the length of the shortest path between i and j, and let $d(ii) = 0$ for convenience. The diameter of the network (g, G) is $D = \max_{i,j\in g} d(ij)$. A signal (in possibly different forms) then takes at most D periods to reach every agent in (g, G). Recall that $X_t = \{x^i_\tau : \forall i \in g, \tau \leq t\}$, and we let $X^i_t = \{x^i_\tau : \forall \tau \leq t\}$ if $t \geq 0$ and $X^i_t = \emptyset$ otherwise. Because it takes $d(ij)$ periods for a signal to travel from j to i, we have

$$q^i_{t+1}(n) = \Pr(s_n | X^1_{t-d(i1)}, ..., X^L_{t-d(iL)}). \tag{12}$$

We call the learning procedure *strongly* Bayesian if $\mathbf{p}^i_t = \mathbf{q}^i_t$ for all i and t.[16]

It is easy to construct examples such that agents agree with the Bayesian posterior beliefs eventually, but the learning is not strongly Bayesian. Consider again the network in Figure 2.

Example **3** (Example 2 continued)**.** *Suppose that agent k' receives another signal $x^{k'}_1 = 0$ in addition to his initial signal $x^{k'}_0 = 1$. Everything else remains the same.*

Agent i and j learn signal $x^{k'}_1 = 0$ at $t = 2$ and incorporate it into their estimates at $t = 3$. Recall that before, at $t = 3$, the estimates in the local shared network g_{ik} is $p^i_3(1) = p^j_3(1) = \phi^1$ and $p^k_3(1) = \frac{(\phi^1)^2}{(\phi^1)^2+(1-\phi^1)^2}$. But agent k thinks that i and j has each received an offsetting signal, and that is why they didn't change their estimates. Thus k's inferred signals are two offsetting signals. In this case, it is just two signal 0 given the symmetry. With the new signal $x^{k'}_1 = 0$, $p^i_4(1) = p^j_4(1) = \frac{1}{2}$, and $p^k_4(1) = \frac{1}{2}$. All agents agree, and because the two signals offset each other, the correct Bayesian posteriors are the priors $\frac{1}{2}$. The learning here, however, is not strongly Bayesian because agent k first thinks that he receives two signal 1, which are fully correlated because they all come from $x^{k'}_0$ and then two signal 0, which again come from $x^{k'}_1$. They merely

[15] This may be particularly important when agents prefer to make an earlier decision because delay is costly.

[16] In Section 7, we show that if agents report the complete travel path of signals and each signal is tagged as in Acemoglu, Bimpikis and Ozdaglar (2014), then $\mathbf{q}^i_t$ is agent i's belief at period t.

cancel out by coincidence. $\diamond$

4.1 Information partition model guarantees strong Bayesian learning

Recall from Section 2.2 that the information partition model is a special case of our setup in which each of an agent's possible signals informs her of one element of her partition. In our setting, each agent starts with a symmetric prior, and her initial signal x_0^i informs her of the element $P^i(s^*)$, which contains the set of states she cannot distinguish from the true state s^*. Similarly, agent l's initial signal x_0^l informs him of $P^l(s^*)$. We assume, in line with the standard information partition model, that agents receive no further signals.[17]

Suppose every agent in (g, G) has information partitions $(\mathcal{P}^i_{i\in g}, S)$ as defined above. 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 under 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 of all the states. As a result, the agents in our model agree, and they reach the correct consensus sooner than in some of the existing models. To see this, consider the following example from Geanakoplos and Polemarchakis (1982).

Example **4.** *There are two agents* 1 *and* 2*. The state space* $S = \{s_1, s_2, ..., s_9\}$*, the ex ante probability of every state is* $1/9$*. Agent* 1*'s partition is* $\mathcal{P}^1 = \{(s_1, s_2, s_3), (s_4, s_5, s_6), (s_7, s_8, s_9)\}$*; agent* 2*'s partition is* $\mathcal{P}^2 = \{(s_1, s_2, s_3, s_4), (s_5, s_6, s_7, s_8), s_9\}$*. The true state is* s_1*. Thus* $P^1(s_1) = \{s_1, s_2, s_3\}$ *and* $P^2(s_1) = \{s_1, s_2, s_3, s_4\}$*. Let event* $A = \{s_3, s_4\}$*.*

Geanakoplos and Polemarchakis (1982) allow agents to announce and revise their posteriors of how likely A is true.[18] Agents know each other's information partitions. At $t = 1$, agent 1 knows A can only be true at $s = s_3$, while agent 2 thinks both can be true. So they will announce

[17] Our result can be easily extended to the case agents receive further signals resulting in finer information partitions.

[18] The posterior of an event given an agent's information is simply $\frac{\Pr(A\cap P^i(s_n))}{\Pr(P^i(s_n))}$ if $s_n \in S$.

1/3 and 1/2. But this is also consistent with the true state is s_4 and $P^1(s_1) = \{s_4, s_5, s_6\}$ and $P^2(s_1) = \{s_1, s_2, s_3, s_4\}$. At $t = 2$, agent 1 will announce 1/3 again because she hasn't learned any new information. Agent 2 notices that agent 1 does not change her posterior to 1. He realizes that the true state cannot be s_4, and thus changes his posterior to 1/3 as well.[19] From $t = 3$ onwards, they agree and the learning is over.

In our model, we don't need the partition to be common knowledge. By our procedure, the agents receive their initial signal at $t = 0$, which informs them $\mathcal{P}^i(s_1)$. At $t = 1$ agent 1 announces her posterior of the whole state distribution $\mathbf{p}_1^1 = \{1/3, 1/3, 1/3, 0, 0, 0, 0, 0, 0\}$. 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 y_0^{21} is distributed as $\{1/3, 1/3, 1/3, 0, 0, 0, 0, 0, 0\}$. Using updating rule (4), we can see that $\mathbf{p}_2^2 = \mathbf{p}_2^1 = \{1/3, 1/3, 1/3, 0, 0, 0, 0, 0, 0\}$. Because no one has any new information, the correct learning takes one, instead of two periods of communication. $\diamond$

In fact, we can show that under our learning procedure, not only agents in any network with information partitions will agree, their learning outcomes are strongly Bayesian.

PROPOSITION **3.** *If all agents in (g, G) have information partitions $(\mathcal{P}^i_{i\in g}, S)$, then their learning outcomes are strongly Bayesian and they reach consensus at $t = D + 1$.*

Proposition 3 holds for all networks because of the special feature of the information partition model: an agent's signal can be imprecise, but never wrong. Each piece of new information eliminates some possible states from an agent's original estimates based on her own initial signal. At $t = D + 1$, the agents' estimates are simply $1/|P^g(s^*)|$ if $s_n \in P^g(s^*) \equiv \cap\{P^l(s^*)\}_{l\in g}$, the intersection of all agents' element of partitions containing state s^*, and 0 otherwise.

To see why, observe that every time agent i infers an informative signal from agent j, the new information agent i learned is simply j's report: $\mathbf{\Delta p}_t^{ij} = \mathbf{p}_t^j$. That is, each unexpected report from a neighbor is a more precise new signal and should be taken into account. More importantly, repeated new signals don't matter. Suppose in some network, $\mathbf{\Delta p}_t^{ij} = \mathbf{\Delta p}_t^{ik}$, where j and k both learn the new signal from a common neighbor. Agent i's estimates are unaffected. She removes the

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

same set of states from what she thinks is possible as she would given one new signal, and assigns equal probabilities to the remaining ones.

The result that the agents' learning outcomes are Bayesian for all networks fails, however, when the information partition model is perturbed.[20] There is a discontinuity in that if the agents have any doubt about the mapping from the signals to their elements of partition, their estimates may depend on the network structure. In particular, the type and the number of inferred signals they learn from their neighbors. Let us revisit the network in Figure 2.

Example **5.** *A partition model with perturbation. The state space $S = \{s_1, s_2, s_3, s_4\}$, the ex ante probability of every state is $1/4$. Agent k''s partition is $\mathcal{P}^{k'} = \{(s_1, s_2), (s_3, s_4)\}$. All the other agents' partitions are simply S. The true state is s_1. Only agent k' sees an informative signal, which indicates the correct element with probability $1-\varepsilon$, and the wrong element with probability ε. Thus when $x_0^{k'} = (s_1, s_2)$, $\mathbf{p}_1^{k'} = (\frac{1-\varepsilon}{2}, \frac{1-\varepsilon}{2}, \frac{\varepsilon}{2}, \frac{\varepsilon}{2})$.*

The learning process is the same as the one in Example 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 learn k''s signal, so $\mathbf{p}_2^i = \mathbf{p}_2^j = \mathbf{p}_2^{k'} = \left(\frac{1-\varepsilon}{2}, \frac{1-\varepsilon}{2}, \frac{\varepsilon}{2}, \frac{\varepsilon}{2}\right)$ and $\mathbf{p}_2^k = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$. At $t = 3$, k learns two copies of k''s signal, so $\mathbf{p}_3^k = \left(\frac{(1-\varepsilon)^2}{2((1-\varepsilon)^2+\varepsilon^2)}, \frac{(1-\varepsilon)^2}{2((1-\varepsilon)^2+\varepsilon^2)}, \frac{\varepsilon^2}{2((1-\varepsilon)^2+\varepsilon^2)}, \frac{\varepsilon^2}{2((1-\varepsilon)^2+\varepsilon^2)}\right)$ while $\mathbf{p}_3^i$, $\mathbf{p}_3^j$ and $\mathbf{p}_3^{k'}$ remain the same from now on. At $t = 4$, k infers two offsetting 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 $\mathbf{p}_4^k$, and thus for any $\varepsilon > 0$, the agents never agree. $\diamond$

4.2 Social quilts guarantee strong Bayesian learning

In light of the result in Proposition 3, we now consider the more general information structure as described in Section 2.1.[21] Since signals can be decomposed by Proposition 2, we often start with the case of one initial signal only. It is easy to show that in complete networks such as the clique in Example 1, agents are able to learn this signal correctly. Whenever a signal reaches agent i in (g, G) in period t, all her neighbors learn this signal in period $t+1$. More importantly, every agent

[20] 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 almost 1, but inform her of some other element(s) with the complementary probability.

[21] Suppose for the remainder of this section that not all agents have information partitions.

thinks that everyone else in their local networks has learned it from agent i, and thus they will not double count this new information. In period $t+2$, everyone believes that there is one and only one copy of this signal and forms the correct estimates. We now generalize this intuition and show that the agents can learn correctly in certain networks even though they only know their local networks.

A *circle* is a path from an agent i to herself through distinct agents. That is, $c = \{i_0, ..., i_k\}$, $i_0 = i_k = i$, $G(i_l, i_{l+1}) = 1$, and $i_l \neq i_m$ for any l and m except $i_0 = i_k$. Moreover, c is a *simple circle* if it contains more than three agents, and $\mathrm{N}_l \cap c = \{i_{l-1}, i_{l+1}\}$ for any $i_l \in c$.[22] That is, other than the two adjacent neighbors, agent l has no other connections across a simple circle. Intuitively, a simple circle does not contain any smaller circles.

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

A social quilt is a tree-like union of cliques (completely connected subnetworks). We first provide a characterization for social quilts.

Lemma **1.** *Network (g, G) is a social quilt if and only if it satisfies NG and does not contain a simple circle.*

Recall that $T > 0$ is the period since which no signal arrives from nature, and s^* is the underlying true state. Then we have the following result.

Proposition **4.** *If (g, G) is a social quilt, then (1) the agents' learning outcomes are strongly Bayesian. Their estimates agree at period $T + D$, and remain constant afterwards. (2) If for any $s_{n'} \neq s_n$, there exists some signal x_m such that $\phi^i_{mn} \neq \phi^i_{mn'}$, then as $T \to \infty$, $p^i_T(s^*) \to 1$ for all i.*

Suppose there is only one signal, x^i_0. If we can show that each agent j learns the signal at, and only at, period $d(ij) + 1$, then the learning is strongly Bayesian. That is, this signal reaches each agent exactly once through the shortest path between i and j for all j. This is because in any social quilt, there is a unique shortest path between any pair of agents. Suppose to the contrary that there are two different shortest paths between agent i and j. If these two shortest paths do

[22] In particular, $i_{-1} = i_{k-1}$ and $i_{k+1} = i_1$.

not overlap, that is, there does not exist an agent who belongs to both paths. Then there must be a circle going from i to j through one path and from j back to i through a different one. By definition, since they belong to the same circle, i and j must be connected, or $d(ij) = 1$. This is a contradiction. The case when these two shortest paths partially overlap is analogous. As a consequence, the signal from agent i reaches agent j through this unique shortest path. Next, we want to show once agent j learns the signal, she will not see this signal again. Suppose the shortest path between i and j is $\{i, ..., k, j\}$, and agent l is a neighbor of agent j who learns the signal from j. By Corollary 2, $G(kl) = 0$ so the signal passes from k to j then to l. Then the shortest path from agent i to agent l must go through j, that is $d(il) = d(ij) + 1$. Otherwise, there must have been a circle involving agent k, j and l with $G(kl) = 0$, which is impossible in a social quilt. As the signal is inferred only by agents further and further away from i, it cannot reach agent j again.

Intuitively, Proposition 4 holds because a social quilt is a tree-like union of cliques, and the signal travels from one clique to another through the tree. In other words, agent k in the same clique as i learns the signal and passes it on to others in the same clique with k, and so on until the signal reaches the "terminal" cliques of the tree. Everyone in the terminal cliques learn the signal and do not have neighbors to pass on the information. Because a social quilt does not have simple circles, the signal does not come back to any clique. Similarly, with multiple signals, agent i's estimates at $t + 1$ include signals observed by each agent j from period 0 to period $t - d(ij)$. So the learning outcome is strongly Bayesian. If the last signal arrives at period $T - 1$, it takes at most $D + 1$ periods to reach all agents. Therefore all learning stop by the end of period $T + D$ with everyone forming the correct estimates.

As 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 portion of these are informative. Conditional on these signals, the ratio of the posterior belief of any state over that of the true state is just the ratio of the priors times the likelihood ratio of observing these signals given each state. The latter itself is a random variable, and by the (weak) Law of Large Numbers, this ratio approaches zero. Therefore as T becomes arbitrarily large, the agents' estimates, which are the Bayesian posterior beliefs based on these signals, must put a probability arbitrarily close to one on the true state.

4.3 When Bayesian learning is impossible

The next question is how well the agents can learn in networks that are not social quilts. Recall from Proposition 4 above that even with only one informative signal, all agents need to learn there is one and only one copy of this signal. In this section we highlight two impediments to Bayesian learning in our model. First, because each agent only knows her local network, the same signal may travel around a circle repeatedly and appear as new information. Second, when NG does not hold, agents have asymmetric information about the local network. Therefore even neighbors in the same local network may never form a consensus, which is necessary for Bayesian learning.

We begin by considering the case when network (g, G) satisfies NG, but it contains at least one simple circle. Suppose that there is only one signal $x_0^i = x_m$. Let $S(x_m)$ be the set of states that maximize μ_{nm}^i, the probability of state being s_n conditional on agent i observing the signal x_m.[23]

Proposition 5. *Suppose that (g, G) satisfies NG and has at least one simple circle. If $x_0^i = x_m$ is the only signal, then the agents' estimates converge to $p_\infty^j(n) = \frac{1}{\#S(x_m)}$ if $s_n \in S(x_m)$ and to 0 otherwise.*

Since (g, G) satisfies NG, we can show that each agent thinks there are one or multiple identical copies of signal x_m at any given time. The difference from Proposition 4 is that inferring new signals does not stop when there is at least one simple circle. The only informative signal travels along the circle repeatedly, and each time it reaches an agent, the agent believes there are at least one more copies of the same signal (the count may be higher if there are multiple simple circles). Because all agents are path-connected, the repeated information will reach agents outside any circles within D periods as well. In the limit, all agents' estimates are equivalent to the Bayesian posterior belief of observing an infinite number of x_m. Consequently, agents believe only in the state(s) most likely given signal x_m.

In light of Proposition 5, a natural hypothesis is with simple circles, the agents' estimates converge to the state(s) that is ex ante most likely given all the signals. Unfortunately, despite a

[23]There may be multiple such states if $\phi_{mn}^i = \phi_{mn'}^i > \phi_{mn''}^i$ for all $n'' \neq n, n'$.

convergence of estimates under each separate signal in the network with simple circles, the overall estimate convergence may fail, even with a very simple network structure. We first show that in the very special case of one simple circle with asymmetric signals, this hypothesis is true. Let $\tilde{s} = \arg\max_{s_n} \Pr(s_n|X_T)$, where $\Pr(s_n|X_T)$ is the Bayesian posterior using all signals.

COROLLARY **3.** *If NG holds and the network contains exactly one simple circle, all agents' estimates converge to $p^i_\infty(\tilde{s}) = 1$ when $\tilde{s}$ is singleton.*

This result holds because signals travel at a linear speed in the case of one simple circle. For example, consider the circle $c = \{i_0, ..., i_k\}$ and one signal $x_0^{i_0}$. At $t = k$, agent i_0 learns two more "new" copies of $x_0^{i_0}$ from her two neighbors in the circle, and another two "new" copies after every k periods. After a sufficiently long period of time, the stronger signal dominates even if it arrives late to the circle. Once a signal dominates in the circle, all agents will receive (in net) more and more copies of this signal even if the agents are themselves not part of the circle. In the end, all agree to the state that is most likely given the signals.

If, however, there are multiple circles, the agents may never form a consensus, let alone form the correct Bayesian posterior beliefs. This is illustrated by the following example.

EXAMPLE **6.** *Exponential information travel and non-convergence of estimates with two circles. Seven agents are connected in two circles as in Figure 3. Continue with the binary state setting with two asymmetric pairs of signals, one observed by agent 3 at $t = 0$, $\Pr(x_0^3 = 0|s = 0) = \Pr(x_0^3 = 1|s = 1) = \phi^3 > \frac{1}{2}$, and another one observed by agent 3 at $t = 2$, $\Pr(x_2^3 = 0|s = 0) = \Pr(x_2^3 = 1|s = 1) = \frac{(\phi^3)^2}{(\phi^3)^2+(1-\phi^3)^2}$. So $x_2^3 = 1$ is more precise than $x_0^3 = 0$.*[24]

Let us first focus on x_0^3 and agent 3. At $t = 0$, agent 3 learns one signal x_0^3, and at $t = 4$, she infers 4 more copies of x_0^3, one from each neighbor. Note that at $t = 5$, each of 3's neighbors also infers 3 copies of x_0^3, because they know they share one copy to agent 3 and the other 3 copies must be new information. At $t = 8$, agent 3 learns $4 \cdot 3$ copies of x_0^3. Iteratively at $t = 4(k + 1)$, agent 3 infers $4 \cdot 3^k$ copies of x_0^3. Signal x_2^3 travels along the network in the same way except a lag of

[24] That is agent 3 can observe 4 possible signals, $\{0L, 1L\}$ with a lower precision ϕ^3 and $\{0H, 1H\}$ with a higher precision $\frac{(\phi^3)^2}{(\phi^3)^2+(1-\phi^3)^2}$.

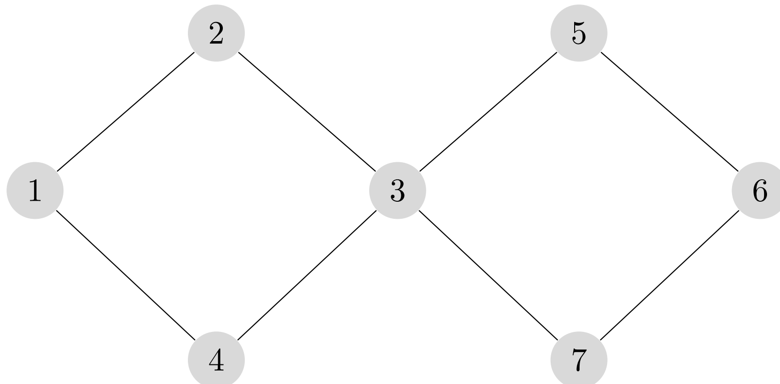

Figure 3: A network with 7 agents in two circles.

two periods. So at $t = 4(k+1)$, agent 3's information set includes $1 + \sum_{l=0}^{k} 4 \cdot 3^l$ copies of x_0^3, and $1 + \sum_{l=0}^{k-1} 4 \cdot 3^l$ copies of x_2^3. Note that when $k \geq 2$,

$$1 + \sum_{l=0}^{k-1} 4 \cdot 3^l < 2 \cdot 4 \cdot 3^{k-1} = \frac{1}{2}(4 \cdot 3^{k-1} + 4 \cdot 3^k) < \frac{1}{2}(1 + \sum_{l=0}^{k} 4 \cdot 3^l).$$

That is the number of x_2^3 is less than half of the number of x_0^3. As one x_2^3 can cancel two x_0^3, that is, $\Pr(s = 1|x_0^3, x_0^3, x_2^3) = \frac{1}{2}$. Agent 3's estimates must put a higher probability on 0, so $p^3_{4(k+1)+1}(1) < \frac{1}{2}$. At $t = 4(k+1) + 2$, agent 3's information set includes the same number of x_0^3 and x_2^3. As x_2^3 is more precise, agent 3's estimates must put a higher probability on 1, so $p^3_{4(k+1)+3}(1) > \frac{1}{2}$. Thus, agent 3's estimates keep oscillating. $\diamond$

In addition, the sequencing of signals may matter in agent's estimates when networks have simple circles. As shown in the previous example, if there are more than one simple circles, information repetition is exponential, giving earlier signals a more important role. The following example shows that the earlier signals may grow so fast that agents can not be persuaded by an arbitrarily large number of signals to the opposite.

Example **7.** *Failure of the Law of Large Number. Consider eight agents connected in a cube, as in Figure 4. Continue with the binary state and binary 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*. Suppose that each agent observes* *a signal of* $x_0^i = 0$ *at* $t = 0$*; and each agent observes a signal of* $x_t^i = 1$ *for* $t \in \{1, ..., T-1\}$.[25] *Then the agents believe the state is* 0 *as* T *approaches infinity:* $lim_{T\to\infty} \Pr(s = 0|X_T) = 1$.

To see why the agents believe the state is 0 even when they receive so many opposing signals from $t = 1$ onward, note that in the beginning of $t = 1$, each agent reports $p_1^i(0) = \phi$. Thus each agent infers three signals of 0 from their neighbors and learns one signal of 1 from the nature. At $t = 2$, each agent's 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 knows that each of her neighbors learns a signal of 0 from her in addition to their own signal of 0. Therefore each agent infers again three signals of 0, one from each neighbor, plus one signal of 1 from the nature. The learning is exactly the same as in period 1, and remains the same in all later periods up to T. Because all agents think they are learning more and more signals of 0, they believe the state is 0 in the limit. $\diamond$

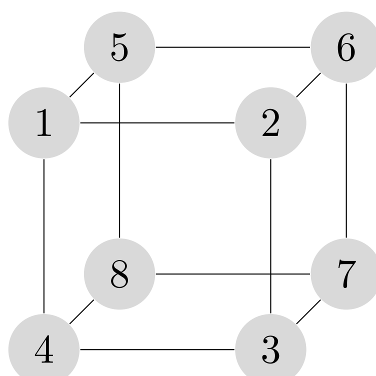

Figure 4: A cube with 8 agents

We now turn our attention toward networks that fail to satisfy NG, in which agents may have asymmetric information even within their local networks. Recall the network in Figure 2: agents i, j, k, k' such that $k, k' \in \mathrm{N_i} \cap \mathrm{N_j}$ and $G(kk') = 0$. While agents k, k' do not think their common friends i, j have common friends that they do not know, i and j know differently. In particular, k and k' cannot tell whether the sources of i and j's information are the same or not, and thus may make mistakes in forming estimates. Let $\Omega = \{k' : \exists\ i, j, k, s.t.\ G(ij) = 1, k, k' \in \mathrm{N}_i \cap \mathrm{N}_j, \text{and } G(kk') = 0\}$

[25]It is an extremely special case that all signals from period 1 to period $T-1$ are 1, however, we will show that agents' estimates put probability close to 1 on state 0. Then in all other cases, agents' estimates put a higher probability on state 0.

be the set of agents whose presence makes that their neighbors' local networks do not satisfy Ni. This set is nonempty if NG does not hold.

Proposition **6.** *Suppose (g, G) doesn't satisfy NG. If any agent $k' \in \Omega$ receives $x_0^{k'} = x_m$, the only informative signal, the agents' estimates cannot agree with the Bayesian posterior beliefs.*

Intuitively, when (g, G) does not satisfy NG, there must exist agents i, j, k, k' such that $k, k' \in \mathrm{N_i} \cap \mathrm{N_j}$ and $G(kk') = 0$. Because kk' are not connected, if an informative signal reaches agent k' (say $x_0^{k'}$), agent k's estimates keep oscillating in the fashion described in Example 2. An extra complication appears when both this type of $\{i, j, k, k'\}$ and simple circles coexist in the network. Then even if there is only one informative signal, it may reach the agents again through circles. We can show, however, these "new" information either reach agents i and j, which are properly incorporated by both agents k and k'; or they make other agents' estimates oscillate as well.

We are now ready to show that social quilts are almost necessary for Bayesian learning in our model. For a small number of signals, it is easy to construct examples where even our agents with local knowledge can form the correct Bayesian posterior beliefs in every period. For instance, consider network (g, G) in which agent 1, 2 and 3 are connected in a triad, and agent 3 is connected to a large set of agents in an unspecified subnetwork. Using our symmetric binary example again, and suppose that agent 1 receives signal $x_0^1 = 1$ and agent 2 receives signal $x_0^2 = 0$. No other agents receive any informative signals. Then clearly, at $t = 2$, the signals cancel and agents $1, 2, 3$ believe the state is 1 with probability $\frac{1}{2}$, which is both the prior and the correct posterior beliefs given the signals. No other agents beliefs change from the priors, and every agent's estimates are correct. This example is highly artificial. We proceed to show that in any network that is not a social quilt (and not all signals are partitional), when the society receives a finite number of signals, the probability that the agents' estimates agree with the Bayesian posterior beliefs is bounded away from one.

Proposition **7.** *In any network (g, G) that is not a social quilt, the probability that the agents' period-T estimates agree with Bayesian posterior beliefs is bounded away from one for any finite T.*

The examples of Bayesian learning in networks that are not social quilts rely on the fact that there are a small number of offsetting signals arriving in a particular sequence. Proposition 7 relies on the fact that when T is finite, the probability of agents receiving these special sequences of signals as a fraction of total possible sequences of signals is bounded away from one. Therefore the agents will not be able to learn correctly. These set of results suggest that in networks that are not social quilts, agents may become overconfident in the wrong state of the world, or they change their estimates constantly as time goes on.[26]

5 Continuous states and signals

In this section we introduce the widely used normal-linear model, in which the agents' signals are independent conditional on the state, with normally distributed noise. We will show that all the main results of our learning procedure hold. The simplicity of this model also allows us to compare conveniently with the myopic learning model in the network literature. It further illustrates the improvement of the agents' learning due to their ability to account for information repetition in their local network.

5.1 The normal-linear model

Consider the model in which the state is drawn from a uniform distribution over the real line: $s \sim U(-\infty, \infty)$. The signal for each agent i is normally distributed such that $x_t^i = s + \varepsilon_t^i$, where $\varepsilon_t^i \sim N(0, \zeta_t^i)$. The noise is independently distributed across agents and time, and $\zeta_t^i \geq 0$ is the precision of agent i's signal in period t (agent i receives an uninformative signal if $\zeta_t^i = 0$). As before, the state distribution is common knowledge, and it is common knowledge that the signals are normally distributed and unbiased, but the precision $\{\zeta_t^i\}_{t\geq 0}$ is each agent i's private information.[27]

The learning procedure is very similar to the main model, so we only outline the differences below. At $t = 0$, each agent simultaneously reports her initial estimates, $\mathbf{p}_0^i = (s_0^i, \pi_0^i)$ where s_0^i is her prior belief of the state and π_0^i is the precision of her prior belief. Given the above (improper) uniform priors, $\mathbf{p}_0^i = (0, 0)$ for all agents. By the end of $t = 0$, agent i learns her signal x_0^i as well as her neighbors' reports, $I_0^i = \{\mathbf{p}_0^h : h \in g_i\} \cup x_0^i$. At the beginning of $t = 1$, agent i updates her

[26]The lack of convergence is partly due to the fact that in the main model, agents treat each piece of information equally and information never decays. We consider the case of discounting in Section 6.2.

[27]As in the main model, allowing the agents to have asymmetric priors does not affect our results.

estimates to $\mathbf{p}_1^i = (s_1^i, \pi_1^i)$, where

$$s_1^i = \mathrm{E}\left[s|I_0^i\right] \text{ and } \pi_1^i = \zeta_0^i + \sum_{j \in g_i} \pi_0^j.$$

Agent i also forms estimates about her neighbors' estimates $\mathbf{p}_1^{ij} = (s_1^{ij}, \pi_1^{ij})$, using the information i thinks j observes,

$$s_1^{ij} = \mathrm{E}\left[s|I_0^{ij}\right] \text{ and } \pi_1^{ij} = \sum_{k \in g_{ij}} \pi_0^k, \text{ where } I_0^{ij} = \{\mathbf{p}_0^k : k \in g_{ij}\},$$

for all $j \in \mathrm{N}_i$. Given the uninformative priors, $\mathbf{p}_1^{ij} = \mathbf{p}_1^{ji} = (0, 0)$. Similarly, agent i forms estimates of agent j's estimates of agent k's estimates when both $j, k \in g_i$, using information $I_0^{ijk} = \{\mathbf{p}_0^h : h \in g_{ijk}\}$. We can define estimates $\mathbf{p}_1^{ijkh}, \mathbf{p}_1^{ijkjh}, \ldots$, in a similar fashion. Initially, all these estimates are the uninformative priors. In each period t, agents update their estimates in a similar way.

Step 1: Identify new information. From agent i's perspective, for every $j \in \mathrm{N}_i$, j has no new information if $\pi_t^j = \pi_t^{ij}$.[28] If $\pi_t^j > \pi_t^{ij}$, j must have had new information, i.e., an inferred signal following a normal distribution with mean y_{t-1}^{ij} and precision ν_{t-1}^{ij}, that i does not know. In this case, agent j must have formed a new estimate of the state by combining information known to i, (s_t^{ij}, π_t^{ij}), with the new information $(y_{t-1}^{ij}, \nu_{t-1}^{ij})$. Because of the nice property of normal distributions, a combination of two normal distributions is still a normal distribution. After incorporating this new information, agent j's updated precision and the mean of her estimate are given by:

$$\pi_t^j = \pi_t^{ij} + \nu_{t-1}^{ij}, \text{ and } s_t^j = \frac{s_t^{ij}\pi_t^{ij} + y_{t-1}^{ij}\nu_{t-1}^{ij}}{\pi_t^j}.$$

Knowing this, agent i can infer the new information as follows:

$$\triangle\pi_t^{ij} = \pi_t^j - \pi_t^{ij}, \text{ and } \triangle s_t^{ij} = \frac{s_t^j\pi_t^j - s_t^{ij}\pi_t^{ij}}{\triangle\pi_t^{ij}}.$$

[28]Lemma 2 below shows that when NG holds, if $\pi_t^j = \pi_t^{ij}$, it must be true that $s_t^j = s_t^{ij}$.

To agent i, the new inferred signal is characterized by $\mathbf{\Delta p}_t^{ij} = (\triangle s_t^{ij}, \triangle\pi_t^{ij})$.

Agent i uses a similar way to identify the new information she thinks j learns from their common friends k, $\mathbf{\Delta p}_t^{ijk} = (\triangle s_t^{ijk}, \triangle\pi_t^{ijk})$ where

$$\triangle\pi_t^{ijk} = \pi_t^k - \pi_t^{ijk}, \text{ and } \triangle s_t^{ijk} = \frac{s_t^k \pi_t^k - s_t^{ijk}\pi_t^{ijk}}{\triangle\pi_t^{ijk}}.$$

The same for $\mathbf{\Delta p}_t^{iji}$ and all high-order estimates.

Step 2: Update own estimates. Agent i uses any new information from nature, (x_t^i, ζ_t^i), as well as those learned from her neighbors, $\mathbf{\Delta p}_t^{ij}$ for all $j \in \mathrm{N}_i$, to update her estimate such that

$$\pi_{t+1}^i = \pi_t^i + \zeta_t^i + \sum_{j\in\mathrm{N}_i} \triangle\pi_t^{ij},$$

and

$$s_{t+1}^i = \frac{s_t^i\pi_t^i + x_t^i\zeta_t^i + \sum_{j\in\mathrm{N}_i} \triangle s_t^{ij}\triangle\pi_t^{ij}}{\pi_{t+1}^i}.$$

Step 3: Update estimates of neighbors' signals. When forming an estimate of neighbor j's posterior, agent i starts with agent j's latest estimate $\mathbf{p}_t^j$ and incorporates the new information i thinks j learns from i, $\mathbf{\Delta p}_t^{iji}$, and those from their common friends, $\mathbf{\Delta p}_t^{ijk}$ for all $k \in \mathrm{N}_i \cap \mathrm{N}_j$.

$$\pi_{t+1}^{ij} = \pi_t^j + \triangle\pi_t^{iji} + \sum_{k\in\mathrm{N}_i\cap\mathrm{N}_j} \triangle\pi_t^{ijk},$$

and

$$s_{t+1}^{ij} = \frac{s_t^j\pi_t^j + \triangle s_t^{iji}\triangle\pi_t^{iji} + \sum_{k\in\mathrm{N}_i\cap\mathrm{N}_j} \triangle s_t^{ijk}\triangle\pi_t^{ijk}}{\pi_{t+1}^{ij}}.$$

Each of agent i's lower-order estimates are thus formed iteratively from her next higher-order estimates. And this proceeds iteratively. $\|$

Very similar proofs to the main model can show that Corollary 1 holds. Thus, when NG holds, agents agree on their second-order estimates of others in the same local network. And step 3 of learning procedure can be simplified as follows.

$$\pi_{t+1}^{ij} = \pi_t^j + \triangle\pi_t^{ji} + \sum_{k\in \mathrm{N}_i\cap \mathrm{N}_j} \triangle\pi_t^{jk},$$

and

$$s_{t+1}^{ij} = \frac{s_t^j\pi_t^j + \triangle s_t^{ji}\triangle\pi_t^{ji} + \sum_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\triangle s_t^{jk}\triangle\pi_t^{jk}}{\pi_{t+1}^{ij}}.$$

Furthermore, our procedure is well defined under our assumptions. In step 1 of the learning procedure, we focus on i's inferred signal when $\pi_t^j \geq \pi_t^{ij}$. Because of the updating process, new information must increase an agent's precision. And the following lemma shows that when NG holds, $\pi_t^j \geq \pi_t^{ij}$ is always true.

LEMMA **2.** *Consider* (g, G) *satisfies NG. It is true that* $\triangle\pi_t^{ij} \geq 0$ *for all* i*,* $j \in \mathrm{N}_i$ *and* t*; when* $s_t^j \neq s_t^{ij}$*, it is true that* $\triangle\pi_t^{ij} > 0$*.*

Intuitively, π_t^{ij} is i's estimate of j's estimate's precision using all the information that i thinks j observes. When NG holds, there does not exist some $k \in \mathrm{N}_i \cap \mathrm{N_j}$ who shares a common friend with j that is not connected to i. So for all the new information i thinks j observes, j indeed treats them as new information, then π_t^j must be weakly higher than π_t^{ij}. Very similar proofs to the main model can show that Proposition 2, 4 and 5 hold in this normal linear signal setting.

What if NG does not hold, and $\pi_t^j < \pi_t^{ij}$? That is, agent i expects agent j to have a higher precision than j actually reports. There are two possible explanations: One logical interpretation is that agent i has over counted the number of signals j has because agent i does not know N_j; or agent j realized she over counted the number of signals she gets from her neighbors so she reduced her precision to correct her mistake. In the former case, a reasonable learning procedure should then allow agent i to update her estimate about j's network, however A2 assumes that our agents behaves as if (g_i, G_i) is the only network. To focus on agents' learning of information, we maintain the assumption that (g, G) satisfies NG in the normal-linear model, and we discuss agents' learning of the network in section 7.4 when we relax assumptions A1 and A2.

5.2 Improvement over myopic learning

Another common way of modeling learning in networks is the myopic learning model in the spirit of DeGroot 1974, DeMarzo, Vayanos and Zwiebel 2003, among others. The essence of myopic learning is that agents update estimates by repeatedly taking weighted averages of their neighbors' estimates, without accounting for any possible repetition of the information. In particular, it allows for two types of repetitions: those arising due to agents listening over time to the same set of neighbors in her local network; and those arising due to common, overlapping sources farther away in the network. In this subsection, we compare these two learning procedures, maintaining the key feature of myopic learning in that agents treat their neighbors' reports as new, independent information. For simplicity and ease of comparison, we use the normal linear setting similar to that in DeMarzo, Vayanos and Zwiebel (2003), and assume that agents only observe signals at the beginning t=0.

More specifically, we modify the myopic learning procedure such that the two learning procedures share the same message space. Namely, each agent reports both her estimate of the state and the precision at the beginning of each period. The crucial feature is that each agent treats that a neighbor j's report as a new signal in each period. This implies that the agents don't remember how they made use of their own signals or their neighbors' reports in the past.[29] Each agent then updates by taking a weighted average of her own reports and her neighbors. Instead of the fixed and subjective weights used in the standard DeGroot learning model, the weights here are the precisions an agent attaches to her neighbors' "new" signals.

Recall that there are L agents in network (g, G). As in the setup above, at $t = 0$, each agent i reports her initial estimate of the state $\mathbf{p}_0^i = (s_0^i, \pi_0^i) = (0, 0)$. Then the agents receive their initial signals (x_0^i, ζ_0^i). At $t = 1$, the agents report $\mathbf{p}_1^i = (s_1^i, \pi_1^i) = (x_0^i, \zeta_0^i)$. Let T_t^{ij} be the weight agent i assigns to agent j's information at time t; and T_t^{ii} be the weight agent i attaches to her own information. Also, let $\mathbf{T_t}$ be the corresponding $L \times L$ matrix where each entry is nonnegative

[29] The amounts to updating using only Step 2 of our learning procedure in both the normal linear model and our main model. Consider a two-agent example where agent 1 and 2 respectively receive a signal x_0^1 and x_0^2. At $t = 1$, their reports are just their signal. At $t = 2$, their estimates are the correct Bayesian estimates conditional on x_0^1 and x_0^2. At $t = 3$, however, agent 1 believes agent 2 received a new signal (the composite of x_0^1 and x_0^2) in period 2 and updates again, and vice versa for agent 2. Since they have the same estimates at $t = 2$ and update using the same "signals" repeatedly, they will reach a consensus based on receiving infinite numbers of x_0^1 and x_0^2 in the end.

and $T_t^{kl} = 0$ means that agents k, l are not directly connected. Notice that since our network is undirected, if $T_t^{ij} > 0, T_t^{ji} > 0$. Bayesian updating for agent i entails:

$$s_2^i = \sum_{j \in g_i} T_1^{ij} s_1^j \quad \text{and} \quad \pi_2^i = \sum_{j \in g_i} \pi_1^j, \text{ where } T_1^{ij} = \frac{\pi_1^j}{\pi_2^i}.$$

Clearly, $\sum_{j \in g_i} T_1^{ij} = 1$. Observe that the weight in the first period is just a function of the precision of the agents' signals. This process continues such that at period $t+1$,

$$s_{m,t+1}^i = \sum_{j \in g_i} T_t^{ij} s_{m,t}^j \quad \text{and} \quad \pi_{m,t+1}^i = \sum_{j \in g_i} \pi_{m,t}^j, \text{ where } T_t^{ij} = \frac{\pi_{m,t}^j}{\pi_{m,t+1}^i}.$$

We use subscript m here to denote myopic learning.

We first show a convergence result similar to that of DeMarzo, Vayanos and Zwiebel (2003).

Proposition **8.** *With myopic learning, agents' estimates of the state converge to a consensus while their estimates of the precision go to infinity.*

The convergence of estimates does not directly follow from the standard DeGroot model because our myopic learning model allows for time-varying weights: T_t^{ij} depends on time t and the precision of each agent's signal. To show convergence, we first need to show that the weights between two connected agents are bounded away from zero. Then, in each period an agent's estimate is a weighted average of her neighbors. If the weights are not too small, then the agents' estimates must move closer over time.

Lemma **3.** *There exists some $\underline{\omega} > 0$ such that $T_t^{ij} \geq \underline{\omega}$ for all i, j s.t. $G(ij) = 1$ and all t.*

In the myopic learning, one agent's precision at period $t+1$ is the sum of her precision and all her neighbors' precision at period t. So one agent's precision grows at a similar rate as her neighbors' precision. Since T_t^{ij} depends on the ratio of j's precision in i's local network, it is bounded above zero. Proposition 8 then follows from the theorem proved by Lorenz (2005).

Theorem **1.** ***Lorenz [2005]*** *Suppose that for all t, the weight matrix $\mathbf{T_t}$ satisfies the following conditions:*

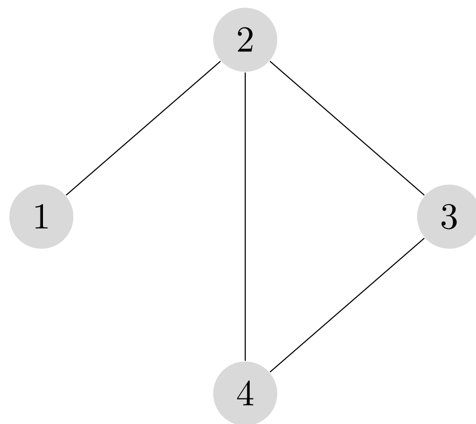

Figure 5: A network that satisfies NG

- $T_t^{ii} > 0$ *for all* i*.*
- $T_t^{ij} > 0$ *if and only if* $T_t^{ji} > 0$ *for all* i *and* j*.*
- *There exists* $\omega > 0$ *such that* $T_t^{ij} \geq \omega$ *if* $T_t^{ij} > 0$ *for all* i *and* j*.*

Then the society can be partitioned into sets of agents such that each group of agents reaches a consensus, and any two agents who place weight on each other infinitely often are in the same group.

Proposition 8 implies that under myopic learning, the agents' estimates never converge to the Bayesian learning outcome regardless of the network structure, because the estimate of the precision goes to infinity. In contrast, Proposition 4 shows that our learning procedure generates Bayesian posterior if the network is a social quilt.

Example **8.** *Comparison with myopic learning in a social quilt.*

Consider 4 agents connected in Figure 5. Suppose at $t = 0$, agent 1 observes a signal $(x_0^1, \zeta_0^1) = (1, 1)$ and agent 2 observes a signal $(x_0^2, \zeta_0^2) = (0, 1)$. If agents use our learning procedure, by $t = 3$, all of them hold the correct Bayesian estimate $(1/2, 2)$ and the learning stops.

Now suppose agents use myopic learning. Their estimates are updated as follows.

- At $t = 1$, $\mathbf{p}_1 = ((1, 1), (0, 1), (0, 0), (0, 0))$

- At $t = 2$, $\mathbf{p}_2 = ((1/2, 2), (1/2, 2), (0, 1), (0, 1))$
- At $t = 3$, $\mathbf{p}_3 = ((1/2, 4), (1/3, 6), (1/4, 4), (1/4, 4))$
- At $t = 4$, $\mathbf{p}_4 = ((2/5, 10), (1/3, 18), (2/7, 14), (2/7, 14))$

In the limit, their estimate of the state must converge to some value smaller than 2/5, while the Bayesian estimate is 1/2. And their estimate of precision goes to infinity. So in the limit, both the estimates of the state and the precision are incorrect. $\diamond$

Furthermore, the following result shows that even when agents are infinitely confident in both models when the network is not a social quilt, the over-confidence is more severe under myopic learning.

PROPOSITION **9.** *For all i, $\lim_{t\to\infty}(\pi^i_{m,t} - \pi^i_t) = \infty$, so agents are a lot more over-confident in a myopic learning.*

Consider 4 agents connected in a circle. Suppose at $t = 0$, agent 1 observes a signal $(x^1_0, \zeta^1_0) = (1, 1)$ and it is the only signal. Then agents' estimates of the state remain at 1, so we just keep track of the precision increases. If agents use our learning procedure, their precision starts with $\pi_1 = (1, 0, 0, 0)$, then 2 and 4 learn the signal from 1 with $\pi_2 = (1, 1, 0, 1)$, then 3 learns 2 signals from 2 and 4 with $\pi_3 = (1, 1, 2, 1)$, then 2 and 4 learn one signal from 3 with $\pi_4 = (1, 2, 2, 2)$, and then 1 learns 2 signals from 2 and 4 with $\pi_5 = (3, 2, 2, 2)$. At time $t = 4\tau+1$, $\pi_t = (2\tau+1, 2\tau, 2\tau, 2\tau)$ with integer $\tau \geq 0$. While if agents use myopic learning, their precision starts with $\pi_{m,1} = (1, 0, 0, 0)$, then 2 and 4 learn the signal from 1 with $\pi_{m,2} = (1, 1, 0, 1)$, then they treat all previous signals independent so $\pi_{m,3} = (3, 2, 2, 2)$. Note that $\pi^i_{m,3} \geq 2$ for all i, and each period each agent summarizes the precisions of her own and her neighbors' precision in the previous period, thus $\pi^i_{m,4} \geq 2 \cdot 3$, and in general $\pi^i_{m,t} \geq 2 \cdot 3^{t-3}$. So the precision grows exponentially. In this example, Proposition 9 can be strengthened to $\lim_{t\to\infty} \frac{\pi^i_t}{\pi^i_{m,t}} = 0$.

6 When all information is not equal

In the main model, we implicitly assume that agents treat all information equally regardless of the sources and the arrival times. Yet for various reasons, agents may weigh different information differently. First, agents may trust some of their neighbors more than others, perhaps due to difference in perceived accuracy or reputation. Second, agents may discount the new information as time goes on, perhaps due to loss of information during imperfect information diffusion or due to the fact that as they become very confident about their estimates of the true state, they pay less attention to the new information. We now extend our main model by letting the agents put different weights on their inferred signals, which is applicable to all the aforementioned settings.

Let w_t^{ij} be the weight agent i puts on the new information she learns from agent j at time t. In the main model, $w_t^{ij} = w_t^{ji} = 1$ for all connected pairs ij and all t. In this section, we maintain the weight each agent i attaches to her own signals as $w_t^{ii} = 1$ for all i and t. In terms of her neighbors, if $w_t^{ij} = 1$, agent i treats information from agent j equally; if $w_t^{ij} > 1$, agent i over values information from j; and if $w_t^{ij} < 1$, agent i under values information from j.

More importantly, we modify our assumption A1 to incorporate weights as local information:

A1' : For any agent $j, k \in \mathrm{N}_i$, i knows how much other agents value her information (w_t^{ki} and w_t^{ji}); and how much her neighbors value each other's information (w_t^{jk} and w_t^{kj}).

The learning process can be modified as follows. In step 1, agent i uses the same method to identify the inferred signals, i.e., y_t^{ij} with conditional distribution $\mathbf{\Delta p}_t^{ij} \equiv \{\alpha_t^{ij}(1), ..., \alpha_t^{ij}(N)\}$.

In step 2, agent i updates her own estimate using the signal from nature with a weight $w_t^{ii} = 1$, and using the inferred signals from step 1, y_t^{ij} with weights w_t^{ij} for all $j \in \mathrm{N}_i$.

$$p_{t+1}^i(n) = \frac{p_t^i(n) \prod_{h \in g_i} (\alpha_t^{ih}(n))^{w_t^{ih}}}{\sum_{n'=1}^N p_t^i(n') \prod_{h \in g_i} (\alpha_t^{ih}(n'))^{w_t^{ih}}}. \tag{13}$$

When $w_t^{ij} = 0$, agent i completely ignores the information from agent j. As w_t^{ij} increases, the inferred signal from agent j becomes more and more influential.[30]

[30] Agents update their estimates with weights using the formula (13) inspired by the idea that each agent treats the weights as numbers of copies of neighbors' signal that she uses to update her estimate.

In step 3, when forming an estimate 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 learns from i, $\mathbf{\Delta p}_t^{iji}(n)$ with a weight w_t^{ji}, and those from their common friends, $\mathbf{\Delta p}_t^{ijk}(n)$ for all $k \in \mathrm{N}_i \cap \mathrm{N}_j$, with weights w_t^{jk}.

$$p_{t+1}^{ij}(n) = \frac{p_t^j(n) \prod_{h \in (g_{ij}/j)} (\alpha_t^{ijh}(n))^{w_t^{jh}}}{\sum_{n'=1}^N p_t^j(n') \prod_{h \in (g_{ij}/j)} (\alpha_t^{ijh}(n'))^{w_t^{jh}}}. \tag{14}$$

Agent i's third-order estimates are formed by Bayes' rule in a similar way:

$$p_{t+1}^{ijk}(n) = \frac{p_t^k(n) \prod_{h \in ((g_{ij} \cap g_{ik})/k)} (\alpha_t^{ijkh}(n))^{w_t^{kh}}}{\sum_{n'=1}^N p_t^k(n') \prod_{h \in ((g_{ij} \cap g_{ik})/k)} (\alpha_t^{ijkh}(n'))^{w_t^{kh}}}. \tag{15}$$

Each of agent i's lower-order estimates are thus formed iteratively from her next higher-order estimates. And this proceeds iteratively. $\|$

It is worth noting that when agents use different weights on different sources of information, Corollary 1 no longer holds. For example, $\mathbf{p}_t^{ij}$ can easily differ from $\mathbf{p}_t^{ji}$, because i and j use different weights on the inferred signal from their common neighbor k. But since agent i knows the weights her neighbors use, we can show that their estimates of a common neighbor's estimates remain the same. The following corollary extends results from Proposition 1 and Corollary 1.

Corollary **4.** *Consider learning with weights. For any $\{k'', ...j, ..., k'\} \subseteq g_{ij}/i$, any $k \in \mathrm{N}_i \cap \mathrm{N}_j$, and all t:*

(0) The highest order of estimates agent i needs to form about her neighbors is $\overline{L}^i + 1$.

(1) $\mathbf{p}_t^{ij} = \mathbf{p}_t^{jij}$.

(2) If (g_i, G_i) satisfies Ni, then $\mathbf{p}_t^{ik} = \mathbf{p}_t^{ijk} = \mathbf{p}_t^{ik''...j...k'k}$, and $\mathbf{p}_t^{iki} = \mathbf{p}_t^{iji} = \mathbf{p}_t^{ik''...j...k'i}$.

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

The first point is a direct extension of Proposition 1. That is if $distinct(i...jh) = distinct(k'...kh)$, then $\mathbf{p}_t^{i...jh} = \mathbf{p}_t^{k'...kh}$. The same set of agents agree on their higher-order estimates of h's estimates. The last agent h in the sequence needs to be the same, because different agents may use different weights on neighbors' information. From agent i's perspective and when $h = i$, the sequence of agents in the estimates can feature agent i twice, for example $\{ik'...ki\}$, and thus the higher-order estimates may have at most $\overline{L}^i + 1$ orders. If (g_i, G_i) satisfies Ni, agent i can divide her neighbors into several disjoint groups. From agent i's perspective, each group is fully connected, so they must see the same set of information and form the same estimates of their common neighbors' estimates. For example, i thinks her estimates of agent k must be the same as j's estimate of agent k, i.e., $\mathbf{p}_t^{ik} = \mathbf{p}_t^{ijk}$. Thus, agent i needs to form at most third-order estimates. If (g_l, G_l) satisfies Nl for every agent $l \in g_i$, each group divided by agent i above are truly fully connected with no other common friends. So members of each group truly share the same information set, then $\mathbf{p}_t^{ik} = \mathbf{p}_t^{jk}$. So the our learning with weights is well defined.

We note a few remarks. First, observe that learning in directed networks is a special case of learning with weights. When ij is a directed link from i to j, it is equivalent to letting $w_t^{ij} = 1$ and $w_t^{ji} = 0$. Second, we can adopt the same proof of Proposition 2 to show that signals can be decomposed in the learning procedure with weights. Moreover, in the normal-linear model, agents simply use weights to adjust the precision of the new information, i.e. $w_t^{ij} \triangle \pi_t^{ij}$.

6.1 Opinion leaders and stubborn agents

We start with the first scenario in which agents treat information differently depending on the source. We say that agent i is a *local opinion leader* if for any $j \in \mathrm{N}_i$, $w_t^{ji} = \frac{1}{\varepsilon}$ and $w_t^{ij} = \varepsilon$ for some small ε. That is agent i's neighbors put a very high weight on agent i's information, and agent i puts a very small weight on them. We call the neighbors of this opinion leader as her followers. Similarly, we say that agent i is *stubborn* if $w_t^{ij} = \varepsilon$ for all $j \in \mathrm{N}_i$. In other words, a stubborn agent is a local opinion leader with no followers. The learning with weights is still well behaved when the network is a social quilt.

Remark **1.** *If (g, G) is a social quilt, the learning with weights stops at period $T + D$. If agent i is a local opinion leader,*

$$p_{T+D}^k(n) = \frac{\prod_{t<T} (\phi_{n,t}^i)^{\frac{1}{\varepsilon}} \prod_{j \neq i, t<T} \phi_{n,t}^j}{\sum_{n'} \prod_{t<T} (\phi_{n',t}^i)^{\frac{1}{\varepsilon}} \prod_{j \neq i, t<T} \phi_{n',t}^j},$$

where $\phi^j_{n,t} = \phi^j_{mn}$ *if* $x^j_t = x_m$*, and* $k \neq i$*. If agent* i *is a stubborn agent,*

$$p^k_{T+D}(n) = \frac{\prod_{j \in B_k(i), t<T} (\phi^j_{n,t})^\varepsilon \prod_{j' \notin B_k(i), t<T} \phi^{j'}_{n,t}}{\sum_{n'} \prod_{j \in B_k(i), t<T} (\phi^j_{n',t})^\varepsilon \prod_{j' \notin B_k(i), t<T} \phi^{j'}_{n',t}},$$

where $B_k(i)$ *is the set of agents, except* i*, whose shortest path to agent* k *includes* i*.*

A local opinion leader can influence the estimates of the entire network. Consider agent k, a neighbor of one follower of the opinion leader. As k only knows the follower but not the opinion leader, she does not know the opinion leader's information gets exaggerated by the followers. She gives a fair weight to this exaggerated signal, and passes it along to her neighbors. So everyone, except the opinion leader, overvalues her information with a weight $\frac{1}{\varepsilon}$. In contrast, a stubborn agent loses her influence in our learning procedure. As i's neighbors know i puts a vanishingly small weight on their information, they can correctly predict i's estimates and avoid listening to her repeatedly. So a stubborn agent can only block information, but she cannot influence others.[31]

6.2 Learning with discounting

Now we consider the second scenario in which agents discount new inferred information over time, and we use the weights w^{ij}_t to represent the discounts in this subsection. Discounts can be caused by loss of information during communication or loss of attention once agents become confident. We show that if the discount w^{ij}_t is sufficiently small, estimates must converge. We say a posterior estimate is *non-degenerate* if all states have strictly positive probability.

PROPOSITION **10.** *If NG holds and the discounts are sufficiently small, the estimates must converge. Moreover, estimates converge to non-degenerate posteriors in the main model with no partitional signals and to those with a finite precision in the normal linear model.*

We remark that with discounts, the limit estimates could be non-Bayesian, and neighbors may disagree in the limit. We can see some stylized fact from the estimates in a social quilt, by extending

[31]One may wonder that the stubborn agent can exaggerate her signal by using a higher weight $w^{ii}_t > 1$ on her information from nature, but if this weight is known to her neighbors, they will discount the weight.

results in Remark 1: If agents use symmetric discounts w since period 1 and the network is a social quilt, the learning stops at period $T+D$, with

$$p^i_{T+D}(n) = \frac{\prod_{j,t\leq T}(\phi^j_{n,t})^{w^{d(ij)}}}{\sum_{n'}\prod_{j,t\leq T}(\phi^j_{n',t})^{w^{d(ij)}}},$$

where $\phi^j_{n,t} = \phi^j_{mn}$ if $x^j_t = x_m$ and $d(ij)$ is the social distance between i and j.

It suggests that in a social quilt, agent i's discount on agent j's signal purely depends on their social distance, and does not depend on the arrival time of the signal. If they are close to each other, j's signal gets a high weight, and if they are far away, j's signal gets a low weight. If i and j are neighbors, the weight they put on each signal x^k_t differ by at most one w, so even though they won't fully agree, the difference is bounded by one layer of discount. While if i and j are far away, when i puts a high weight on signal x^k_t, k must be close to i, and thus far away to j, so clearly j will put a very low weight on x^k_t. Thus i and j's estimates can be very different in our learning model with discounts, while consensus prevails in myopic learning and Bayesian learning models. Our learning model with discounts can lead to possible opinion polarization.

7 Extensions and discussions

7.1 Communication and richer message spaces

Agents only report their most up-to-date estimates about the distribution of the state in our simple learning procedure. Because the sources of information are never reported, the agents in our model can identify new information in a limited way, but they cannot learn more about the network structure.[32] In general, there is a limit to the agents' learning because the network is not common knowledge, and the agents don't have common priors over the network structure. How much can agents learn if they are allowed to communicate more than their own estimates? This section considers how our learning procedure can be extended to a larger message space, allowing the agents to communicate in some fashion about the source of their information. Not surprisingly, a

[32] This is also due to our assumption A2, whose role is discussed in section 7.4.

larger message space improves learning since the agents have perfectly aligned interests. But the amount of improvement depends on the message space allowed.

As a benchmark, suppose that the agents report all the information they have about their source of information. This is similar to Acemoglu, Bimpikis and Ozdaglar (2014), in which each signal is tagged and thus the agents are able to learn all the signals received over time.[33]

The largest message space is a tag system, in which the time and travel path of a signal, as well as the distribution of the signal, are reported. Recall that $\mathrm{N}_i^d = \{j : d(ij) = d\}$, that is, N_i^d is agent i's neighbors d steps away from her. Also recall that $\phi^i_{mn} = \Pr(x_m|s_n)$ for agent i. Let $\left(x_t^{k_d,\dots,k'_d,\dots,i}, \phi^{k_d}_{mn}\right)$ be the tag of a signal agent k_d, who are d steps away from agent i, received at time t. Moreover, for any given signal, there may be multiple paths through which the signal can reach an agent. We use $\left(x_t^{k_d,i}, \phi^{k_d}_{mn}\right)$ to denote the collection of all such paths. Then a message from agent i to a neighbor at time $t+1$ consists of:

$$\mathbf{m}^i_t = \left\{ \left(x^i_t, \phi^i_{mn}\right), \left\{ \left(x^{k_1,i}_{t-1}, \phi^{k_1}_{mn} \right) \right\}_{k_1 \in \mathrm{N}_i^{(1)}}, \dots, \left\{ \left(x^{k_d,i}_{t-d}, \phi^{k_d}_{mn} \right) \right\}_{k_d \in \mathrm{N}_i^{(d)}}, \dots \right\}. \tag{16}$$

Observe that the path of how a signal reaches agent i is reported, as well as the signal distribution of the agent who received that signal. Because the number of agents is finite, all the agents learn a new signal within D periods of the signal's arrival. Clearly, this would lead to Bayesian learning outcomes both in terms of signals and in terms of network structure.

7.2 When the network is common knowledge

There are two parts to our learning procedure: learning can only occur locally which is a restriction on the message space; and local knowledge of the network since agents do not have common knowledge of the network. To highlight the role of local learning, we briefly study how our learning procedure performs if the network is common knowledge. In particular, we show that the agents reach the correct Bayesian posterior beliefs in a reasonable amount of time even though we severely restrict the message space.

[33] Their paper differs from ours in two ways. First, they focus on the endogenous network formation and the role of information hub. Second, agents communicate signals, not the posteriors of their estimates.

Recall that the finite set of agents, $\{1, 2, ..., L\}$, are connected in an undirected network (g, G). Let (g, G) be common knowledge among agents. Also, $\mathrm{N}_i^d = \{j : G(ij) = d\}$ is the set of agent i's distance d neighbors and D is the maximum shortest distance between these agents. We use the normal-linear model here for simplicity. Further, new information only arrives at $t = 0$: Each agent i observes a signal and its precision, (x^i, ζ^i).[34] Recall that (g, G) is path-connected such that there is a path from every agent i to every agent j in the network. Then it is easy to see that using the most general message space (16) in Section 7.1, we have:

Remark **2.** *All agents learn the signals* $\{x^i, \zeta^i\}_{i=1,...,L}$ *since time* $t = D$.

In fact, agent i's estimate at time t, s_t^i, equals the Bayesian posterior beliefs based on signals received by agents within a distance t to agent i if $t < D$. Clearly, the agents have the correct estimates once they learn everyone's signals.

Using our learning procedure, however, agents only report their posteriors in each period and form the best estimate of their neighbor's estimates given their information. It is immediate that this is not without loss of generality. Consider a simple 4-agent network in Figure 5, in which agent 1 is connected to a triad 234. Then 1 will never learn 3 and 4 signals separately because 1 can only infer the new information based on 2's estimate, that is $\boldsymbol{\Delta}\mathbf{p}_2^{12}$ a Bayesian combination of 3 and 4's signals.

But we do have correct Bayesian learning outcomes. An agent's estimate may feature repetitions because she cannot distinguish individual signals beyond her immediate neighbors. We lay out one possible estimate formation process, while the estimate formation can be some other process as long as all agents follow the same one and it is common knowledge. At time t, agent i's information set consists all the information she observes, for example $I_t^i = \{\mathbf{p}_\tau^j : j \in g_i, \tau < t\}$. Then agent i forms her estimate by using a particular combination of signals from her information set I_t^i, say $\sum_j a_t^{ij}(x^j, \zeta^j)$ being agent i's estimate update formula. It must satisfies the following restrictions. First, a_t^{ij} must be a positive integer unless agent i does not hear about (x^j, ζ^j). We can order vectors $a_t^i = (a_t^{i1}, ..., a_t^{iL})$ by "Lexicographical order," that is $a_t^i > b_t^i$ if $a_t^{i1} > b_t^{i1}$ or if $a_t^{i1} = b_t^{i1}$ and $a_t^{i2} > b_t^{i2}$ or so on. Second, a_t^i must be the smallest vector such that $\sum_j a_t^{ij}(x^j, \zeta^j)$ is in I_t^i. It is

[34]We don't index the signals by time here since no new signals arrive afterwards.

easy to see that if agent i knows (x^j, ζ^j) at time t, a_t^{ij} must be 1.

Then her estimate is

$$\pi_t^i = \sum_{j=1}^{L} a_t^{ij} \zeta^j, \quad s_t^i = \frac{1}{\pi_t^i} (\sum_{j=1}^{L} a_t^{ij} x^j \zeta^j)$$

So if she knows all the individual signals, her estimate is the Bayesian posterior. And since the network is common knowledge, everyone's estimate update formula is also common knowledge.

LEMMA **4.** *The followings are true:*

L1. If no one change their information set at some time t, then the learning process completes (no one change their estimates after t).

L2. An agent learns the Bayesian estimate after at most L changes of her information set.

L3. If $a_t^i \neq a_t^j$ and $G(ij) = 1$, at least one of them learns something new.

Using these properties, we can show that the learning process always converges to a consensus, in which everyone holds the Bayesian estimate.

PROPOSITION **11.** *There exists some period $t < L^2$, such that all agents' estimates equal to the Bayesian posterior. In the normal linear case, for any agent i,*

$$\pi_t^i = \sum_{j=1}^{n} \zeta^j, \quad s_t^i = \frac{1}{\pi_t^i} \left(\sum_{j=1}^{n} x^j \zeta^j \right)$$

By L1, before the learning process completes, there must be at least one agent changing her estimate update formula in each period. And by L2, any agent won't change more than L times. So the learning process must completes within L^2 periods. L3 suggests that once it completes, it must be a consensus. Lastly, the consensus must be the Bayesian estimate because agent i knows (x^i, ζ^i) so in her estimate a_t^{ii} must be 1.

EXAMPLE **9.** *Learning in a network with 8 agents and with common knowledge of the network.*

Suppose 8 agents are connected in a network in Figure 6. At time 0, they see their own information and truthfully report it. At each time, they learn from their neighbors' previous reports and announced their updated estimates. We focus on the learning of agent 1.

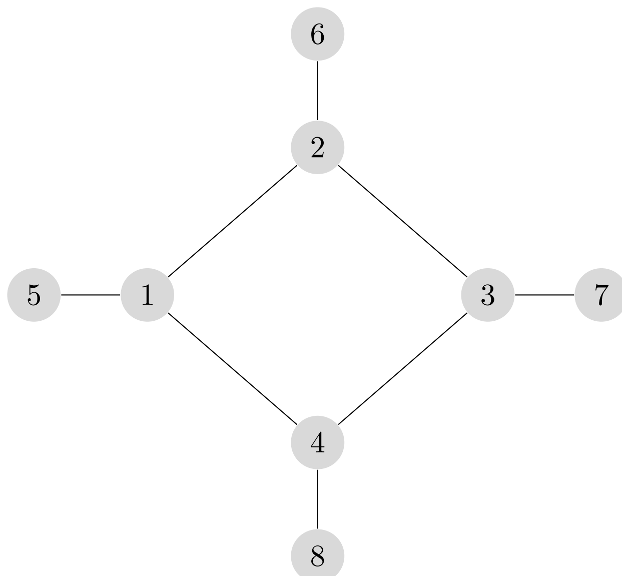

Figure 6: A network of 8 agents.

- At $t = 1$, agent 1 observes $(x^1, \zeta^1; x^5, \zeta^5; x^2, \zeta^2; x^4, \zeta^4)$, and use Bayesian rule to update her estimate to

$$\pi_2^1 = \zeta^1 + \zeta^5 + \zeta^2 + \zeta^4, \quad s_2^1 = \frac{1}{\pi_2^1}(x^1\zeta^1 + x^5\zeta^5 + x^2\zeta^2 + x^4\zeta^4)$$

- At $t = 2$, agent 1 learns $\zeta^3 + \zeta^6$ from agent 2 and $\zeta^3 + \zeta^8$ from agent 4. This is an example in which agent 1 cannot distinguish ζ^3, and cannot use a simple Bayesian procedure to update her estimate. So agent 1 follows the method above to update her estimate, and so agent 1 double counts ζ_3.

$$\pi_3^1 = \pi_2^1 + 2\zeta^3 + \zeta^6 + \zeta^8, \quad s_3^1 = \frac{1}{\pi_3^1}(s_2^1\pi_2^1 + 2x^3\zeta^3 + x^6\zeta^6 + x^8\zeta^8)$$

Similarly, agent 2 double counts ζ^4,

$$\pi_3^2 = \zeta^2 + \zeta^6 + \zeta^1 + \zeta^3 + 2\zeta^4 + \zeta^5 + \zeta^7, \quad s_3^2 = \frac{1}{\pi_3^2}(x^2\zeta^2 + x^6\zeta^6 + x^1\zeta^1 + x^3\zeta^3 + 2x^4\zeta^4 + x^5\zeta^5 + x^7\zeta^7)$$

- At $t = 3$, agent 1 knows $(\zeta^1, \zeta^5, \zeta^2, \zeta^4, \zeta^3 + \zeta^6, \zeta^3 + \zeta^8)$ and knows agent 2 double counts ζ^4. Based on π_3^2, agent 1 can infer ζ^7. So the updated estimate is

$$\pi_4^1 = \pi_3^1 + \zeta^7, \quad s_4^1 = \frac{1}{\pi_4^1}(s_3^1\pi_3^1 + x^7\zeta^7)$$

  Similarly, agent 2 can infer ζ^8.

- At $t = 4$, agent 1 learns ζ_8 from agent 2 which completes her knowledge of all signals. So the updated estimate is the full Bayesian estimate,

$$\pi_5^1 = \sum_{i=1}^{8} \zeta^i, \quad s_5^1 = \frac{1}{\pi_5^1}(\sum_{i=1}^{8} x^i\zeta^i)$$

  So do agents 2-4.

- At $t = 5$, agents $5 - 8$ update to the Bayesian estimate. And the learning completes.

By the end of $t = 5$, all agents hold the Bayesian estimate. $\diamond$

7.3 Endogenous network formation

One may wonder where the network comes from, and if agents build their connections endogenously, can they form a social quilt that generating correct Bayesian learning? In this subsection, we add a simple network formation game before the communication stage, and show that when agents are patient, they can form a social quilt network.

Suppose in period $t = -1$, agents form their communication network endogenously. In particular, agents enter the society sequentially based on their index, that is agent 1 comes first and agent L comes last. We follow the island-connection model in Jackson and Rogers (2005) that agents belong to H different groups, which present either their geographic locations or their types such as gender, occupation, race and so on. The cost of a link within each type is low, and the cost of a link across two types is relatively high, $C > c$. For simplicity, we assume the cost of within type connection is zero, $c = 0$. The sequence of agents and their types are common knowledge, but to be consistent with A2, the existing network is not observable. When agent i in group h^i enters the society, she has no information of the existing network and simply chooses the number of links to form to each group, $(l_1^i, ..., l_H^i)$. The l_h^i links will be formed uniformly randomly to agents in group h entering earlier than agent i. Denote the resulting network in the end of period $t = -1$ as (g, G).

Agents get values from holding a correct estimate of the states in each period. Let the Bayesian posterior belief $\beta(n) = \Pr(s = s_n | X_T)$, in which $X_T = \{x_t^i : t < T\}$. So the correct Bayesian posterior belief is $\mathbf{B} = (\beta(1), ..., \beta(N))$. Suppose $||\cdot||$ is a distance measure between two distributions of states. For example, we can use the standard KL divergence to measure the difference between the Bayesian posterior distribution $\mathbf{B}$ and an agent i's estimates $\mathbf{p}_t^i$:

$$\mathrm{D_{KL}}(\mathbf{B}, \mathbf{p}_t^i) = \sum_{n=1}^{N} \beta(n) log \frac{\beta(n)}{p_t^i(n)},$$

Lastly, recall information arrival stops since period T. Agent i's utility takes the form:

$$u^i = E_{X_T} \sum_{t=0}^{\infty} \left[\delta^t \left(v(-||\mathbf{B}(X_T), \mathbf{p}_t^i(X_T, G)||\right) - C \sum_{h \neq h^i} l_h^i \right]$$

where v is a strictly increasing value function, $C > (T + L)v(0)$ is a relatively high cost to sustain a cross-group link, and $\delta \in (0, 1)$ is the discount factor.

Proposition **12.** *For any C, there exist $\delta^* \in (0, 1)$ such that when $\delta > \delta^*$, there is a Nash equilibrium in which agents form a social quilt.*

Intuitively, as the cost of within-group connections is minimized, agents want to form a local clique by connecting to everyone in the same group. In addition, to gather complete information, each group wants to connect to the rest of the world; while the cost of cross-group connections is relatively high, each group would want to form only one connection to the outside. The resulting network is a social quilt, i.e., a tree-like union of cliques. Although the model is different, our equilibrium network could be similar to the efficient network in Jackson and Rogers (2005).

7.4 Knowledge and behavioral assumptions

Assumption A1 assumes that all agents know their immediate neighbors and connections among these neighbors. Intuitively, this is appropriate if the agents know their immediate social circles well. Instead, we may restrict the agents' knowledge of the network to just their neighbors, but not the connections among them. Clearly, with this limited knowledge of one's local network, agents can no longer discount information from their common neighbors. They only avoid repeatedly listening to themselves and to their immediate neighbors. Even in this case, however, the agents can still form the correct Bayesian estimates if the network is a tree with no circle. This is because no information can reach an agent again through connections of her own neighbors.

We also assume that agents follow A2: Agents are agnostic over the part of the network they do not know. Without sufficient information to make the correct inference of the whole network, agents are bound to make mistakes regardless of the learning procedure. We can, however, allow agents to learn in a limited fashion about their close neighbors, say those who are two links away from them. Doing so will reduce, but not eliminate mistakes if the agents cannot form the correct Bayesian estimates. But the information processing will become more complex for the agents.

7.5 Concluding remarks

We propose a simple and tractable learning procedure for agents in social networks. We show that even with knowledge only of the local network, the agents may form the correct Bayesian estimates in certain networks. We also show that this procedure performs just as well when the network is common knowledge, and better than myopic learning even when the agents fail to learn the Bayesian outcomes.

This procedure may help us better understand experimental data in which people are able to learn in a more sophisticated manner than myopic learning, but still make mistakes with respect to the information available. It can also help us remedy social isolation and opinion polarization at the lowest cost, for instance, by identifying and providing information to the key members in each network.

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.

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

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. 1993. "The Economics of Rumors." *The Review of Economic Studies*, 60: 309–327.

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

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 Inffuence, and Uni-Dimensional Opinions." *Quarterly Journal of Economics*, 118(3): 909–968.

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.

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

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.

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., and Brian W. Rogers. 2005. "The Economics of Small Worlds." *Journal of the European Economic Association*, 106(2): 265–295.

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.

Lorenz, Jan. 2005. "A stabilization theorem for dynamics of continuous opinions." *Physica A*, 355(1): 217–223.

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

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

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

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.

Munshi, Kaivan. 2003. "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market,." *Quarterly Journal of Economics*, 118(2): 549–597.

Munshi, Kaivan. 2004. "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." *Journal of Development Economics*, 73(1): 209–245.

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.

Appendix: Proofs

Proof to Proposition 1: For any sequence of fully connected (possibly repeated) agents $\{i...j\}$, let $distinct(i...j)$ be the set of all distinct agents of the original sequence. First, we show that for any higher-order estimates, if $distinct(i...j) = distinct(k'...k)$ and they are fully connected, then $\mathbf{p}_t^{i...j} = \mathbf{p}_t^{k'...k}$. Intuitively, this is because the information set they rely on is the same: $I_t^{i...j} = I_t^{k'...k} = \{\mathbf{p}_\tau^l : l \in g_{i...j}, \tau \leq t\}$. By definition,

$$g_{i...j} = g_i \cap ... \cap g_j = g_{k'} \cap ... \cap g_k = g_{k'...k},$$

because $g_h \cap g_h = g_h$, only the distinct agents matter.

We now prove this formally by induction. At $t = 0$, $I_0^{i...j} = I_0^{k'...k} = \{1/N, ..., 1/N\}$. Clearly, $\mathbf{p}_1^{i...j} = \mathbf{p}_1^{k'...k}$. Next, suppose this is true at period t. Then at $t+1$, the new information is inferred as

$$\alpha_t^{i...jl}(n) = \frac{p_t^l(n)}{p_t^{i...jl}(n)} \Bigg/ \sum_{n'} \frac{p_t^l(n')}{p_t^{i...jl}(n')} = \frac{p_t^l(n)}{p_t^{k'...kl}(n)} \Bigg/ \sum_{n'} \frac{p_t^l(n')}{p_t^{k'...kl}(n')} = \alpha_t^{k'...kl}(n).$$

This is because $distinct(i...jl) = distinct(k'...kl)$, then $\mathbf{p}_t^{i...jl} = \mathbf{p}_t^{k'...kl}$. Also, as we show that using the normalized or unnormalized αs does not affect the estimate update, we use the unnormalized version below for simplicity,

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

because $distinct(i...jk) = distinct(k'...kj)$. Recall in step 3,

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

So, if $distinct(i...j) = distinct(k'...k)$, then $\mathbf{p}_t^{i...j} = \mathbf{p}_t^{k'...k}$, equal to the estimates involving only distinct agents.

Suppose agents $\{i, j, k_1, k_2, ..., k_z\}$ is the largest fully connected subset of g_i. Then the number of distinct agents in any sequence $\{i...\}$ of the high-order estimates must be at most $z+2$. Suppose we can find some high-order estimates $\mathbf{p}_t^{i\cdots}$ such that there are $z+3$ distinct agents, Notice that if $G(kk') = 0$, agent i's higher-order estimate cannot feature agent k, k' together because k and k' don't know the existence of each other. So the $z+3$ agents must be all in g_i and be fully connected with each other, contracting the fact that $\{i, j, k_1, k_2, ..., k_z\}$ is the largest fully connected subset of g_i. That is for any high-order estimates $\mathbf{p}_t^{i\cdots}$, it is equal to the estimates involving only $distinct(i...)$ with an order at most $z+2$. Recall that $\overline{L}^i$ is the number of agents in the largest fully connected subset of g_i, that is $\overline{L}^i = z+2$. It is straightforward to show that the highest order of estimates agent i needs to form about her neighbors is $\overline{L}^i$. ∎

Proof of Corollary 1: (1) It is a direct corollary of Proposition 1.

(2) If (g_i, G_i) satisfies Ni, g_{ij} is fully connected, so for any $k \in g_{ij}/i$, $g_{ij} = g_{ik}$. I_t^{ij} is the same as I_t^{ik}, implying $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik}$. Formally, we prove the case by induction. Recall that agent i calculates $\mathbf{p}_1^{ij}$ based on her information set $I_0^{ij} = \{\mathbf{p}_0^l\}_{l\in g_{ij}}$. She forms her estimates of j's estimates of k's estimates $\mathbf{p}_1^{ijk}$ based on her information set $I_0^{ijk} = \{\mathbf{p}_0^l\}_{l\in g_{ij}\cap g_{ik}} = I_0^{ij}$. That is, all the $t=0$ reports that she can observe that k can observe that j observes. Similarly, she calculates $\mathbf{p}_1^{ik...j...k'}$ based on $I_0^{ik...j...k'} = \{\mathbf{p}_0^l\}_{l\in g_{ij}} = I_0^{ij}$. Since $p_0^l(n) = 1/N$ for all $l \in g_{ij}$, and agent i thinks that the agents in g_{ij} update using the same information, all the higher-order estimates are $p_1^{ik...j...k'}(n) = 1/N$ as well. Let agent i's estimates be $\mathbf{p}_1^{g_{ij}}$, where the superscript refers to i's estimates of agents in g_{ij}.

Next, suppose that this is true at period t: $\mathbf{p}_t^{ij} = \mathbf{p}_t^{ik} = \mathbf{p}_t^{ijk} = \mathbf{p}_t^{ik...j...k'} = \mathbf{p}_t^{g_{ij}}$. At period $t+1$, agent i observe the reports from all others in g_{ij}. Then by the updating rules given in (5) and (6), we can see that the numerator of agent i's estimates $p_{t+1}^{ij}(n)$ is the same as that of her estimates $p_{t+1}^{ik}(n)$:

$$\alpha_t^{ijl}(n) = \frac{p_t^l(n)}{p_t^{ijl}(n)} \Bigg/ \sum_{n'} \frac{p_t^l(n')}{p_t^{ijl}(n')} = \frac{p_t^l(n)}{p_t^{ikl}(n)} \Bigg/ \sum_{n'} \frac{p_t^l(n')}{p_t^{ikl}(n')} = \alpha_t^{ikl}(n).$$

This is because $\mathbf{p}_t^{ijl} = \mathbf{p}_t^{ikl}$ by induction. Also, using the unnormalized αs for simplicity, we have

$$p_t^j(n)\alpha_t^{ijk}(n) = p_t^j(n)\frac{p_t^k(n)}{p_t^{ijk}(n)} = p_t^k(n)\frac{p_t^j(n)}{p_t^{ikj}(n)} = p_t^k(n)\alpha_t^{ikj}(n),$$

this is because $\mathbf{p}_t^{ijk} = \mathbf{p}_t^{ikj}$ by induction. Similarly for the denominator. Thus $\mathbf{p}_{t+1}^{ik} = \mathbf{p}_{t+1}^{ik}$ are equal. Moreover, because $\mathbf{p}_{t+1}^{ijk}$ and high-order estimates $\mathbf{p}_{t+1}^{ik...j...k'}$ are calculated using the same information, all the higher-order estimates are the same.

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

Proof of Corollary 2: By definition,

$$\begin{aligned}
\alpha_{t+1}^{ij}(n) &= \frac{p_{t+1}^j(n)}{p_{t+1}^{ij}(n)} \Big/ \sum_{n'} \frac{p_{t+1}^j(n')}{p_{t+1}^{ij}(n')} \\
&= \frac{\dfrac{p_t^j(n)\prod_{h\in g_j}\alpha_t^{jh}(n)}{p_t^j(n)\prod_{h\in g_{ij}/j}\alpha_t^{ijh}(n)} \cdot \dfrac{\sum_{n'} p_t^j(n')\prod_{h\in g_{ij}/j}\alpha_t^{ijh}(n')}{\sum_{n'} p_t^j(n')\prod_{h\in g_j}\alpha_t^{jh}(n')}}{\sum_{n'} \dfrac{p_t^j(n')\prod_{h\in g_j}\alpha_t^{jh}(n')}{p_t^j(n')\prod_{h\in g_{ij}/j}\alpha_t^{ijh}(n')} \cdot \dfrac{\sum_{n'} p_t^j(n')\prod_{h\in g_{ij}/j}\alpha_t^{ijh}(n')}{\sum_{n'} p_t^j(n')\prod_{h\in g_j}\alpha_t^{jh}(n')}} \\
&= \prod_{l\in((g_j/g_i)\cup j)} \alpha_t^{jl}(n) \prod_{h\in g_{ij}/j} \frac{\alpha_t^{jh}(n)}{\alpha_t^{ijh}(n)} \Big/ \sum_{n'} \prod_{l\in((g_j/g_i)\cup j)} \alpha_t^{jl}(n') \prod_{h\in g_{ij}/j} \frac{\alpha_t^{jh}(n')}{\alpha_t^{ijh}(n')}. \qquad (17)
\end{aligned}$$

The last equality is true because if (g_i, G_i) satisfies Ni, $\alpha_t^{jh}(n) = \alpha_t^{ijh}(n)$ for all s_n and $j, h \neq j$, the ratios in the right hand side of expression (17) becomes 1, and we have expression (2) as given in the lemma. ∎

Proof of Proposition 2: We prove it by induction. The initial information set $\{x_0^{a,i}, x_0^{b,i}\}$ is simply $\{x_0^i, \emptyset\}$. That is, agent i is uninformed in one of (X_t^a, X_t^b), and learns x_0^i in the other. Therefore $\{\mathbf{p}_1^{a,i}, \mathbf{p}_1^{b,i}\} = \{\mathbf{p}_1^i, (\frac{1}{N}, ..., \frac{1}{N})\}$, and it is easy to check that (9) holds. Moreover, no one learns information from their neighbors at $t = 0$, so $\mathbf{p}_1^{ij} = \mathbf{p}_1^{a,ij} = \mathbf{p}_1^{b,ij} = (\frac{1}{N}, ..., \frac{1}{N})$, and the same for any higher order estimates. So (9), (10), and (11) hold at $t = 1$.

Suppose the lemma holds at time t, and we want to prove it also holds at time $t + 1$. In Step

1, recall that the inferred signal under X_t^a and X_t^b is respectively

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

Further, using (9) and (10), we have:

$$\begin{aligned}
\alpha_t^{ij}(n) &= \frac{p_t^{j}(n)}{p_t^{ij}(n)} \Big/ \sum_{n'} \frac{p_t^{j}(n')}{p_t^{ij}(n')} \\
&= \frac{p_t^{a,j}(n)p_t^{b,j}(n)}{\sum_{n'=1}^N p_t^{a,j}(n')p_t^{b,j}(n')} \cdot \frac{\sum_{n'=1}^N p_t^{a,ij}(n')p_t^{b,ij}(n')}{p_t^{a,ij}(n)p_t^{b,ij}(n)} \Big/ \sum_{n'} \frac{p_t^{j}(n')}{p_t^{ij}(n')} \\
&= \alpha_t^{a,ij}(n)\alpha_t^{b,ij}(n) \cdot \frac{\sum_{n'=1}^N p_t^{a,ij}(n')p_t^{b,ij}(n')}{\sum_{n'=1}^N p_t^{a,j}(n')p_t^{b,j}(n')} \Big/ \sum_{n'} \frac{p_t^{j}(n')}{p_t^{ij}(n')}. \qquad (18)
\end{aligned}$$

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

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

The second equality holds by (9) and (18), 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 (9) holds at time $t+1$. Similarly, in

Step 3,

$$
\begin{aligned}
p_{t+1}^{ij}(n) &= \frac{p_t^j(n)\alpha_t^{iji}(n)\prod_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\alpha_t^{ijk}(n)}{\sum_{n'=1}^N p_t^j(n')\alpha_t^{iji}(n')\prod_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\alpha_t^{ijk}(n')}\\
&= \frac{p_t^{a,j}(n)p_t^{b,j}(n)\alpha_t^{a,iji}(n)\alpha_t^{b,iji}(n)\prod_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\alpha_t^{a,ijk}(n)\alpha_t^{b,ijk}(n)}{\sum_{n'=1}^N p_t^{a,j}(n')p_t^{b,j}(n')\alpha_t^{a,iji}(n')\alpha_t^{b,iji}(n')\prod_{k\in \mathrm{N}_i\cap \mathrm{N}_j}\alpha_t^{a,ijk}(n')\alpha_t^{b,ijk}(n')}\\
&= \frac{p_{t+1}^{a,ij}(n)p_{t+1}^{b,ij}(n)}{\sum_{n'=1}^N p_{t+1}^{a,ij}(n')p_{t+1}^{b,ij}(n')}.
\end{aligned}
$$

Thus (10) also holds at time $t+1$. Lastly,

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

Thus (11) also holds at time $t+1$. By Proposition 1 the level of higher-order estimates is finite, so the proof is complete. ∎

Proof of Lemma 1: By definition, a social quilt does not contain a simple circle. Also, a social quilt must satisfy NG. To see this, note that for any agent i and any $j \in \mathrm{N}_i$, if there exist agents $k, k' \in \mathrm{N}_i \cap \mathrm{N}_j$, then $\{k, i, k', j, k\}$ must be a circle. By the definition of social quilts, $G(kk') = 1$. So Ni holds for any agent i, and thus NG holds.

Next, NG implies that if any three agents have a common neighbor, these four agents must form a clique. Consider a circle of four agents. No simple circle means that there must be at least one link across two nonadjacent agents, say agent i_0 and i_2. Then by NG, these four agents must be a clique. To continue, suppose a circle of l agents is a clique, and another agent l' is connected to two agents in the circle. By NG again, agent l' must be connected to all the other agents in the network. Therefore any circle must be part of a clique, which is the definition of social quilt. ∎

Proof of Proposition 3: Consider agent i and her local network (g_i, G_i). Let $|.|$ represent the number of states in any subset of S, $|S| = N$ and $|\emptyset| = 0$.

At $t = 1$, $p_1^i(n) = 1/|P^i(s^*)|$ if $s_n \in P^i(s^*)$, and 0 otherwise. By definition (12), this is the correct Bayesian belief for i at $t = 1$. 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 $s_n \in P^j(s^*)$, $\alpha_1^{ij}(n) = 1/|P^j(s^*)|$ and 0 otherwise. Thus the inferred signal y_0^{ij} has the same distribution as j's estimates: $\mathbf{\Delta p}_1^{ij} = \mathbf{p}_1^j$. Similarly, $\mathbf{\Delta p}_1^{kj} = \mathbf{p}_1^j$.

At $t = 2$, by Step 2 of our learning procedure, agent i's estimates are

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

if s_n belongs to the intersection $P_1^{g_i}(s^*) \equiv \cap\{P^h(s^*)\}_{h \in g_i}$, and $p_2^i(n) = 0$ otherwise. Similarly, the second-order estimates are $p_2^{ij}(n) = 1/|P_1^{g_{ij}}(s^*)|$, for s_n belonging to the intersection $P_1^{g_{ij}}(s^*) \equiv \cap\{P^h(s^*)\}_{h \in g_{ij}}$, and $p_2^{ij}(n) = 0$ otherwise. Clearly, $P_1^{g_i}(s^*) \subseteq P_1^{g_{ij}}(s^*)$, that is agent 1's information is finer than the information observed by both i and j. And so on for higher-order estimates.

If $\mathbf{p}_2^j \neq \mathbf{p}_2^{ij}$, there must be some states $s_n \in P_1^{g_{ij}}(s^*)$ that have zero probability under $\mathbf{p}_2^j$. As at period 1, the inferred signal y_1^{ij} has the same distribution as j's estimates, $\mathbf{\Delta p}_2^{ij} = \mathbf{p}_2^j$. Let $P_2^{g_i}(s^*)$ be the set of states agent i thinks are still possible at the beginning of $t = 3$, then $P_2^{g_i}(s^*) \subset P_1^{g_i}(s^*)$. It is important to notice that, because $P_1^{g_j}(s^*) \equiv \cap\{P^h(s^*)\}_{h \in g_j}$,

$$P_2^{g_i}(s^*) \equiv \cap\{P_1^{g_h}(s^*)\}_{h \in g_i} = \cap\{P^l(s^*)\}_{l \in \{g_h : h \in g_i\}}.$$

That is, $P_2^{g_i}(s^*)$ is the intersection of the original element of partitions of all agent i's $d = 1$ and $d = 2$ neighbors. Therefore $p_3^i(n) = 1/|P_2^{g_i}(s^*)|$ if $s_n \in P_2^{g_i}(s^*)$, and 0 otherwise.

Because there are no new signals, by at most $t = D + 1$, all the initial signals have reached agent i through her neighbors in a similar way. Their estimates are simply $1/|P^g(s^*)|$ if $s_n \in P^g(s^*) \equiv \cap\{P^l(s^*)\}_{l \in g}$, the intersection of all agents' element of partitions containing state s^*, and

0 otherwise. ∎

Proof of Proposition 4: We begin with a property of social quilts: if $d(ij) = d$, then there must be a unique path with 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 $i, i_1, i_2, ..., i_{d-1}, j$ and $i, j_1, j_2, ..., j_{d-1}, j$, and $i = i_0, j = i_d$. Then there must exist two numbers k and h, $0 \leq k < h \leq d$ such that:

$$\begin{cases} i_l = j_l & \text{if } l \leq k; \\ i_l \neq j_l & \text{if } l \in (k, h); \\ i_l = j_l & \text{if } l \geq h. \end{cases}$$

Clearly, $\{i_k, i_{k+1}, ..., i_h, j_{h-1}, ..., j_k\}$ must be a circle, going from i_k to herself through distinct agents. By definition, $i_l \neq j_l$ for any $l \in (k, h)$, and since $d(ii_l) = l$ and $d(ij_{l'}) = l'$, $i_l \neq j_{l'}$ whenever $l \neq l'$. But in a social quilt, any two agents in a circle are connected. Thus agent i_k and i_h must be connected and there is a unique shortest path between them, which is a contradiction.

The second property of social quilt is that if l is the last person on the shortest path from i to k, such that $d(ik) = d(il) + 1$ and $G(kl) = 1$, then for any j with $G(jk) = 1$ and $G(jl) = 0$, the shortest path from i to j must go through l and k, so $d(ij) = d(ik)+1$. To see this, note that since $G(jk) = 1$, the maximum 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 and l. Thus we have a simple circle involving jkl and $G(jl) = 0$, which is a contradiction to the definition of social quilts. Therefore $d(ij) = d(ik) + 1$.

(1) By Proposition 2, if we can show that agents' estimates agree with the Bayesian posterior beliefs with one signal, then it is also true with multiple signals. Without loss of generality, let agent i receive an initial signal $x_0^i = x_m$. Recall that $\mu_{nm}^i = \Pr(s_n|x_m)$ for agent i. We want to show that each agent j infers the signal x_m 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 $d(ij) = t - 1$ and $d(ik) = t - 2$, and $\alpha_t^{jk}(n) = 1/N$ otherwise. Notice that this implies

that agent j learns the signal and change his estimates at, and only at, $t = d(ij) + 1$.

We now prove this by induction on time t. First, this holds at $t = 2$. If $d(ij) = 1$, or $j \in \mathrm{N}_i$, then every agent j infers the signal from agent i's report in period 1 such that $\alpha_2^{ji}(n) = \mu_{nm}^i$. No other agents (including agent i) infers any new signal from their neighbors, $\alpha_2^{jk}(n) = 1/N$ if $d(ij) \neq 1$. If $\alpha_2^{jk}(n) = \mu_{nm}^i$, then since social quilts satisfy NG, by Corollary 2,

$$\alpha_2^{jk}(n) = \frac{\alpha_1^{ki}(n)}{\sum_{n'} \alpha_1^{ki}(n')},$$

and thus $\alpha_1^{ki}(n) = \mu_{nm}^i$. But since agent i is the only one who receives a signal from nature, $k = i$, and thus $d(ji) = 1$.

Next, suppose this holds at period t, we want to show it also holds at $t + 1$. By expression (2) in Corollary 2, $\alpha_{t+1}^{jk}(n)$ depends on $\alpha_t^{kl}(n)$ for all $l \in g_k/g_j$. By induction, $\alpha_t^{kl}(n) = \mu_{nm}^i$ if and only if $d(ik) = t - 1$ and $d(il) = t - 2$. As $G(jl) = 0$, by the second property above, it must be true that $d(ij) = t$. So we prove $\alpha_{t+1}^{jk}(n) = \mu_{nm}^i$ if and only if $d(ij) = t$ and $d(ik) = t - 1$.

As one 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 $D + 1$. So the learning is strongly Bayesian with signal x_0^i. When there are multiple signals, Proposition 2 ensures that the learning remains strongly Bayesian.

(2) We prove the result for any agent $i \in g$. To begin with, suppose that agent i's information structure is partitional. Together with the assumption that for any $s_{n'} \neq s_n$, there exists some signal x_m such that $\phi_{mn}^i \neq \phi_{mn'}^i$, we can show that each element of i's partition must include only one state, that is, $\mathcal{P}^i = \{(s_1), (s_2), ..., (s_N)\}$. Thus i can learn the true state from her initial signal.

Next, suppose that agent i's information structure is not partitional: $\phi_{mn}^i \in (0, 1)$ for some x_m, s_n. Let $\hat{X}_h^i$ denote a set of h informative signals for agent i: $\hat{X}_h^i = \{\hat{x}_1^i, ..., \hat{x}_h^i\}$ such that $\hat{x}_l^i \neq x_\emptyset$ for all $l = 1, ..., h$. If $\Pr(\hat{x}_l^i | s_{n''}) = 0$ for some $s_{n''}$ and $\hat{x}_l^i$, the agent believes that state $s_{n''}$ cannot be the true state. For the following proof, we limit attention to the remaining states such that $\Pr(s_n | \hat{x}_l^i) \neq 0$ for all $l = 1, ..., h$. Let the probability of observing these informative signals conditional on the true state be $\Pr(\hat{X}_h^i | s^*)$. Similarly, let the probability of observing these signals conditional on state s_n be $\Pr(\hat{X}_h^i | s_n)$. We use the standard KL divergence to measure the difference between these two probabilities. Formally, since all these signals are independent conditional on the state, we have:

$$\mathrm{D_{KL}}(s_n) = \mathrm{E}\left(log\frac{\Pr(\hat{X}_h^i|s^*)}{\Pr(\hat{X}_h^i|s_n)}\right) = \sum_{l=1}^{l=h}\Pr(\hat{x}_l^i|s^*)log\frac{\Pr(\hat{x}_l^i|s^*)}{\Pr(\hat{x}_l^i|s_n)},$$

which is non-negative. It is zero if $\Pr(\hat{x}_l^i|s_n) = \Pr(\hat{x}_l^i|s^*)$ for every l, which is ruled out by our assumption. Then by Bayes' rule: for any state s_n,

$$\frac{\Pr(s_n|\hat{X}_h^i)}{\Pr(s^*|\hat{X}_h^i)} = \frac{\Pr(\hat{X}_h^i|s_n)}{\Pr(\hat{X}_h^i|s^*)}\cdot\frac{\Pr(s_n)}{\Pr(s^*)}.$$

This implies that

$$\frac{1}{h}log\frac{\Pr(s_n|\hat{X}_h^i)}{\Pr(s^*|\hat{X}_h^i)} = -\frac{1}{h}\sum_{l=1}^{h}log\frac{\Pr(\hat{x}_l^i|s^*)}{\Pr(\hat{x}_l^i|s_n)}.$$

By the weak law of large numbers, the right hand side is the average of a sequence of random variables and converges to its expected value $-\mathrm{D_{KL}}(s_n)$ if h is sufficiently large. Since $\mathrm{D_{KL}}(s_n) > 0$,

$$log\frac{\Pr(s_n|\hat{X}_h^i)}{\Pr(s^*|\hat{X}_h^i)} \to -\infty, \text{ or } \frac{\Pr(s_n|\hat{X}_h^i)}{\Pr(s^*|\hat{X}_h^i)} \to 0$$

if h is sufficiently large. Thus if there are enough number of informative signals, agent i believes that the true state is s^* with probability 1.

We now turn to the case where the number of possible informative signals is arbitrarily large. The above analysis shows that if agent i observes a large number of informative signals, she can learn the true state with probability arbitrarily close to 1. That is for agent i, for any $\varepsilon > 0$, there exist a $\hat{T}^i$ such that if $t' \geq \hat{T}^i$, then

$$\Pr(\Pr(s^*|\hat{X}_{t'}^i) > 1-\varepsilon) > 1-\varepsilon.$$

We now show if agent i observes a sufficiently large number of signals, the number of informative signals must be larger than $\hat{T}^i$ with probability very close to 1. Recall that $\psi_0^i \in (0,1)$ is the probability agent i observes an uninformative signal in every period. For any $\hat{T}^i$, these exists an integer τ^i such that

$$(1-(\psi_0^i)^{\tau^i})^{\hat{T}^i} > 1-\varepsilon. \tag{19}$$

Here $1-(\psi_0^i)^{\tau^i}$ is the probability that agent i observes at least one informative signal in τ^i periods. Let $T^i = \tau^i \hat{T}^i$. Then the probability that agent i observes at least $\hat{T}^i$ informative signals by T^i is higher than $(1-(\psi_0^i)^{\tau^i})^{\hat{T}^i}$.

Recall that X_t^i is the set of all signals agent i observes until time t. For any $t > T^i$,

$$\begin{aligned}
\Pr(\Pr(s^*|X_t^i) > 1-\varepsilon) &\geq \sum_{t'=\hat{T}^i}^{t} \Pr(\Pr(s^*|\hat{X}_{t'}^i) > 1-\varepsilon)\Pr(\hat{X}_{t'}^i) && (20)\\
&> (1-\varepsilon)\sum_{t'=\hat{T}^i}^{t} \Pr(\hat{X}_{t'}^i) && (21)\\
&> (1-\varepsilon)^2, && (22)
\end{aligned}$$

where $\Pr(\hat{X}_{t'}^i)$ is the probability that out of t signals, t' of them are informative. So $\sum_{t'=\hat{T}^i}^{t}\Pr(\hat{X}_{t'}^i)$ is the probability agent i observes at least $\hat{T}^i$ informative signals until period $t > T^i$, which is higher than $1-\varepsilon$ by (19). Thus, $\Pr(\Pr(s^*|X_t^i) > 1-\varepsilon) > (1-\varepsilon)^2$.

Lastly, consider any subset of agents $g' = \{i_1, ..., i_z\}$. We can repeat the same process above and get a threshold $T^{g'}$, such that when $t \geq T^{g'}$, $\Pr(s^*|X_t^{i_1}, ..., X_t^{i_z}) > 1-\varepsilon$ with probability at least $(1-\varepsilon)^2$. Let $T^G = \max_{g' \subset \{1,...,L\}} T^{g'}$, then when $t \geq T^G$, for any subset of agents $g' = \{i_1, ..., i_z\}$, $\Pr(s^*|X_t^{i_1}, ..., X_t^{i_z}) > 1-\varepsilon$ with probability at least $(1-\varepsilon)^2$. Thus, as T becomes arbitrarily large, all agents' estimates after $T+D$ put a probability arbitrarily close to 1 on the true state. ∎

Proof of Proposition 5: For any simple circle $c = \{i_0, i_1, ..., i_k\}$, there are two separate cases: an agent $i \in c$ or $i \notin c$.

Case 1: $i \in c$. Without loss of generality, assume $i = i_0 = i_k$. First, $\alpha_2^{i_1 i_0}(n) = \alpha_2^{i_{k-1} i_0}(n) = \mu_{nm}^i$. Let $\mu_{nm}^i(\eta)$ be the Bayesian posterior of state s_n after seeing η copies of identical x_m, so $\mu_{nm}^i(\eta) = \frac{(\mu_{nm}^i)^\eta}{\sum_{n'}(\mu_{n'm}^i)^\eta}$. By Corollary 2, when NG holds, all inferred signals must have the distribution equal to $\mu_{nm}^i(\eta)$ for some non-negative η. Let $\alpha_t^{jk}(n) = \mu_{nm}^i(\eta_t^{jk})$. Corollary 2 can be rewritten as

$$\eta_{t+1}^{jk} = \sum_{l \in g_k / g_j} \eta_t^{kl}. \tag{23}$$

Because i_{j+2} and i_j are not connected, i_2 infers at least one x_m from i_1 at $t = 3$, i_3 infers at least one x_m from i_2 at $t = 4$, and so on. $\eta_{k+1}^{i_k i_{k-1}} \geq 1$ when the signal completes a round trip along the simple circle. Then agent i infers x_m in every k periods.

Case 2: $i \notin c$, without loss of generality, assume $i_0 = i_k$ is the first one of the simple circle (or one of the first ones) with a positive $\eta_t^{i_0 j}$ for some $j \in \mathrm{N}_{i_0}$. By NG, j cannot be connected with i_1 and i_{k-1} at the same time. Suppose $G(ji_1) = 0$, then $\eta_{t+1}^{i_1 i_0} \geq 1$. We can repeat the same process as above and show agent i_0 infers x_m repeatedly.

Lastly, we show that if some agent j (agent i_0 in the above two cases) learns an infinite number of x_m, all agents in the network must learn an infinite number of x_m. Let $p_t^h(n) = \mu_{nm}^i(\eta_t^h)$. Then we want to show that if η_t^j goes to ∞ for some j, η_t^h must go to ∞ for any agent h. We begin with h is j's neighbor, and we claim that $\eta_{t+1}^h \geq \eta_t^j$. By Corollary 1, $\mathbf{p}_{t+1}^{jh} = \mathbf{p}_{t+1}^{hj}$, that is

$$\eta_t^h + \eta_t^{hj} + \sum_{l \in \mathrm{N}_h \cap \mathrm{N}_j} \eta_t^{hl} = \eta_t^j + \eta_t^{jh} + \sum_{l \in \mathrm{N}_h \cap \mathrm{N}_j} \eta_t^{jl}.$$

Since the left-hand-side (LHS) is smaller than $\eta_{t+1}^h = \eta_t^h + \sum_{l \in \mathrm{N}_h} \eta_t^{hl}$ and the right-hand-side (RHS) is above η_t^j, we have $\eta_{t+1}^h \geq \eta_t^j$. Then if $d(jh) = 2$, there is a path from h to j, for example $G(hh') = G(h'j) = 1$, then $\eta_{t+2}^h \geq \eta_{t+1}^{h'} \geq \eta_t^j$. Recall that D is the diameter of the network, then $\min_{h \in g} \eta_{t+D}^h \geq \eta_t^j$, so all η_t^h goes to ∞. The expression of $p_\infty^j(n)$ in the proposition is simply $\lim_{\eta \to \infty} \mu_{nm}^i(\eta)$. ∎

Proof of Corollary 3: Suppose the only simple circle is $c = \{i_0, i_1, ..., i_l\}$ with distinct agents except $i_0 = i_l$. First, we focus on one signal $x_0^i = x_m$. c is the only simple circle in the network, which means the network outside of c must have social quilt structure—NG and no simple circle. Then the first time this signal arrives at the circle, it must be learned by either only one agent (say i_0), or two connected agents from the same source (say i_0 and i_1 learn from their common friend).

If not, suppose i_0 and i_k learns the signal at the same time, either $k \neq 1, l-1$ or i_k learns from a different source, then there is another simple circle inside the circle from i to i_0, i_0 to i_k through c, and i_k to i. It is a contradiction to the fact c is the only simple circle. In both cases, the signal is then learned by others in the circle in two opposite directions, that is $i_1 i_2 i_3 ...$ and $i_l i_{l-1} i_{l-2} ....$ In addition, agents in the circle do not learn any other new signal from agents outside the circle. Using the intuition from the proof of Proposition 4, if the signal goes back to the circle for a second time, it must involves another simple circle which contradicts the assumption. Let τ be the first time the signal arrives at i_0, the other case with the signal arriving at $i_0 i_1$ is very similar. Then $p_\tau^{i_0} = \mu_{nm}^i$, and $p_{\tau+lk}^{i_0} = \mu_{nm}^i(2l+1)$, that is after every l periods, agent i_0 gets two more copies of x_m due to the two opposite trips of the signal.

Now consider the case with multiple signals. For an abstract signal x, if the first time it arrives at the circle is $\tau(x)$ and to agent i_0, then it gets counted as $2\left\lfloor \frac{t-\tau(x)}{l} \right\rfloor + 1$ copies by the agent i_0 at $t > \tau(x)$. If all signals arrives at the circle at the same time and at the same agent, then $p_t^{i_0}(\tilde{s})$ must converge to 1. Note that the different arrival time and different arrival locations can only cause a difference bounded by a finite number of repeated signals, i.e., the difference in copies due to arrival time is bounded by $T + D$ periods and the difference in copies due to arrival locations is bounded by l periods. So it won't change the fact that $p_t^{i_0}(\tilde{s})$ converges to 1, and so do all agents in the circle. As the network outside the circle has a social quilt structure, they also learn $\tilde{s}$. ∎

Proof of Proposition 6: When (g, G) does not satisfy NG, we focus on the learning of one triad of agents ijk and show that if there is one informative signal x_m from agent k', these agents cannot form Bayesian posteriors.

At $t = 2$, agent i and j learn the signal, $p_2^i(n) = p_2^j(n) = \mu_{nm}^{k'}$. At $t = 3$, $\alpha_3^{ki}(n) = \alpha_3^{kj}(n) = \mu_{nm}^{k'}$, but $\alpha_3^{ij}(n) = \alpha_3^{ji}(n) = \alpha_3^{ik}(n) = \alpha_3^{ij}(n) = 1$. Thus $p_3^i(n) = p_3^j(n) = \mu_{nm}^{k'}$, but $p_3^k(n) = \mu_{nm}^{k'}(2)$. From $t = 4$ onwards until there is new information reaching the triad ijk, $\alpha_t^{ij}(n) = \alpha_t^{ji}(n) = \alpha_t^{ik}(n) = \alpha_t^{ij}(n) = 1$. Thus $p_t^i(n) = p_t^j(n) = \mu_{nm}^{k'}$. But for agent k, for even periods, $\alpha_t^{ki}(n) = \alpha_t^{kj}(n) = \mu_{nm}^{k'}(-2)$, and thus $p_t^k(n) = 1/2$. Similarly, in odd periods, $p_t^k(n) = \mu_{nm}^{k'}(2)$. Clearly, agent k's beliefs are not Bayesian.

Suppose at $t = \tau$, new information reaches the triad ijk. Observe that even though there is only one initial signal, this signal may travel through multiple paths before reaching the agents again. We denote the signal as y_τ. There are two cases. First, this information reaches only one of i, j, k, say agent j, or all three of them. Then all agents observe it and update in period $\tau + 1$. Thus $p^i_{\tau+2}(n) = p^j_{\tau+2}(n) = \Pr(s_n|x^{k'}_0, y_\tau)$. By the signal decomposition result, agent k's estimates still oscillate between periods.

In the second case, y_τ reaches two of the agents via a common source, say agent j and k. Then we have, $\alpha^{ji}_{\tau+1}(n) = \alpha^{jk}_{\tau+1}(n) = 1$, $\alpha^{ij}_{\tau+1}(n) = \alpha^{ik}_{\tau+1}(n) = \Pr(s_n|y_\tau)$. Therefore from period $\tau + 1$ onwards, $p^j_{\tau+1}(n) = \Pr(s_n|x^{k'}_0, y_\tau)$. But for agent i, her estimates oscillate between $p^i_{\tau+2}(n) = \Pr(s_n|x^{k'}_0, y_\tau, y_\tau)$ and $p^i_{\tau+3}(n) = \Pr(s_n|x^{k'}_0)$. For agent k, his estimates oscillate between $\Pr(s_n|x^{k'}_0, x^{k'}_0, y_\tau)$ and $\Pr(s_n|y_\tau)$, which are clearly not Bayesian. Similarly, if the agents receive more signals from a common neighbor, some or all agents' estimates will oscillate between periods, and thus cannot be Bayesian.

The only possible signal that can stop agent k from oscillating is when i and j gets a signal y_τ from k' and when y_τ is the exact opposite of $x^{k'}_0$.[35] To see this, note that one possible estimates are $\Pr(s_n|x^{k'}_0, y_\tau)$ for i and j, and $\Pr(s_n|x^{k'}_0, x^{k'}_0, y_\tau, y_\tau)$ for agent k. If y_τ is the exact opposite of $x^{k'}_0$, then $p^i_{\tau+2}(n) = p^j_{\tau+2}(n) = p^k_{\tau+2}(n) = 1/N$. However, recall that $x^{k'}_0$ is the only signal from nature. Agent k' can get y_τ because he is part of a circle, so k' is going to keep on receiving new signals. While k's oscillating may pause for a few periods, it resumes when k' receives another new signal, and thus the learning cannot be Bayesian. ▮

Proof of Lemma 2: This can be shown by induction. At $t = 0$, $\triangle\pi^{ij}_1 = \pi^j_1 - \pi^{ij}_1 = \pi^j_1 \geq 0$ since $\pi^{ij}_1 = 0$. Suppose $\triangle\pi^{ij}_t \geq 0$ at time t, and we want to show it also holds at time $t + 1$.

$$\begin{aligned}
\triangle\pi^{ij}_{t+1} &= \pi^j_{t+1} - \pi^{ij}_{t+1} \\
&= \left[\pi^j_t + \zeta^j_t + \textstyle\sum_{k\in \mathrm{N}_j} \triangle\pi^{jk}_t\right] - \left[\pi^j_t + \triangle\pi^{ji}_t + \textstyle\sum_{k\in \mathrm{N}_i\cap \mathrm{N}_j} \triangle\pi^{jk}_t\right] \\
&= \zeta^j_t + \textstyle\sum_{k\in \mathrm{N}_j/g_i} \triangle\pi^{jk}_t \geq 0.
\end{aligned}$$

[35] It would also work if i and j gets this signal y_τ from one of their common friends k'' who is not connected to k. If $G(k'k'') = 1$, the case is identical. If $G(k'k'') = 0$, this signal will stop k from oscillating but cause k' to oscillate.

So $\triangle\pi_{t+1}^{ij} = 0$ only when $\zeta_t^j = 0$ and $\triangle\pi_t^{jk} = 0$ for all $k \in \mathrm{N}_j/g_i$. Then,

$$s_{t+1}^j = \frac{s_t^j\pi_t^j + x_t^j\zeta_t^j + \sum_{k\in \mathrm{N}_j}\triangle s_t^{jk}\triangle\pi_t^{jk}}{\pi_{t+1}^j} = \frac{s_t^j\pi_t^j + \sum_{k\in \mathrm{N}_j\cap g_i}\triangle s_t^{jk}\triangle\pi_t^{jk}}{\pi_{t+1}^{ij}} = s_{t+1}^{ij}.$$

That is when $s_t^j \neq s_t^{ij}$, it is true that $\triangle\pi_t^{ij} > 0$. ∎

Proof of Lemma 3: Without loss of generality, we assume the network is one component, otherwise, we can repeat the same proof for each component. Let $\overline{\pi}_{m,t} = \max_i \left(\pi_{m,t}^i\right)$ be the highest precision of all agents' estimates at time t and $\underline{\pi}_{m,t} = \min_i \left(\pi_{m,t}^i\right)$ the lowest. Let $\overline{L} = \max_i L_i$ be maximum number of neighbors an agent has, or the highest degree, in a component. The diameter of the network, D, is the longest shortest path between any two agents in the component. Then $\pi_{m,t+1}^i = \sum_{l\in \mathrm{N}_i\cup\{i\}}\pi_{m,t}^j$ implies that $\overline{\pi}_{m,t+1} \leq (\overline{L}+1)\overline{\pi}_{m,t}$. Moreover, for any $t > D$,

$$\overline{\pi}_{m,t+1} \leq (\overline{L}+1)^{D+1}\overline{\pi}_{m,t-D}.$$

It follows that the weight

$$T_t^{ij} \geq \frac{\underline{\pi}_{m,t}}{\overline{\pi}_{m,t+1}} \geq \frac{\underline{\pi}_{m,t}}{(\overline{L}+1)^{D+1}\overline{\pi}_{m,t-D}}.$$

We now show that $\underline{\pi}_{m,t} \geq \overline{\pi}_{m,t-D}$. Suppose agent l holds the highest precision at time $t-D$. By the learning procedure, all agents within a distance of τ to l must hold a precision weakly higher than $\overline{\pi}_{m,t-D}$ at time $t-D+\tau$. Since D is the diameter, all agents are within a distance of D to l. All their precision must be weakly higher than $\overline{\pi}_{m,t-D}$ at time t. So $\underline{\pi}_{m,t} \geq \overline{\pi}_{m,t-D}$. Therefore for $t > D$,

$$T_t^{ij} \geq \frac{1}{(\overline{L}+1)^{D+1}} \equiv \omega'.$$

Lastly, let $\underline{\omega}$ be the minimum of ω' and all the weights when $t \leq D$, then the lemma holds. ∎

Proof of Proposition 9: Let $\triangle\pi_t^i = \pi_{t+1}^i - \pi_t^i$ be the new information agent i learns from neighbors at time t, and $\triangle\pi_{m,t}^i$ is analogous. Let $\underline{\pi} = \min_{\pi_1^i>0}(\pi_1^i)$ be the smallest positive initial precision, so $\underline{\pi} > 0$. We claim that when $t > D$ and for any i, $\triangle\pi_{m,t}^i \geq \triangle\pi_t^i + \underline{\pi}$. In our learning procedure, $\triangle\pi_t^{ij} = \pi_t^j - \pi_t^{ij}$, while in myopic learning $\triangle\pi_{m,t}^{ij} = \pi_{m,t}^j$. It is easy to see that as long as $t > D$, for the same agent i and network (g, G), it must be that $\pi^i_{m,t} > \pi^i_t$. To see this, note that $\pi^i_{m,t} = \pi^i_t$ the first time agent i receives any information, whether from her own signal x^i_t or her neighbor's report. From the next period on, agent i does not discount her own signal and those she already learned from j, and thus $\pi^i_{m,t} > \pi^i_t$.

When $t > D$, information must have reached all agents, and thus $\pi^{ij}_t \geq \underline{\pi}$ for any connected ij.

$$\triangle\pi^i_{m,t} - \triangle\pi^i_t = \sum_{j \in \mathrm{N}_i} \triangle\pi^{ij}_{m,t} - \sum_{j \in \mathrm{N}_i} \triangle\pi^{ij}_t > \sum_{j \in \mathrm{N}_i} \pi^{ij}_t \geq \underline{\pi}.$$

It follows that

$$\pi^i_{m,\tau} - \pi^i_\tau > \sum_{t=D+1}^{\tau} (\triangle\pi^i_{m,t} - \triangle\pi^i_t) > (\tau - D)\underline{\pi}.$$

Hence the difference between the agent's estimate of her precision goes to infinity as time goes to infinity. ∎

Proof of Corollary 4: First, using the proof of Proposition 1, we can show that if $distinct(i...jh) = distinct(k'...kh)$, then $\mathbf{p}^{i...jh}_t = \mathbf{p}^{k'...kh}_t$. The last agent h in the sequence needs to be the same, because different agents may use different weights on neighbors' information. From agent i's perspective and when $h = i$, the sequence of agents in the estimates can feature agent i twice, for example $\{ik'...ki\}$, and thus the higher-order estimates may have at most $\overline{L}^i + 1$ orders.

We then prove $\mathbf{p}^{ik}_t = \mathbf{p}^{ijk}_t = \mathbf{p}^{ik''...j...k'k}_t$ by induction. At $t = 0$, by the same argument in the proof of Corollary 1, all these estimates are based on the same information set I^{ij}_0, so $p^{ik}_0(n) = p^{ijk}_0(n) = p^{ik''...j...k'k}_0(n) = 1/N$. Next, suppose that this is true at time t. Then $\alpha^{ikh}_t(n) = \alpha^{ijkh}_t(n)$ for any $h \in g_{ij}/k$. By (14), (15) and Ni, $p^{ik}_{t+1}(n) = p^{ijk}_{t+1}(n)$, and it is equal to all the higher-order estimates $p^{ik''...j...k'k}_{t+1}(n)$. We can use an analogous inductive proof to show $\mathbf{p}^{iki}_t = \mathbf{p}^{iji}_t = \mathbf{p}^{ik''...j...k'i}_t$.

Then we prove that $\mathbf{p}^{ik}_t = \mathbf{p}^{jk}_t$ if g_l satisfies Nl for every agent $l \in g_i$. Note that $g^{ik} = g^{jk}$, so the agents in these two subnetworks observes the same information set, so $p^{ik}_t(n) = p^{jk}_t(n)$. ∎

Proof of Proposition 10: Let $\overline{L} = \max_i L_i$. We prove a slightly stronger result: if the discounts $w^{ij}_t \leq \frac{1}{\overline{L}+1}$ for all $t \geq \tau$ where $\tau \geq T$, then the estimates must converge. That is we only need the discounts to be small enough in later periods. Let's start with the normal-linear model. Suppose at time t, the maximal precision of inferred signals is $\triangle\overline{\pi}_t = \max_{G(ij)=1} \pi_t^{ij}$. In the beginning of time $t+1$, by definition $\pi_t^{ij} \leq \triangle\overline{\pi}_t$, so the total discounted new information i learns from all neighbors is $\sum_{j\in \mathrm{N}_i} w_t^{ij}\pi_t^{ij} \leq \frac{\underline{L}}{\underline{L}+1}\triangle\overline{\pi}_t$. Notice that the new information's precision j infers from i in the next period must be smaller than $\frac{\overline{L}}{\underline{L}+1}\triangle\overline{\pi}_t$, because j may already know some of the new information i learns due to common friends. So the upper bound on the precision of the inferred signals one learns from another in period $t+1$ must satisfy $\triangle\overline{\pi}_{t+1} \leq \frac{\overline{L}}{\underline{L}+1}\triangle\overline{\pi}_t$. Then the precision of total new information one agent learns since period τ is below $\sum_{h=1}^{\infty}(\frac{\overline{L}}{\underline{L}+1})^h\triangle\overline{\pi}_\tau$, so it is finite.

Once the precision converges to a finite value (say π_∞^i for agent i), the estimate of state must also converge. Let $\overline{x} = \max_{i,t} x_t^i$ and $\underline{x} = \min_{i,t} x_t^i$, denote the upper bound and lower bound of all signals. They are well-defined because of the finiteness of signals, i.e., no new information since time T. For any $\varepsilon > 0$, there exits some $t(\varepsilon) > T$ such that $\pi_{t(\varepsilon)}^i \geq \frac{\pi_\infty^i}{2}$ and $\pi_\infty^i - \pi_{t(\varepsilon)}^i < \frac{\varepsilon}{\overline{x}-\underline{x}} \cdot \frac{\pi_\infty^i}{2}$. For any two periods $t' > t \geq t(\varepsilon)$,

$$|s_t^i - s_{t'}^i| = |s_t^i - \frac{s_t^i\pi_t^i + \triangle s_{t',t}^i(\pi_{t'}^i - \pi_t^i)}{\pi_{t'}^i}| = \frac{|s_t^i - \triangle s_{t',t}^i|(\pi_{t'}^i - \pi_t^i)}{\pi_{t'}^i} < \varepsilon,$$

where $(\triangle s_{t',t}^i, \pi_{t'}^i - \pi_t^i)$ is the aggregate new inferred signal agent i learns between periods t and t', so $\triangle s_{t',t}^i \in [\underline{x}, \overline{x}]$. Thus, s_t^i must converge.

The convergence in the main model can be proved in the same fashion. At time t, let $\overline{\alpha}_t = \max_{n,n',G(ij)=1} \frac{\alpha_t^{ij}(n)}{\alpha_t^{ij}(n')}$ be the highest ratio among all distributions of inferred signals. So $\overline{\alpha}_t \geq 1$. The aggregate new information i learns at time t after the discount is collinear to a vector α_t^i such that the highest ratio of all its elements must be below $(\overline{\alpha}_t)^{\frac{\overline{L}}{\underline{L}+1}}$. This is because each new signal agent i learns from her neighbor is discounted by $w_t^{ij} \leq \frac{1}{\underline{L}+1}$ and agent i has at most $\overline{L}$ neighbors. Then, the new information agent i learns at time $t+h$ after the discount is collinear to a vector α_{t+h}^i such that the highest ratio of all its elements are bounded below $(\overline{\alpha}_t)^{(\frac{\overline{L}}{\underline{L}+1})^{h+1}}$. For any $\varepsilon > 0$, there exists some time $t(\varepsilon) = \tau + k(\varepsilon)$ such that $(\overline{\alpha}_\tau)^{\sum_{h=k(\varepsilon)}^{\infty}(\frac{\overline{L}}{\underline{L}+1})^{h+1}} < 1+\varepsilon$. Then for any $t' > t > t(\varepsilon)$,

$$|p_t^i(n) - p_{t'}^i(n)| = |p_t^i(n)(1 - \frac{\prod_{t\leq t''<t'}\alpha_{t''}^i(n)}{\sum_{n'} p_t^i(n')\prod_{t\leq t''<t'}\alpha_{t''}^i(n')})| < \varepsilon.$$

The last inequality is true by the following arguments.

$$\frac{\prod_{t\le t''<t'}\alpha^i_{t''}(n)}{\sum_{n'}p^i_t(n')\prod_{t\le t''<t'}\alpha^i_{t''}(n')}=\frac{1}{\sum_{n'}p^i_t(n')\prod_{t\le t''<t'}\frac{\alpha^i_{t''}(n')}{\alpha^i_{t''}(n)}}\ge\frac{1}{\sum_{n'}p^i_t(n')(1+\varepsilon)}\ge 1-\varepsilon,$$

and

$$\frac{\prod_{t\le t''<t'}\alpha^i_{t''}(n)}{\sum_{n'}p^i_t(n')\prod_{t\le t''<t'}\alpha^i_{t''}(n')}\le\frac{\prod_{t\le t''<t'}\alpha^i_{t''}(n)}{\min_{n'}\prod_{t\le t''<t'}\alpha^i_{t''}(n')}=\max_{n'}\left(\frac{\prod_{t\le t''<t'}\alpha^i_{t''}(n)}{\prod_{t\le t''<t'}\alpha^i_{t''}(n')}\right)\le 1+\varepsilon.$$

Thus, $\mathbf{p}^i_t$ must converge. As $\sum_{h=k(\varepsilon)}^{\infty}(\frac{\overline{L}}{\overline{L}+1})^{h+1}$ is finite, the highest ratio in one's estimate does not go to infinity, so the estimate is non-degenerate. ▮

Proof of Lemma 4: We start with L1. If no one change their information set at some time t, it means no one learns new information in the previous period. While their estimates remain the same, no one learns new information in this period, and so on. Thus, the learning stops. Then we prove L2. Each time an agent changes her information set, she must learn a new equation of signals. As there are L pairs of state-precision variables, an agent needs at most L pairs of equations to learn all signals and after that her information set remains the same. Lastly we prove L3. Without loss of generality, suppose $a^i_t > a^j_t$ and $G(ij)=1$. Agent i must not know $\sum_k a^{jk}_t(x^k,\zeta^k)$. Because if she knew, a^i_t is not the smallest vector she can use, and it contradicts the estimate formation process. So agent i must learn something new. ▮

Proof of Proposition 12: We construct one equilibrium as follows. Each agent chooses to connect to all existing agents in the same group, as the cost is zero. Without loss of generality, assume agent $L-H+1$ to L belongs to group 1 to group H, then agent $L-H+2$ to L each also chooses to form one connection to group 1. It is easy to see that the resulting network (g,G) is a social quilt.

Then, we show when δ is sufficiently high, (g,G) is an equilibrium outcome. Conditional on all others in the group do not form cross-group connections, agent $L-H+2$ to L wants to form one cross-group link. Because without the link, agents will never learn the full information. As δ is sufficiently high, holding a correct estimate in the long run outweighs the cost of one link. Also, as the cost of forming a cross-group link is high, i.e., $C > (T + L)v(0)$, forming more than one link is not profitable. Because the biggest difference it makes is increasing the value from the first $T + L$ periods, while the cost is surely more than the benefit. ∎