OPERATIONS RESEARCH
Vol. 64, No. 3, May–June 2016, pp. 564–584
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2015.1364

Preferences, Homophily, and Social Learning

Ilan Lobel, Evan Sadler
Stern School of Business, New York University, New York, New York 10012
{ilobel@stern.nyu.edu, esadler@stern.nyu.edu}

We study a sequential model of Bayesian social learning in networks in which agents have heterogeneous preferences, and neighbors tend to have similar preferences—a phenomenon known as homophily. We find that the density of network connections determines the impact of preference diversity and homophily on learning. When connections are sparse, diverse preferences are harmful to learning, and homophily may lead to substantial improvements. In contrast, in a dense network, preference diversity is beneficial. Intuitively, diverse ties introduce more independence between observations while providing less information individually. Homophilous connections individually carry more useful information, but multiple observations become redundant.

Subject classifications: network/graphs; probability; games/group decisions.
Area of review: Special Issue on Information and Decisions in Social and Economic Networks.
History: Received June 2013; revision received June 2014; accepted December 2014. Published online in *Articles in Advance* April 7, 2015.

1. Introduction

We rely on the actions of our friends, relatives, and neighbors for guidance in many decisions. This is true of both minor decisions, like where to go for dinner or what movie to watch tonight, and major, life-altering ones, such as whether to go to college or get a job after high school. This reliance may be justified because others' decisions are informative: our friends' choices reveal some of their knowledge, which potentially enables us to make better decisions for ourselves. Nevertheless, in the vast majority of real-life situations, we view the implicit advice of our friends with at least some skepticism. In essentially any setting in which we would expect to find social learning, preferences influence choices as much as information does. For instance, upon observing a friend's decision to patronize a particular restaurant, we learn something about the restaurant's quality, but her decision is also influenced by her preference over cuisines. Similarly, a stock purchase signals both company quality and risk preferences. In each case, two individuals with the same information may reasonably choose opposite actions.

The structure of our social network determines what information we can obtain through social ties. This structure comprises not only the pattern of links between individuals, but also patterns of preferences among neighbors. Since differences in preferences impact our ability to infer information from our friends' choices, these preference patterns affect the flow of information. In real-world networks, homophily is a widespread structural regularity: individuals interact much more frequently with others who share similar characteristics or preferences.[1] To understand the effects of network structure on information transmission, we must therefore consider how varying preference distributions, and the extent of homophily, influence social learning.

In this paper, we study a sequential model of social learning to elucidate how link structure and preference structure interact to determine learning outcomes. We adopt a particularly simple representation of individual decisions, assuming binary states and binary actions, to render a clear intuition on the role of the network. Our work builds directly on that of Acemoglu et al. (2011) and Lobel and Sadler (2015), introducing two key innovations. First, we significantly relax the assumption of homogeneous preferences, considering a large class of preference distributions among the agents. Although all agents prefer the action matching the underlying state, they can differ arbitrarily in how they weigh the risk of error in each state. Second, we allow correlations between the link structure of the network and individual preferences, enabling our study of homophily.

Our principal finding is that network link density determines how preference heterogeneity and homophily will impact learning. In a relatively sparse network, diverse preferences are a barrier to information transmission: heterogeneity between neighbors introduces an additional source of signal noise. However, we find that sufficiently strong homophily can ameliorate this noise, leading to learning results that are comparable to those in networks with homogeneous preferences. In a dense network, the situation is quite different. If preferences are sufficiently diverse, dense connectivity facilitates learning that is highly robust, and too much homophily may lead to inefficient herding. Homophilous connections offer more useful information, but the information provided through additional ties quickly becomes redundant. Conversely, diverse ties provide less information individually, but there is more independence between observations. Hence, the value of a homophilous connection versus a diverse connection depends on the other social information available to an individual.

Our formal analysis begins with sparse networks, meaning networks in which there is a uniform bound on how many neighbors any agent can observe. We first present a specialized example to highlight key mechanisms underlying our results. There are two distinct reasons why preference heterogeneity reduces the informational value of an observation. One source of inefficiency is uncertainty about how our neighbors make trade-offs. Perhaps surprisingly, there is additional inefficiency simply by virtue of the opposing trade-offs agents with different preferences make. Consider a diner who prefers Japanese food to Italian. When choosing between restaurants, this diner might default to a Japanese restaurant unless there is strong evidence that a nearby Italian option is of higher quality. Observing this individual enter a Japanese restaurant is a much weaker signal of quality than observing the same individual enter an Italian restaurant. Consequently, this observation is unlikely to influence a person who strongly prefers Italian food. When connections are scarce, long-run learning depends upon individual observations carrying a greater amount of information, and homophily reduces inefficiency from both sources.

Further results show that preference heterogeneity is generically harmful in a sparse network. Theorem 1 shows that we can always render asymptotic learning impossible via a sufficiently extreme preference distribution. Moreover, we find that the improvement principle, whereby agents are guaranteed higher ex ante utility than any neighbor, breaks down with even a little preference diversity. The improvement principle is a cornerstone of earlier results in models with homogeneous preferences (Banerjee and Fudenberg 2004, Acemoglu et al. 2011), but we show that there can be no improvement principle unless strong preferences occur much less frequently than strong signals. Homophily reduces the inefficiency from preference diversity by rescuing the improvement principle: if an agent can identify a neighbor with preferences very close to her own, then she can use her signal to improve upon that neighbor's choice. We define the notion of a strongly homophilous network, in which each agent can find a neighbor with arbitrarily similar preferences in the limit as the network grows, and we find that learning in a strongly homophilous network is comparable to that in a network with homogeneous preferences. Moreover, adding homophily to a network that already satisfies the improvement principle will never disrupt learning.

In contrast, when networks are dense, preference heterogeneity plays a positive role in the learning process. When agents have many sources of information, the law of large numbers enables them to learn as long as there is some independence between the actions of their neighbors. If the support of the preference distribution is sufficiently broad, there are always some types that act on their private information, creating the needed independence. Goeree et al. (2006) find this leads to asymptotic learning in a complete network even with bounded private beliefs. Theorem 4 generalizes this insight to a broad class of network structures, showing that learning robustly succeeds within a dense cluster of agents as long as two conditions are satisfied: there is sufficient diversity within the cluster, and agents in the cluster are aware of its existence. Since homophily reduces the diversity of preferences among an agent's neighbors, it has the potential to interfere with learning in a dense network. Example 4 shows that if homophily is especially extreme, with agents sorting themselves into two isolated clusters based on preferences, informational cascades can emerge, and learning is incomplete. We might imagine a network of individuals with strongly polarized political beliefs, in which agents on each side never interact with agents on the other. Interestingly, a small amount of bidirectional communication between the two clusters is sufficient to overcome this failure. The seminal paper of Bikhchandani et al. (1992) emphasized the fragility of cascades, showing that introducing a little outside information can quickly reverse them. Our result is similar in spirit, showing that in the long run, the negative impact of homophily on learning in dense networks is also fragile.

We make several contributions to the study of social learning and the broader literature on social influence. First, we show clearly how preference heterogeneity has distinct effects on two key learning mechanisms. This in turn helps us understand how the network affects long-run learning. Most of the social learning literature assumes that all individuals in society have identical preferences, differing only in their knowledge of the world.[2] Smith and Sorensen (2000) offer one exception, showing that, in a complete network, diverse preferences generically lead to an equilibrium in which observed choices become uninformative. However, a key assumption behind this result is that preferences are nonmonotonic in the state: different agents may change their actions in opposite directions in response to the same new information. This excludes the case in which utilities contain a common value component, which is our focus. Closer to our work, Goeree et al. (2006) consider a model with private and common values in a complete network, obtaining a result that we generalize to complex dense networks. Models that situate learning in the context of investment decisions have also considered preference heterogeneity. This work finds that heterogeneity hinders learning and may increase the incidence of cascades, leading to pathological spillover effects (Cipriani and Guarino 2008). Our innovation is the incorporation preference heterogeneity and complex network structure into the same model. We obtain results for general classes of networks, providing new insights on the underlying learning mechanisms.

In tandem with our study of preference diversity, we shed light on the impact of homophily on social learning. Despite its prevalence, few papers on learning have considered homophily. Golub and Jackson (2012) study group-based homophily in a non-Bayesian framework, finding that homophily slows down the convergence of beliefs in a network. To the best of our knowledge, ours is the first paper on learning in a context with preference-based homophily. We offer a nuanced understanding of its effects, showing two ways homophily can impact information transmission: homophily increases the information content of observations while decreasing independence between observations. The first effect suggests one reason why empirical studies find that homophilous ties have more influence on behavior. For instance, Centola (2011) provides experimental evidence that positive health behaviors spread more slowly in networks with diverse types, whereas they propogate faster and more broadly in networks with greater homophily. In an empirical study of learning about health plan choices, Sorensen (2006) finds that employees learn more from peers in a similar demographic. Similarly, Conley and Udry (2010) find that the farmers in their study learn more about agricultural techniques from other farmers with similar wealth levels.

We can also fruitfully compare our results to those in the empirical literature on tie strength and social influence. There has been much debate on the relative importance of weak versus strong ties in propagating behaviors. Notable work has found that an abundance of weak ties bridging structural holes often accounts for the majority of diffusion or information flows in a network (see, for instance, Granovetter 1973, Burt 2004, and Bakshy et al. 2012). Others have emphasized the importance of strong ties that carry greater bandwidth, especially when the information being transmitted is complex (see Hansen 1999, Centola and Macy 2007, and Aral and Walker 2014). Aral and Van Alstyne (2011) analyze this "diversity-bandwidth" trade-off in behavior diffusion, arguing that weak, diverse ties are more likely to provide access to novel information, but strong, homophilous ties transmit a greater volume of information, potentially resulting in more novel information flowing through strong ties. Our work offers an alternative theoretical basis for this trade-off, showing how rationality implies that decisions made by a homophilous contact convey less noisy information and reflect similar priorities to our own. Furthermore, our analysis of the impact of network structure provides guidance on which contexts should lead strong or weak ties to be most influential.

We first describe our model before analyzing learning in sparse networks. Our analysis emphasizes the difficulties that preference diversity creates and the benefits that homophily confers. We next provide general results on dense networks, finishing with a brief discussion. Most proofs are contained in the main body of the paper, with more technical results given in the appendix.

2. Model

Each agent in a countably infinite set sequentially chooses between two actions, 0 and 1. We index agents by the order of their decisions, with agent $m < n$ choosing an action prior to agent n's decision. The decision of agent n is denoted by x_n. Each agent has a random, privately observed type $t_n \in (0, 1)$, and the payoff to agent n is a function of n's decision, n's type, and the underlying state of the world $\theta \in \{0, 1\}$. For simplicity, we assume that a priori the two possible states are equally likely. Agent n's payoff is

$$u(t_n, x_n, \theta) = \begin{cases} (1-\theta) + t_n, & \text{if } x_n = 0, \\ \theta + (1 - t_n), & \text{if } x_n = 1. \end{cases}$$

This utility function has two components: the first depends on whether $x_n = \theta$ and is common to all types, whereas the second is independent of the state, depending only on the private type t_n and the action chosen x_n. Agents balance between two objectives: each agent n wishes to choose $x_n = \theta$, but the agent also has a type-dependent preference for a particular action. The type t_n determines how agent n makes this trade-off. If $t_n = \frac{1}{2}$, agent n is neutral between the two actions and chooses solely based on which action is more likely to realize $x_n = \theta$. Higher values of t_n correspond to stronger preferences toward action 0, and lower values of t_n correspond to stronger preferences toward action 1. Restricting types to the interval $(0, 1)$ implies that no agent is precommitted to either action.

Agent n is endowed with a private signal s_n, a random variable taking values in an arbitrary metric space $\mathcal{S}$. Conditional on the state θ, each agent's signal is independently drawn from a distribution $\mathbb{F}_\theta$. The pair of measures $(\mathbb{F}_0, \mathbb{F}_1)$ is common knowledge, and we call these measures the signal structure. We assume that $\mathbb{F}_0$ and $\mathbb{F}_1$ are not almost everywhere equal, so agents have a positive probability of receiving an informative signal. The private belief $p_n \equiv \mathbb{P}(\theta = 1 \mid s_n)$ of agent n is a sufficient statistic for the information contained in s_n. We write $\mathbb{G}_i(r) \equiv \mathbb{P}(p_n \leqslant r \mid \theta = i)$ for the state-conditional private belief distribution. We assume the private beliefs have full support over an interval $(\underline{\beta}, \bar{\beta})$, where $0 \leqslant \underline{\beta} < \frac{1}{2} < \bar{\beta} \leqslant 1$. As in previous papers,[3] we say private beliefs are unbounded if $\underline{\beta} = 0$ and $\bar{\beta} = 1$ and bounded if $\underline{\beta} > 0$ and $\bar{\beta} < 1$.

In addition to the signal s_n, agent n also has access to social information. Agent n observes the actions of the agents in her neighborhood $B(n) \subseteq \{1, 2, \ldots, n-1\}$. That is, agent n observes the value of x_m for all $m \in B(n)$. The neighborhood $B(n)$ is a random variable, and the sequence of neighborhood realizations describe a social network of connections between the agents. We call the probability $q_n = \mathbb{P}(\theta = 1 \mid B(n),\ x_m \text{ for } m \in B(n))$ agent n's social belief.

We represent the structure of the network as a joint distribution $\mathbb{Q}$ over all possible sequences of neighborhoods and types; we call $\mathbb{Q}$ the network topology and assume $\mathbb{Q}$ is common knowledge. We allow correlations across neighborhoods and types since this is necessary to model homophily, but we assume $\mathbb{Q}$ is independent of the state θ and the private signals. The types t_n share a common marginal distribution $\mathbb{H}$, and the distribution $\mathbb{H}$ has full support in some range $(\underline{\gamma}, \bar{\gamma})$, where $0 \leqslant \underline{\gamma} \leqslant \frac{1}{2} \leqslant \bar{\gamma} \leqslant 1$. The range $(\underline{\gamma}, \bar{\gamma})$ provides one measure of the diversity of preferences. If both $\underline{\gamma}$ and $\bar{\gamma}$ are close to $\frac{1}{2}$, then all agents are nearly indifferent between the two actions prior to receiving information about the state θ. On the other hand, if $\underline{\gamma}$ is close to 0 and $\bar{\gamma}$ is close to 1, then there are agents with strong biases toward particular actions. Analogous to our definition for private beliefs, we say preferences are unbounded if $\underline{\gamma} = 0$ and $\bar{\gamma} = 1$ and bounded if $\underline{\gamma} > 0$ and $\bar{\gamma} < 1$.

For algebraic simplicity, we often impose the following symmetry and density assumptions, but where they save little effort, we refrain from invoking them.

Assumption 1. *The private belief distributions are antisymmetric*: $\mathbb{G}_0(r) = 1 - \mathbb{G}_1(1-r)$ *for all* $r \in [0, 1]$; *the marginal type distribution* $\mathbb{H}$ *is symmetric around* $1/2$; *the distributions* $\mathbb{G}_0$, $\mathbb{G}_1$, *and* $\mathbb{H}$ *have densities.*

Agent n's information set, denoted by $I_n \in \mathcal{I}_n$, consists of the type t_n, the signal s_n, the neighborhood realization $B(n)$, and the decisions x_m for every agent $m \in B(n)$. Agent n's strategy σ_n is a function mapping realizations of I_n to decisions in $\{0, 1\}$. A strategy profile σ is a sequence of strategies for each agent. We use σ_{-n} to denote the set of all strategies other than agent n's, $\sigma_{-n} = \{\sigma_1, \ldots, \sigma_{n-1}, \sigma_{n+1}, \ldots\}$, and we can represent the strategy profile as $\sigma = (\sigma_n, \sigma_{-n})$. Given a strategy profile σ, the sequence of actions $\{x_n\}_{n \in \mathbb{N}}$ is a stochastic process with measure $\mathbb{P}_\sigma$. We analyze the perfect Bayesian equilibria of the social learning game, denoting the set of equilibria by Σ.

We study asymptotic outcomes of the learning process. We interpret these outcomes as measures of long-run efficiency and information aggregation. We say that asymptotic learning occurs if, in the limit as n grows, agents act as though they have perfect knowledge of the state. Equivalently, asymptotic learning occurs if actions converge in probability on the true state:

$$\lim_{n \to \infty} \mathbb{P}_\sigma(x_n = \theta) = 1.$$

This is the strongest limiting result achievable in our framework. Asymptotic learning implies almost sure convergence of individual beliefs, but almost sure convergence of actions need not occur: agents must continue to act based on their signals to achieve full information aggregation. Whether asymptotic learning obtains is the main focus of our analysis, but we also comment on learning rates to shed some light on shorter term outcomes.

We conclude this section with two basic lemmas required in our subsequent analysis. These lemmas provide useful properties of belief distributions and characterize best response behavior.

Lemma 1. *The private belief distributions* $\mathbb{G}_0$ *and* $\mathbb{G}_1$ *satisfy the following properties*:

(a) *For all* $r \in (0, 1)$, *we have* $d\mathbb{G}_0/d\mathbb{G}_1(r) = (1-r)/r$.

(b) *For all* $0 < z < r < 1$, *we have* $\mathbb{G}_0(r) \geqslant ((1-r)/r) \cdot \mathbb{G}_1(r) + ((r-z)/2)\mathbb{G}_1(z)$.

(c) *For all* $0 < r < w < 1$, *we have* $1 - \mathbb{G}_1(r) \geqslant (r/(1-r))(1 - \mathbb{G}_0(r)) + ((w-r)/2)(1 - \mathbb{G}_0(w))$.

(d) *The ratio* $\mathbb{G}_0(r)/\mathbb{G}_1(r)$ *is nonincreasing in* r *and is strictly larger than* 1 *for all* $r \in (\underline{\beta}, \bar{\beta})$.

(e) *Under Assumption* 1, *we have* $d\mathbb{G}_0(r) = ((1-r)/r) \cdot d\mathbb{G}_0(1-r)$, *and* $d\mathbb{G}_1(r) = (r/(1-r))d\mathbb{G}_1(1-r)$.

Proof. Parts (a) through (d) comprise Lemma 1 in Acemoglu et al. (2011). Part (e) follows immediately from part (a) together with Assumption 1. ☐

Lemma 2. *Let* $\sigma \in \Sigma$ *be an equilibrium, and let* $I_n \in \mathcal{I}_n$ *be a realization of agent* n*'s information set. The decision of agent* n *satisfies*

$$x_n = \begin{cases} 0, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid I_n) < t_n, \\ 1, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid I_n) > t_n, \end{cases}$$

and $x_n \in \{0, 1\}$ *otherwise. Equivalently, the decision of agent* n *satisfies*

$$x_n = \begin{cases} 0, & \text{if } p_n < \dfrac{t_n(1-q_n)}{t_n(1-q_n) + q_n(1-t_n)}, \\[2ex] 1, & \text{if } p_n > \dfrac{t_n(1-q_n)}{t_n(1-q_n) + q_n(1-t_n)}, \end{cases}$$

and $x_n \in \{0, 1\}$ *otherwise.*

Proof. Agent n maximizes her expected utility given her information set I_n and the equilibrium $\sigma \in \Sigma$. Therefore, she selects action 1 if $\mathbb{E}_\sigma[\theta + (1 - t_n) \mid I_n] > \mathbb{E}_\sigma[(1-\theta) + t_n \mid I_n]$, where $\mathbb{E}_\sigma$ represents the expected value in a given equilibrium $\sigma \in \Sigma$. The agent knows her type t_n, so this condition is equivalent to $\mathbb{E}_\sigma[\theta \mid I_n] > t_n$. Since θ is an indicator function and t_n is independent of θ, we have $\mathbb{E}_\sigma[\theta \mid I_n] = \mathbb{P}_\sigma(\theta = 1 \mid s_n, B(n), x_m \text{ for all } m \in B(n))$, proving the clause for $x_n = 1$. The proof for $x_n = 0$ is identical with the inequalities reversed. The second characterization follows immediately from the first and an application of Bayes' rule. ☐

An agent chooses action 1 whenever the posterior probability that the state is 1 is higher than her type, and she chooses action 0 whenever the probability is lower than her type. Therefore, agent n's type t_n can be interpreted as the minimum belief agent n must have that the true state is $\theta = 1$ before she will choose action $x_n = 1$.

3. Sparsely Connected Networks

We call a network sparse if there exists a uniform bound on the size of all neighborhoods; we say the network is M-sparse if, with probability 1, no agent has more than M neighbors. We begin our exploration of learning in sparse networks with an example.

Example 1. Suppose the signal structure is such that $\mathbb{G}_0(r) = 2r - r^2$ and $\mathbb{G}_1(r) = r^2$. Consider the network topology $\mathbb{Q}$ in which each agent observes her immediate predecessor with probability 1, agent 1 has type $t_1 = \frac{1}{5}$ with probability 0.5 and type $t_1 = \frac{4}{5}$ with probability 0.5. Any other agent n has type $t_n = 1 - t_{n-1}$ with probability 1.

In this network, agents fail to asymptotically learn the true state, despite having unbounded beliefs and satisfying the connectivity condition from Acemoglu et al. (2011).

Without loss of generality, suppose $t_1 = \frac{1}{5}$. We show inductively that all agents with odd indices err in state 0 with probability at least $\frac{1}{4}$, and likewise agents with even indices err in state 1 with probability at least $\frac{1}{4}$. For the first agent, observe that $\mathbb{G}_0(\frac{1}{5}) = \frac{9}{25} < \frac{3}{4}$, so the base case holds. Now suppose the claim holds for all agents of index less than n, and n is odd. The social belief q_n is minimized if $x_{n-1} = 0$, taking the value

$$\frac{\mathbb{P}_\sigma(x_{n-1}=0 \mid \theta=1)}{\mathbb{P}_\sigma(x_{n-1}=0 \mid \theta=1)+\mathbb{P}_\sigma(x_{n-1}=0 \mid \theta=0)} \geqslant \frac{1/4}{1/4+1} = \frac{1}{5}.$$

It follows from Lemma 2 that agent n will choose action 1 whenever $p_n > \frac{1}{2}$. We obtain the bound

$$\mathbb{P}_\sigma(x_n = 1 \mid \theta = 0) \geqslant 1 - \mathbb{G}_0\left(\frac{1}{2}\right) = \frac{1}{4}.$$

An analogous calculation proves the inductive step for agents with even indices. Hence, all agents err with probability bounded away from zero, and asymptotic learning fails. Note this failure is quite different from the classic results on herding and informational cascades. Agents continue acting on their private signals in perpetuity, and both actions are chosen infinitely often. Here, the difference in preferences between neighbors confounds the learning process, bounding the informational content of any observation even though there are signals of unbounded strength.

This example sheds light on the mechanisms that drive our results in this section. Given a simple network structure in which each agent observes one neighbor, long-run learning hinges on whether an individual can improve upon a neighbor's decision. We can decompose an agent's utility into two components: utility obtained through copying her neighbor and the improvement her signal allows over this. The improvement component is always positive when private beliefs are unbounded, creating the possibility for improvements to accumulate toward complete learning over time. With homogeneous preferences, this is exactly what happens because copying a neighbor implies earning the same utility as that neighbor, so utility is strictly increasing along a chain of agents. In our example, differing preferences mean that an agent earns less utility than her neighbor if she copies. The improvement component is unable to make up this loss, and asymptotic learning fails.

Although we shall model far more general network structures, in which agents have multiple neighbors and are uncertain about their neighbors' types, this insight is important throughout this section. Learning fails if there is a chance that an agent's neighbors have sufficiently different preferences because opposing risk trade-offs obscure the information that is important to the agent. Homophily in the network allows individuals to identify neighbors who are similar to themselves, and an improvement principle can operate along homophilous connections.

3.1. Failure Caused by Diverse Preferences

This subsection analyzes the effects of diverse preferences in the absence of homophily. We assume that all types and neighborhoods are independently distributed.

Assumption 2. *The neighborhoods* $\{B(n)\}_{n\in\mathbb{N}}$ *and types* $\{t_n\}_{n\in\mathbb{N}}$ *are mutually independent.*

Definition 1. Define the ratio $R_t^\epsilon \equiv t(1-\epsilon)/(t(1-\epsilon) + (1-t)\epsilon)$. Given private belief distributions $\{\mathbb{G}_\theta\}$, we say that the preference distribution $\mathbb{H}$ is *M-diverse* with respect to beliefs if there exists ϵ with $0 < \epsilon < \frac{1}{2}$ such that

$$\int_0^1 \left[\mathbb{G}_0(R_t^\epsilon) - \left(\frac{1-\epsilon}{\epsilon}\right)^{1/M} \mathbb{G}_1(R_t^{1-\epsilon})\right] \mathrm{d}\mathbb{H}(t) \leqslant 0.$$

A type distribution with M-diverse preferences has at least some minimal amount of probability mass located near the endpoints of the unit interval. We could also interpret this as a polarization condition, ensuring there are many people in society with strongly opposed preferences. As M increases, the condition requires that more mass is concentrated on extreme types. For any $N < M$, an M-diverse distribution is also N-diverse. Our first theorem shows that without any correlation between types and neighborhoods, M-diversity precludes learning in an M-sparse network.

Theorem 1. *Suppose the network is M-sparse, and let Assumptions 1 and 2 hold. If preferences are M-diverse with respect to beliefs, then asymptotic learning fails. Moreover, for a given signal structure, the probability that an agent's action matches the state is uniformly bounded away from 1 across all M-sparse networks.*

Proof. Given the ϵ in Definition 1, we show the social belief q_n is contained in $[\epsilon, 1-\epsilon]$ with probability 1 for all n, which immediately implies all of the stated results. Proceed inductively; the case $n = 1$ is clear with $q_1 = \frac{1}{2}$. Suppose the result holds for all $n \leqslant k$. Let $B(k+1) = B$ with $|B| \leqslant M$, and let $\mathbf{x}_B$ denote the random vector of observed actions. If $\mathbf{x} \in \{0, 1\}^{|B|}$ is the realized vector of decisions that agent $k+1$ observes, we can write the social belief q_{k+1} as $1/(1+l)$, where l is a likelihood ratio:

$$l = \frac{\mathbb{P}_\sigma(\mathbf{x}_B = \mathbf{x} \mid \theta = 0)}{\mathbb{P}_\sigma(\mathbf{x}_B = \mathbf{x} \mid \theta = 1)} = \prod_{m \in B} \frac{\mathbb{P}_\sigma(x_m = \mathbf{x}_m \mid \theta = 0,\ x_i = \mathbf{x}_i,\ i < m)}{\mathbb{P}_\sigma(x_m = \mathbf{x}_m \mid \theta = 1,\ x_i = \mathbf{x}_i,\ i < m)}.$$

Conditional on a fixed realization of the social belief q_m, the decision of each agent m in the product terms is independent of the actions of the other agents. Since $q_m \in [\epsilon, 1-\epsilon]$ with probability 1, we can fix the social belief of agent m at the end points of this interval to obtain bounds. Each term of the product is bounded above by

$$\frac{\int_0^1 \mathbb{G}_0(R_t^{\epsilon})\, \mathrm{d}\mathbb{H}(t)}{\int_0^1 \mathbb{G}_1(R_t^{1-\epsilon})\, \mathrm{d}\mathbb{H}(t)} \leqslant \left(\frac{1-\epsilon}{\epsilon}\right)^{1/M}.$$

This in turn implies that $q_{k+1} \geqslant \epsilon$. A similar calculation using a lower bound on the likelihood ratio shows that $q_{k+1} \leqslant 1-\epsilon$. □

Regardless of the signal or network structure, we can always find a preference distribution that will disrupt asymptotic learning in an M-sparse network. Moreover, because of the uniform bound, this is a more severe failure than what occurs with bounded beliefs and homogeneous preferences. Acemoglu et al. (2011) show that asymptotic learning often fails in an M-sparse network if private beliefs are bounded, but the point at which learning stops depends on the network structure and may still be arbitrarily close to complete learning. In this sense, the challenges introduced by preference heterogeneity are more substantial than those introduced by weak signals.

An analysis of 1-sparse networks allows a more precise characterization of when diverse preferences cause the improvement principle to fail. In the proof of our result, we employ the following convenient representation of the improvement function under Assumptions 1 and 2. Recall the notation $R_t^y \equiv t(1-y)/(t(1-y)+(1-t)y)$.

LEMMA 3. *Suppose Assumptions 1 and 2 hold. Define* $\mathscr{Z}\colon [\frac{1}{2}, 1] \to [\frac{1}{2}, 1]$ *by*

$$\mathscr{Z}(y) = y + \int_0^1 \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\Big]\, \mathrm{d}\mathbb{H}(t). \qquad (1)$$

In any equilibrium σ, *we have* $\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{m\}) = \mathscr{Z}(\mathbb{P}_\sigma(x_m = \theta))$.

PROOF. Observe under Assumption 1 we have $\mathbb{P}_\sigma(x_n = \theta) = \mathbb{P}_\sigma(x_n = \theta \mid \theta = i)$ for each $i \in \{0, 1\}$, and $\int_0^1 f(t)\, \mathrm{d}\mathbb{H}(t) = \int_0^1 f(1-t)\, \mathrm{d}\mathbb{H}(t)$ for any f. Taking $y = \mathbb{P}_\sigma(x_m = \theta)$, we have $\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{m\})$ equal to

$$\int_0^1 \sum_{i=0}^{1} \mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{m\},\ x_m = i,\ \theta = 0,\ t_n = t)$$

$$\cdot \mathbb{P}_\sigma(x_m = i \mid \theta = 0)\, \mathrm{d}\mathbb{H}(t)$$

$$= \int_0^1 \sum_{i=0}^{1} \mathbb{P}(p_n \leqslant R_t^{q_n}) \mathbb{P}_\sigma(x_m = i \mid \theta = 0)\, \mathrm{d}\mathbb{H}(t)$$

$$= \int_0^1 \Big[y\mathbb{G}_0(R_t^{1-y}) + (1-y)\mathbb{G}_0(R_t^y)\Big]\, \mathrm{d}\mathbb{H}(t)$$

$$= y + \int_0^1 \Big[(1-y)\mathbb{G}_0(R_t^y) - y(1 - \mathbb{G}_0(R_t^{1-y}))\Big]\, \mathrm{d}\mathbb{H}(t)$$

$$= y + \int_0^1 \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_{1-t}^y)\Big]\, \mathrm{d}\mathbb{H}(t)$$

$$= y + \int_0^1 \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\Big]\, \mathrm{d}\mathbb{H}(t),$$

as desired. □

THEOREM 2. *Suppose the network is 1-sparse, and let Assumptions 1 and 2 hold. If either of the following conditions is met, then asymptotic learning fails.*

(a) *The preference distribution satisfies*

$$\liminf_{t \to 0} \frac{\mathbb{H}(t)}{t} > 0.$$

(b) *For some* $K > 1$ *we have*

$$\lim_{r \to 0} \frac{\mathbb{G}_0(r)}{r^{K-1}} = c > 0, \quad \textit{and}$$

$$\int_0^1 \frac{t^{K-1}(K - (2K-1)t)}{(1-t)^K}\, \mathrm{d}\mathbb{H}(t) < 0.$$

PROOF. Each part follows from the existence of an $\epsilon > 0$ such that $\mathscr{Z}(y)$ given in Lemma 3 satisfies $\mathscr{Z}(y) < y$ for $y \in [1-\epsilon, 1)$. This immediately implies that learning is incomplete.

Assume the condition of part (a) is met. Divide by $1-y$ to normalize; for any $\epsilon > 0$ we obtain the bound

$$\int_0^1 \Big[\mathbb{G}_0(R_t^y) - \frac{y}{1-y}\mathbb{G}_1(R_t^y)\Big]\, \mathrm{d}\mathbb{H}(t)$$

$$\leqslant \int_0^{1-\epsilon} \mathbb{G}_0(R_t^y)\, \mathrm{d}\mathbb{H}(t) + \int_{1-\epsilon}^1 \mathbb{G}_0(R_t^y)\, \mathrm{d}\mathbb{H}(t)$$

$$- \frac{y}{1-y} \int_y^1 \mathbb{G}_1(R_t^y)\, \mathrm{d}\mathbb{H}(t)$$

$$\leqslant \mathbb{G}_0(R_{1-\epsilon}^y) + \mathbb{H}(\epsilon) - \mathbb{G}_1\left(\frac{1}{2}\right)\frac{y\mathbb{H}(1-y)}{1-y}.$$

Choosing ϵ sufficiently small, the second term is negligible, while for large enough y, the first term approaches zero, and the third is bounded above by a constant less than zero. Hence the improvement term in Equation (1) is negative in $[1-\epsilon, 1)$ for sufficiently small ϵ.

The proof of part (b) is presented in the appendix. □

These conditions are significantly weaker than 1-diversity. We establish a strong connection between the tail thicknesses of the type and signal distributions. The improvement principle fails to hold unless the preference distribution has sufficiently thin tails; for instance, part (a) implies learning fails under any signal structure if types are uniformly distributed on $(0, 1)$. As the tails of the belief distributions become thinner, the type distribution must increasingly concentrate around $\frac{1}{2}$ in order for learning to remain possible. In many cases, little preference diversity is required to disrupt the improvement principle.

Even if preferences are not so diverse that learning fails, we might expect heterogeneity to significantly slow the learning process. Interestingly, once the distribution of preferences is concentrated enough to allow an improvement principle, learning rates are essentially the same as with homogeneous preferences. Lobel et al. (2009) show in a line network with homogeneous preferences that the probability of error decreases as $n^{-1/(K+1)}$, where K is a tail thickness parameter; we adapt their techniques to obtain the following result.

Proposition 1. *Suppose* $B(n) = \{n-1\}$ *with probability* 1 *for all* $n \geqslant 2$, *and let Assumptions* 1 *and* 2 *hold. Suppose an improvement principle holds, meaning that the function* $\mathcal{Z}$ *from Lemma* 3 *satisfies* $\mathcal{Z}(y) > y$ *for all* $y \in [\frac{1}{2}, 1)$. *If for some* $K > 1$ *we have*

$$\lim_{r\to 0} \frac{\mathbb{G}_0(r)}{r^{K-1}} = c > 0, \quad \text{and}$$

$$\int_0^1 \frac{t^{K-1}(K-(2K-1)t)}{(1-t)^K}\, \mathrm{d}\mathbb{H}(t) > 0,$$

then the probability of error decreases as $n^{-1/(K+1)}$:

$$\mathbb{P}_\sigma(x_n \neq \theta) = O(n^{-1/(K+1)}).$$

Proof. See appendix. □

Once an improvement principle holds, the size of the improvements, and hence the rate of learning, depends on the tails of the signal structure. Early work on learning suggests that diverse preferences can slow learning (Vives 1993, 1995), even if information still aggregates in the long run. Perhaps surprisingly, Proposition 1 shows that as long as an improvement principle still holds, asymptotic learning rates are identical to the homogeneous preferences case. This result suggests that the learning speed we find if preferences are not too diverse, or if we have homophily as in the next section, is comparable to that in the homogeneous preferences case.

3.2. The Benefits of Homophily

In this subsection, we extend our analysis to sparse networks with homophily: we consider networks in which agents are more likely to connect to neighbors with similar preferences to their own. We find that homophily resolves some of the learning challenges preference heterogeneity introduces. Analyzing homophily in our model presents technical challenges because we cannot retain the independence assumptions from the last subsection; we need to allow correlations to represent homophily. In a model with homogeneous preferences, Lobel and Sadler (2015) highlight unique issues that arise when neighborhoods are correlated. This creates information asymmetries leading different agents to have different beliefs about the overall structure of connections in the network. To avoid the complications this creates, and focus instead on issues related to preferences and homophily, we assume the following.

Assumption 3. *For every agent* n, *the neighborhood* $B(n)$ *is independent of the past neighborhoods and types* $\{(t_m, B(m))\}_{m<n}$.

This assumption allows significant correlations in the sequence of types and neighborhoods $\{(t_n, B(n))\}$, but the neighborhoods $\{B(n)\}$ by themselves form a sequence of independent random variables. This representation of homophily is a technical contribution of our paper. Instead of first realizing independent types and then realizing links as a function of these types, we reverse the order of events to obtain a more tractable problem. Although the assumption is not costless, we retain the ability to explore the role of preferences and homophily in a rich class of network topologies. In many cases, we can translate a network of interest into an identical (or at least similar) one that satisfies Assumption 3. For instance, Examples 2 and 3 in §4 show two different, but equivalent ways to model a network with two complete subnetworks, one with low types and one with high types. The first example is perhaps more intuitive, but the second, which satisfies Assumption 3, is entirely equivalent.

If agent m is a neighbor of agent n, then under our previous assumption, the distribution of t_n conditioned on the value of t_m is $\mathbb{H}$, regardless of agent m's type. With homophily, we would expect this conditional distribution to concentrate around the realized value of t_m. We define a notion of *strong homophily* to capture the idea that agents are able to find neighbors with similar types to their own in the limit as the network grows large.

To formalize this, we recall terminology introduced by Lobel and Sadler (2015). We use $\hat{B}(n)$ to denote an extension of agent n's neighborhood, comprising agent n's neighbors, her neighbors' neighbors, and so on.

Definition 2. A network topology $\mathbb{Q}$ features *expanding subnetworks* if, for all positive integers K, $\limsup_{n\to\infty} \mathbb{Q}(|\hat{B}(n)| < K) = 0$.

Let $\mathbb{N}_n = \{1, 2, \ldots, n-1\}$. A function $\gamma_n\colon 2^{\mathbb{N}_n} \to \mathbb{N}_n \cup \{\varnothing\}$ is a *neighbor choice function* for agent n if for all sets $B_n \in 2^{\mathbb{N}_n}$ we have either $\gamma_n(B_n) \in B_n$ or $\gamma_n(B_n) = \varnothing$.

A *chosen neighbor topology*, denoted by $\mathbb{Q}_\gamma$, is derived from a network topology $\mathbb{Q}$ and a sequence of neighbor choice functions $\{\gamma_n\}_{n\in\mathbb{N}}$. It consists only of the links in $\mathbb{Q}$ selected by the neighbor choice functions $\{\gamma_n\}_{n\in\mathbb{N}}$.

Without the expanding subnetworks condition, there is some subsequence of agents acting based on a bounded number of signals, which clearly precludes asymptotic learning. We interpret this as a minimal connectivity requirement. Neighbor choice functions allow us to make precise the notion of identifying a neighbor with certain attributes. We call a network topology *strongly homophilous* if we can form a minimally connected chosen neighbor topology in which neighbors' types become arbitrarily close. That is, individuals can identify a neighbor with similar preferences, and the subnetwork formed through these homophilous links is itself minimally connected.

Definition 3. The network topology $\mathbb{Q}$ is *strongly homophilous* if there exists a sequence of neighbor choice functions $\{\gamma_n\}_{n\in\mathbb{N}}$ such that $\mathbb{Q}_\gamma$ features expanding subnetworks, and for any $\epsilon > 0$ we have

$$\lim_{n\to\infty} \mathbb{Q}(|t_n - t_{\gamma_n(B(n))}| > \epsilon) = 0.$$

Theorem 3. *Suppose private beliefs are unbounded and Assumption 3 holds. If $\mathbb{Q}$ is strongly homophilous, asymptotic learning obtains.*

Proof. See appendix. ☐

This theorem is proved by repeatedly applying an improvement principle along the links of the chosen neighbor topology. If, in the limit, agents are nearly certain to have neighbors with types in a small neighborhood of their own, they will asymptotically learn the true state. The reasoning behind this result mirrors that of our previous negative result. With enough homophily, we ensure the neighbor shares the agent's priorities with regard to tradeoffs, so copying this neighbor's action entails no loss in ex ante expected utility. Unbounded private beliefs are then sufficient for improvements to accumulate over time. We note that the condition in Theorem 3 does not require the sequence of type realizations to converge. As we now illustrate, sufficient homophily can exist in a network in which type realizations are mutually independent and distributed according to an arbitrary $\mathbb{H}$ with full support on $(0, 1)$.

To better understand our positive result, consider its application to the following simple class of networks. We shall call a network topology $\mathbb{Q}_\kappa$ a *simple κ-homophilous* network if it has the following structure. Let κ be a nonnegative real-valued parameter, and let $\mathbb{H}$ be a type distribution with a density. The types are mutually independent and are generated according to the distribution $\mathbb{H}$. Given a realization of $\{t_m\}_{m\leqslant n}$, define i^* such that

$$t_n \in \left[\mathbb{H}^{-1}\left(\frac{i^*-1}{n-1}\right), \mathbb{H}^{-1}\left(\frac{i^*}{n-1}\right)\right), \tag{2}$$

and let m_i denote the index of the agent with ith smallest type among agents $1, 2, \ldots, n-1$. Define the weights $\{w_i\}_{i<n}$ by $w_{i^*} = n^\kappa$, $w_i = 1$ if $i \neq i^*$, and let

$$\mathbb{Q}_\kappa(B(n) = \{m_i\}) = \frac{w_i}{\sum_{j<n} w_j}.$$

For example, let $\mathbb{H}$ be the uniform distribution on $(0, 1)$, and suppose agent $n = 101$ has randomly drawn type $t_{101} = 0.552$. By Equation (2), we have $i^* = 55$ since $0.552 \in [0.55, 0.56)$. Suppose agent $n = 80$ has the 55th lowest type among the first 100 agents. Then, agent 101 will observe the decision of agent 80 with probability $101^\kappa/(101^\kappa + 99)$ and observe any other agent with probability $1/(101^\kappa + 99)$. Since agent 80, with the 55th lowest type among the first hundred agents, is likely to have a type near $t_{101} = 0.552$, this network exhibits homophily.

The parameter κ neatly captures our concept of homophily. A value of $\kappa = 0$ corresponds to a network with no homophily, in which each agent's neighborhood is a uniform random draw from the past, independent of realized types. As κ increases, agent n's neighborhood places increasing weight on the past agent whose type rank most closely matches agent n's percentile in the type distribution. Moreover, it is easy to check that $B(n)$ contains a past agent drawn uniformly at random, independent of the history, so Assumption 3 is satisfied. We obtain a sharp characterization of learning in these networks as a function of the parameter κ.

Proposition 2. *Suppose private beliefs are unbounded, Assumption 1 holds, and $\liminf_{t\to 0} \mathbb{H}(t)/t > 0$. Asymptotic learning obtains in a simple κ-homophilous network if and only if $\kappa > 1$.*

Proof. For the forward implication, two observations establish that Theorem 3 applies. First,

$$\lim_{n\to\infty} \mathbb{Q}_\kappa(B(n) = \{i^*\}) = \frac{n^\kappa}{n - 2 + n^\kappa} = 1$$

whenever $\kappa > 1$. Second, it follows from the Glivenko-Cantelli Theorem that for any $\epsilon > 0$, we have

$$\lim_{n\to\infty} \mathbb{Q}_\kappa(|t_n - t_{i^*}| > \epsilon) = 0.$$

Given any $\epsilon > 0$, we can find $N(\epsilon)$ with $\mathbb{Q}_\kappa(B(n) \neq \{i^*\}) \leqslant \epsilon/2$ and $\mathbb{Q}_\kappa(|t_n - t_{i^*}| > \epsilon) \leqslant \epsilon/2$ for all $n \geqslant N(\epsilon)$. Defining the functions $\{\gamma_n\}_{n\in\mathbb{N}}$ in the natural way, we have

$$\mathbb{Q}_\kappa(|t_n - t_{\gamma_n(B(n))}| > \epsilon) \leqslant \epsilon$$

for all $n \geqslant N(\epsilon)$. It is a simple exercise to show that $\mathbb{Q}_\kappa$ features expanding subnetworks.

For the converse result, fix $n > m$. Define $y = \mathbb{P}_\sigma(x_m = \theta)$,

$$P_i(t) = \mathbb{P}_\sigma(x_m = \theta \mid \theta = i,\ B(n) = \{m\},\ t_n = t), \quad \text{and}$$

$$q_i(t) = \frac{P_i(t)}{P_i(t) + 1 - P_{1-i}(t)}.$$

Using Assumption 1 and following calculations similar to the proof of Lemma 3, we obtain

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{m\}) = y + \int_0^1 \left[(1 - P_1(t))\mathbb{G}_0(R_t^{q_1(t)}) - P_1(t)\mathbb{G}_1(R_t^{q_1(t)})\right] d\mathbb{H}(t).$$

We can uniformly bound $P_i(t)$, and hence $q_i(t)$, using y as follows. Define $p^* = (n-1)/(n-2+n^\alpha)$, and note that the distribution of agent n's neighbor can be expressed as a compound lottery choosing uniformly from the past with probability p^* and choosing agent k^* otherwise. Conditioned on the event that a uniform draw from the past was taken, the probability that $x_m = \theta$ is equal to y when conditioned on either possible value of θ. Thus we have

$$p^* y \leqslant P_i(t) \leqslant 1 - p^*(1-y)$$

for each i and for all $t \in (0,1)$. As long as p^* is uniformly bounded away from zero, we can make a slight modification to the argument of Theorem 2 part (a) to show that learning fails. This is precisely the case if and only if $\alpha \leqslant 1$. □

With $\kappa > 1$, the likelihood of observing the action of a neighbor with a similar type grows fast as n increases. While types across the entire society could be very diverse—$\mathbb{H}$ could be the uniform distribution on $(0,1)$, for instance—agents almost always connect to those who are similar to themselves. If the density of the type distribution is bounded away from zero near the end points 0 and 1, the threshold $\kappa = 1$ is sharp. With any less homophily, there is a nontrivial chance of connecting to an agent with substantially different preferences, and asymptotic learning fails.

Our last proposition in this section complements our results on the benefits from homophily. We show that, in a certain sense, making a network more homophilous never disrupts social learning. Under Assumption 3, we can describe a network topology via a sequence of neighborhood distributions together with a sequence of conditional type distributions $\mathbb{H}_{t_n \mid t_1, \ldots, t_{n-1}, B(1), \ldots, B(n)}(t)$. These conditional distributions may depend on the history of types $t_1, \ldots, t_{n-1}$ and the history of neighborhoods $B(1), \ldots, B(n)$. To simplify notation, we represent the conditional type distribution of agent n by $\mathbb{H}_n$.

Definition 4. Let $\mathbb{Q}$ and $\mathbb{Q}'$ be two network topologies that satisfy Assumption 3. We say that $\mathbb{Q}$ and $\mathbb{Q}'$ are *modifications* of one another if they share the same neighborhood distributions and the same marginal type distribution $\mathbb{H}$.

If the network topologies $\mathbb{Q}$ and $\mathbb{Q}'$ are modifications of each other, they represent very similar networks. They will have identical neighborhood distributions and equally heterogeneous preferences. However, the conditional distributions $\mathbb{H}_n$ and $\mathbb{H}'_n$ can differ, so the networks may vary significantly with respect to homophily. This will allow us to make statements about the impact of increasing homophily while keeping other properties of the network topology constant. Given two network topologies that are modifications of each other, we can take convex combinations to create other modifications of the original two.

Definition 5. Let $\mathbb{Q}$ and $\mathbb{Q}'$ be modifications of one another with conditional type distributions $\mathbb{H}_n$ and $\mathbb{H}'_n$, respectively. Given a sequence $\lambda = \{\lambda_n\}_{n\in\mathbb{N}} \in [0,1]^{\mathbb{N}}$, we define the *$\lambda$-mixture* of $\mathbb{Q}$ and $\mathbb{Q}'$ as the modification $\mathbb{Q}^{(\lambda)}$ of $\mathbb{Q}$ with conditional type distributions $\mathbb{H}_n^{(\lambda)} = \lambda_n \mathbb{H}_n + (1-\lambda_n)\mathbb{H}'_n$.

We consider networks in which asymptotic learning obtains by virtue of an improvement principle, and we find that mixing such networks with strongly homophilous modifications preserves learning.

Definition 6. We say that a network topology $\mathbb{Q}$ *satisfies the improvement principle* if any agent can use her signal to achieve strictly higher utility than any neighbor. Formally, this means there is a continuous, increasing function[4] $\mathscr{Z}\colon [1,3/2] \to [1,3/2]$ such that $\mathscr{Z}(y) > y$ for all $y \in [1,3/2)$, and $\mathbb{E}_\sigma[u(t_n, x_n, \theta) \mid B(n) = \{m\}] \geqslant \mathscr{Z}(\mathbb{E}_\sigma[u(t_m, x_m, \theta)])$ for all $m < n$.

This definition expresses the idea that an agent in the network can combine her private signal with information provided by one of her neighbors to arrive at a better decision than that neighbor.

Proposition 3. *Suppose private beliefs are unbounded, and Assumption 3 holds. Assume the network topology $\mathbb{Q}$ features expanding subnetworks and satisfies the improvement principle. Suppose $\mathbb{Q}'$ is any strongly homophilous modification of $\mathbb{Q}$. Asymptotic learning obtains in any λ-mixture of $\mathbb{Q}$ and $\mathbb{Q}'$.*

Proof. See appendix. □

Proposition 3 shows if we take any network topology satisfying the natural sufficient conditions for learning in a sparse network—expanding subnetworks and the applicability of the improvement principle—and consider a convex combination between this network and one with strong homophily, agents will still learn. One interpretation of this statement is that the addition of homophily to a network never harms learning that is already successful without homophily.

Our findings in this section establish a positive role for homophily in social learning. Without homophily, a sparse network with diverse preferences struggles to aggregate information because the meaning of a decision is difficult to interpret, and different agents make different trade-offs. Homophily counters both sources of inefficiency, leading to better outcomes as the degree of homophily increases. In the next section, we see the other side of homophily's impact on learning. Networks with dense connections offer an opportunity to overcome the difficulty of learning with bounded private beliefs. The ability to observe a diverse neighborhood is a major driver of this positive result, and an excess of homophily can stifle the learning process.

4. Densely Connected Networks

A classic result by Banerjee (1992) and Bikhchandani et al. (1992) shows that informational cascades occur in a complete network when private beliefs are bounded and preferences are homogeneous. More recently, Goeree et al. (2006) find that unbounded preferences can remedy the situation, proving a strong positive result in a comparable setting. Intuitively, if the distribution of preferences has full support on $(0, 1)$ and types are independent, then agents always emerge whose preferences roughly balance against their social information. These agents must rely on their private signals to choose an action, so new information is revealed to the rest of the network. Even if an individual decision provides little information, dense connections allow learning to operate through a law-of-large-numbers mechanism in addition to the improvement principle. The presence of this second learning mechanism leads to robust learning.

In contrast to our results for sparse networks, preference diversity appears decisively beneficial with full observation. In this section, we show this insight is generic in a broad class of dense network components that we call *clusters*.

Definition 7. A cluster $\mathscr{C}$ is a sequence of stopping times $\{\alpha_i\}_{i\in\mathbb{N}}$ with respect to the filtration generated by $\{(t_n, B(n))\}_{n\in\mathbb{N}}$ that satisfy

$$\lim_{i\to\infty} \mathbb{Q}(\alpha_k \in B(\alpha_i)) = 1 \quad \text{for all } k \in \mathbb{N}.$$

A cluster is a generalization of the concept of a (randomly generated) clique. Any complete subnetwork—a clique—with infinitely many members is a cluster. A subset $\mathscr{C}$ of the agents such that any member of $\mathscr{C}$ is observed with probability approaching one by later members of $\mathscr{C}$ is also a cluster. Clusters may exist deterministically, with the indices α_i being degenerate random variables as in the complete network, or they may arise stochastically. For instance, suppose types are i.i.d. and all agents with types above 0.8 are connected with each other through some correlation between neighborhoods and types. In this case, the stopping time α_i would refer to the ith agent with type above 0.8 and this group would form a cluster. Note that a cluster potentially has far fewer edges than a clique; the ratio of the number of edges in a cluster to the number of edges in a clique can approach zero as n grows.

As in §3, we introduce an independence assumption for tractability.

Assumption 4. *For each stopping time α_i in the cluster $\mathscr{C}$, conditional on the event $\alpha_i = n$, the type t_n is generated according to the distribution $\mathbb{H}_{\alpha_i}$ independently of the history $t_1, \ldots, t_{n-1}$ and $B(1), \ldots, B(n)$.*

The assumption above represents a nontrivial technical restriction on the random agents that form a cluster. If the types are i.i.d., then any stopping time will satisfy the assumption, but if the types are correlated, the condition needs to be verified. We still have a great deal of flexibility to represent different network structures. Examples 2 and 3 demonstrate two ways to represent essentially the same network structure, both of which comply with Assumption 4.

Example 2. Suppose types are i.i.d., and suppose the network topology is such that any agent n with type $t_n \geqslant \frac{1}{2}$ observes any previous agent m with $t_m \geqslant \frac{1}{2}$ and any agent n with type $t_n < \frac{1}{2}$ similarly observes any previous agent m with $t_m < \frac{1}{2}$. The sequence of stopping times $\{\alpha_i\}_{i\in\mathbb{N}}$, where α_i is the ith agent with type at least $\frac{1}{2}$, forms a cluster, and the agents with types below $\frac{1}{2}$ form another.

In the example above, a cluster of agents forms according to realized types; neighborhood realizations are correlated such that high type agents link to all other high type agents. A similar clustering of types can be achieved in a network with deterministic neighborhoods and correlated types instead of correlated neighborhoods.

Example 3. Let $\mathbb{H}_-$ and $\mathbb{H}_+$ denote the distribution $\mathbb{H}$ conditional on a type realization less than or at least $\frac{1}{2}$, respectively. Suppose agents are partitioned into two disjoint fixed subsequences $C = \{c_i\}_{i\in\mathbb{N}}$ and $D = \{d_i\}_{i\in\mathbb{N}}$ with $c_1 = 1$. Conditional on $t_1 < \frac{1}{2}$, agents in C realize types independently according to $\mathbb{H}_-$ with agents in D realizing types independently from $\mathbb{H}_+$; if $t_1 \geqslant \frac{1}{2}$, the type distributions for the subsequences are switched. Let the network structure be deterministic, with all agents in C observing all previous agents in C, and likewise for the subsequence D. The members of C form one cluster, and the members of D form another.

This section's main finding is a sufficient condition for learning within a cluster $\{\alpha_i\}_{i\in\mathbb{N}}$. We say that *asymptotic learning occurs within a cluster* $\{\alpha_i\}_{i\in\mathbb{N}}$ if we have

$$\lim_{i\to\infty} \mathbb{P}_\sigma(x_{\alpha_i} = \theta) = 1.$$

We require the cluster satisfies two properties.

Definition 8. A cluster is *identified* if there exists a family of neighbor choice functions $\{\gamma_i^k\}_{i,k\in\mathbb{N}}$ such that for each k, we have

$$\lim_{i\to\infty} \mathbb{Q}(\gamma_i^k(B(\alpha_i)) = \alpha_k) = 1.$$

That is, a cluster is identified only if, in the limit as the network grows large, agents in a given cluster can identify which other individuals belong to the same cluster.

A cluster is *uniformly diverse* if the following two conditions hold.

(a) For any interval $I \subseteq (0, 1)$ there exists a constant $\epsilon_I > 0$ such that $\mathbb{P}(t_{\alpha_i} \in I) \geqslant \epsilon_I$ for infinitely many i.

(b) There exists a finite measure μ such that the Radon-Nykodim derivative $\mathrm{d}\mu/\mathrm{d}\mathbb{H}_{\alpha_i} \geqslant 1$ almost surely for all i.

The first part of the uniform diversity condition generalizes the notion of unbounded preferences. It is unnecessary for any particular member of a cluster to have a type drawn from a distribution with full support; we simply need infinitely many members of the cluster to fall in any given interval of types. The second part of the condition is technical, requiring the existence of a finite measure on $(0,1)$ that dominates the type distribution of all agents in the cluster. Without this condition, it would be possible to construct pathological situations by having the type distributions $\mathbb{H}_{\alpha_i}$ concentrate increasing mass near the end points of $(0,1)$ as i grows. Together, these conditions are sufficient for asymptotic learning in clusters regardless of the signal structure or the structure of the rest of the network.

Theorem 4. *Let Assumption 4 hold. Suppose $\{\alpha_i\}_{i\in\mathbb{N}}$ is an identified, uniformly diverse cluster. Asymptotic learning obtains within the cluster.*

Proof. See appendix. ☐

With enough preference diversity, there is sufficient experimentation within a cluster to allow full information aggregation, even with bounded private beliefs. A key and novel element in our characterization is the concept of identification. Agents need to be able to tell with high probability who the other members of their cluster are to ensure asymptotic learning within the cluster. We can immediately apply this result to any network without homophily where the entire network forms a cluster.

Corollary 1. *Suppose preferences are unbounded, and the types $\{t_n\}_{n\in\mathbb{N}}$ form an i.i.d. sequence that is also independent of the neighborhoods. If for each m we have*

$$\lim_{n\to\infty}\mathbb{Q}(m\in B(n))=1, \tag{3}$$

asymptotic learning obtains.

Proof. Define the stopping indices $\alpha_i=i$, and define the neighbor choice functions

$$\gamma_n^i(B(n))=\begin{cases} i & \text{if } i\in B(n)\\ \varnothing & \text{otherwise.}\end{cases}$$

Theorem 4 applies. ☐

This corollary substantially generalizes the asymptotic learning result of Goeree et al. (2006) for the complete network; it obtains a positive result that applies even to networks with correlated neighborhoods and gaps, as long as Equation (3) is satisfied. Both identification and uniform diversity play a vital role in Theorem 4. Without identification, agents may be unable to interpret the information available to them, and without uniform diversity, inefficient herd behavior may appear.

Example 4. Suppose the types $\{t_n\}_{n\in\mathbb{N}}$ form an i.i.d. sequence, $\mathbb{H}$ is symmetric around $\frac{1}{2}$, and private beliefs are bounded. Consider a network in which agents sort themselves into three clusters C_i for $i\in\{1,2,3\}$. Let $1\in C_1$, and let c_n denote the number of agents in C_1 with index less than n. Each agent n is contained in C_1 with probability $1/2^{c_n}$, and these agents observe the entire history of action. Otherwise, if $t_n<\frac{1}{2}$ we have $n\in C_2$, and n observes all prior agents in C_2 and only those agents for sure. Likewise, if $t_n\geqslant\frac{1}{2}$ we have $n\in C_3$, and n observes all prior agents in C_3 and only those agents for sure.

The cluster C_1 is uniformly diverse, but unidentified, and learning fails. The clusters C_2 and C_3 are identified, but the lack of diversity means informational cascades occur with positive probability; asymptotic learning fails here too.

In Example 4, agents sort themselves into three clusters: a small uniformly diverse cluster, a large cluster of low types, and a large cluster of high types. Since private beliefs are bounded, if the low type cluster approaches a social belief close to one, no agent will ever deviate from action 1. Likewise, if the high type cluster approaches a social belief close to zero, no agent will deviate from action 0. Both of these clusters have a positive probability of inefficient herding. If this occurs, the agents in C_1 observe a history in which roughly half of the agents choose each action, and this is uninformative. Since the other members of C_1 make up a negligible share of the population and cannot be identified, their experimentation fails to provide any useful information to agents in C_1.

This example clearly suggests a negative role for homophily in a dense network. Densely connected networks offer an opportunity to overcome the difficulty presented by bounded private beliefs, but if homophily is strong enough that no identified cluster is uniformly diverse, the benefits of increased connectivity are lost. However, an inefficient herding outcome requires extreme isolation from types with different preferences, and even slight exposure is enough to recover a positive long-run result.

Consider the network introduced in Example 3. For a given marginal type distribution $\mathbb{H}$ with unbounded preferences, let $\mathbb{H}_-$ denote the distribution $\mathbb{H}$ conditional on a type realization less than $\frac{1}{2}$, and $\mathbb{H}_+$ the distribution $\mathbb{H}$ conditional on a type realization of at least $\frac{1}{2}$. Agents are partitioned into two disjoint deterministic subsequences $C=\{c_i\}_{i\in\mathbb{N}}$ and $D=\{d_i\}_{i\in\mathbb{N}}$ with $r_1=1$. Conditional on $t_1<\frac{1}{2}$, agents in C realize types independently according to $\mathbb{H}_-$, while agents in D realize types independently according to $\mathbb{H}_+$, with the distributions switched conditional on $t_1\geqslant\frac{1}{2}$. We assume that $c_i\in B(c_j)$ and $d_i\in B(d_j)$ for all $i<j$, so clearly C and D are identified clusters. If the two clusters are totally isolated from one another, meaning agents in C never observe an agent in D, and vice versa, an informational cascade can occur if the private beliefs are bounded. However, the outcome changes drastically if agents in each

cluster occasionally observe a member of the other cluster. Only a small amount of bidirectional communication between the clusters is needed to disrupt herding, leading to asymptotic learning in the entire network.

Proposition 4. *Assume private beliefs are bounded. Consider a network topology* $\mathbb{Q}$ *with deterministic neighborhoods that is partitioned into two clusters as described above. Asymptotic learning occurs in both clusters if and only if*

$$\sup\{i: \exists\, j,\ c_i \in B(d_j)\} = \sup\{i: \exists\, j,\ d_i \in B(c_j)\} = \infty.$$

Proof. See appendix. □

The key requirement in this statement is another type of minimal connectivity. For each cluster, the observations of the other cluster cannot be confined to a finite set. This means that periodically there is some agent in each cluster who makes a new observation of the other. This single observation serves as an aggregate statistic for an entire cluster's accumulated social information, so it holds substantial weight even though it is generated by a cluster of agents with significantly different preferences. This suggests that informational cascades in two cluster networks are difficult to sustain indefinitely, even in the presence of homophily, though if observations across groups are very rare, convergence may be quite slow.

5. Conclusions

Preference heterogeneity and homophily are pervasive in real-world social networks, so understanding their effects is crucial if we want to bring social learning theory closer to how people make decisions in practice. Preference diversity, homophily, and network structure impact learning in complex ways, and the results presented in this paper provide insight on how these three phenomena interact. Underlying our analysis is the idea that learning occurs through at least two basic mechanisms, one based on the improvement principle and the other based on a law-of-large-numbers effect. The structure of the network dictates which mechanisms are available to the agents, and the two mechanisms are affected differently by preference heterogeneity and homophily.

The improvement principle is one way people can learn about the world. An individual can often combine her private information with that provided by her neighbor's action to arrive at a slightly better decision than her neighbor. In sparsely connected networks, the improvement principle is the primary means through which learning occurs. However, this mechanism is, to some extent, quite fragile: it requires the possibility of extremely strong signals and a precise understanding of the decisions made by individual neighbors. Preference heterogeneity introduces noise in the chain of observations, and our results show this noise challenges the operation of the improvement principle. The problems that arise are due to both uncertainty regarding how a neighbor makes her decision and to different tradeoffs neighbors face. Homophily ameliorates both issues, making it unambiguously helpful to this type of learning.

Dense connectivity allows a far more robust learning mechanism to operate. Having strong information from individual neighbors becomes less important when one can observe the actions of many neighbors. As long as each observation provides some new information, the law of large numbers ensures that learning succeeds. Diverse preferences provide a degree of independence to the observations, facilitating this means of learning. Homophily, since it reduces preference diversity within a neighborhood, has the potential to interfere, but our results suggest this learning mechanism is resilient enough to be unaffected unless homophily is particularly extreme.

Overall, we find that preference heterogeneity and homophily play positive, complementary roles in social learning as typical complex networks include both sparse and dense components. If private signals are of bounded strength, preference diversity is necessary to ensure learning in the dense parts of the network via the law-of-large-numbers mechanism. If homophily is also present, the information accumulated in the dense parts of the network can spread to the sparse parts via the improvement principle mechanism. The combination of preference heterogeneity and homophily should generally benefit information aggregation in complex, realistic networks.

Acknowledgments

This paper incorporates results from "Heterogeneity and Social Learning," a working paper with Daron Acemoglu and Asu Ozdaglar. The authors thank Daron and Asu for their permission to use earlier results here. The authors are grateful to Alireza Tahbaz-Salehi for very helpful conversations on early versions of this work, and to Sinan Aral for guidance on the homophily literature. The authors gratefully acknowledge financial support from the NET Institute, www.NETinst.org. The authors also thank the National Science Foundation [Grant CCF-1216004] for its financial support. Any errors are those of the authors.

Appendix

We first prove part (b) of Theorem 2 from §3.1 and establish a version of the improvement principle to prove the results in §3.2. We conclude by applying the theory of martingales to establish results from §4.

Proof of Theorem 2 Part (b)

Our first task is to show the assumption on $\mathbb{G}_0$ implies a similar condition on $\mathbb{G}_1$:

$$\lim_{r\to 0}\frac{\mathbb{G}_0(r)}{r^{K-1}} = c \quad\Longrightarrow\quad \lim_{r\to 0}\frac{\mathbb{G}_1(r)}{r^{K}} = c\left(1-\frac{1}{K}\right).$$

An application of Lemma 1 and integration by parts gives

$$\begin{aligned}\mathbb{G}_1(r) &= \int_0^r \mathrm{d}\mathbb{G}_1(s) = \int_0^r \frac{s}{1-s}\,\mathrm{d}\mathbb{G}_0(s)\\ &= \frac{r}{1-r}\mathbb{G}_0(r) - \int_0^r \frac{1}{(1-s)^2}\mathbb{G}_0(s)\,\mathrm{d}s.\end{aligned}$$

Using our assumption on $\mathbb{G}_0$, given any $\epsilon > 0$, we may find r_ϵ such that for all $r \leqslant r_\epsilon$,

$$c(1-\epsilon)r^{K-1} \leqslant \mathbb{G}_0(r) \leqslant c(1+\epsilon)r^{K-1}.$$

Thus, for any $r \leqslant r_\epsilon$ we have

$$\begin{aligned}\int_0^r \frac{c(1-\epsilon)s^{K-1}}{(1-s)^2}\,\mathrm{d}s &\leqslant \int_0^r \frac{1}{(1-s)^2}\mathbb{G}_0(s)\,\mathrm{d}s\\ &\leqslant \int_0^r \frac{c(1+\epsilon)s^{K-1}}{(1-s)^2}\,\mathrm{d}s.\end{aligned}$$

Now compute,

$$\begin{aligned}\int_0^r \frac{s^{K-1}}{(1-s)^2}\,\mathrm{d}s &= \frac{r^K}{1-r} - (K-1)\int_0^r \frac{s^{K-1}}{(1-s)}\,\mathrm{d}s\\ &= \frac{r^K}{1-r} - (K-1)\int_0^r \sum_{i=0}^{\infty} s^{K-1+i}\,\mathrm{d}s\\ &= \frac{r^K}{1-r} - (K-1)\frac{r^K}{K} + O(r^{K+1}).\end{aligned}$$

It follows that for any $r \leqslant r_\epsilon$,

$$\begin{aligned}&\frac{r}{1-r}\mathbb{G}_0(r) - c(1+\epsilon)\left(\frac{r^K}{1-r} - (K-1)\frac{r^K}{K}\right) + O(r^{K+1})\\ &\quad \leqslant \mathbb{G}_1(r), \quad \text{and}\\ &\mathbb{G}_1(r) \leqslant \frac{r}{1-r}\mathbb{G}_0(r) - c(1-\epsilon)\left(\frac{r^K}{1-r} - (K-1)\frac{r^K}{K}\right)\\ &\quad + O(r^{K+1}).\end{aligned}$$

Dividing through by r^K and letting r go to zero, we have

$$c\left(1 - \frac{1+\epsilon}{K}\right) \leqslant \lim_{r\to 0}\frac{\mathbb{G}_1(r)}{r^K} \leqslant c\left(1 - \frac{1-\epsilon}{K}\right)$$

for any ϵ, proving the result.

We now show that

$$\int_0^1 \frac{t^{K-1}(K-(2K-1)t)}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t) < 0,$$

implies the existence of $\epsilon > 0$ such that $\mathscr{Z}(y) < y$ for all $y \in [1-\epsilon, 1)$. Choose a function $h(\epsilon)$ such that $h(\epsilon) > 0$ whenever $\epsilon > 0$, and

$$\lim_{\epsilon\to 0}\int_0^{1-\epsilon} \frac{h(\epsilon)}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t) = 0.$$

For a given $\epsilon > 0$, the argument $R^y_{1-\epsilon}$ goes to zero as y approaches 1, so for sufficiently large y and all $t \leqslant 1-\epsilon$,

$$\begin{aligned}&(1-y)\mathbb{G}_0(R^y_t) - y\mathbb{G}_1(R^y_t)\\ &\quad \leqslant c\left[(1-y)(1+h(\epsilon))(R^y_t)^{K-1} - y(1-h(\epsilon))\frac{K-1}{K}(R^y_t)^K\right]\\ &\quad = (1-y)^K\Bigg(ct^{K-1}\bigg[(1+h(\epsilon))(t(1-y)+y(1-t))\\ &\qquad\qquad - yt(1-h(\epsilon))\frac{K-1}{K}\bigg]\Bigg)\cdot((t(1-y)+y(1-t))^K)^{-1}.\end{aligned}$$

For y sufficiently close to 1, the portion of the improvement term

$$\int_0^{1-\epsilon} (1-y)\mathbb{G}_0(R^y_t) - y\mathbb{G}_1(R^y_t)\,\mathrm{d}\mathbb{H}(t) \tag{4}$$

is then bounded above by

$$\begin{aligned}&\frac{c}{K}(1-y)^K\int_0^{1-\epsilon}(t^{K-1}[K(1+h(\epsilon))(t(1-y)+y(1-t))\\ &\qquad\qquad - (K-1)(1-h(\epsilon))ty])\\ &\quad \cdot((t(1-y)+y(1-t))^K)^{-1}\,\mathrm{d}\mathbb{H}(t).\end{aligned} \tag{5}$$

For $t \in [0, 1-\epsilon]$, the integrand in Equation (5) converges uniformly as y approaches 1 to

$$\frac{t^{K-1}(K-(2K-1)t + h(\epsilon)(K-t))}{(1-t)^K}.$$

Therefore, for any $\epsilon' > 0$ and y sufficiently close to 1,

$$\begin{aligned}&\int_0^{1-\epsilon}(t^{K-1}[K(1+h(\epsilon))(t(1-y)+y(1-t)) - (K-1)\\ &\qquad \cdot(1-h(\epsilon)ty)])\cdot((t(1-y)+y(1-t))^K)^{-1}\,\mathrm{d}\mathbb{H}(t)\\ &\leqslant \epsilon' + \int_0^{1-\epsilon}\frac{t^{K-1}(K-(2K-1)t+h(\epsilon)(K-t))}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t).\end{aligned}$$

As ϵ approaches zero, this converges to

$$\epsilon' + \int_0^1 \frac{t^{K-1}(K-(2K-1)t)}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t). \tag{6}$$

If the integral in Equation (6) is negative, then we can choose some $\epsilon' > 0$, $\epsilon > 0$, and a corresponding y^* such that for any $y \in [y^*, 1)$, the integral in Equation (5) is negative. Therefore,

$$\int_0^{1-\epsilon}\left[(1-y)\mathbb{G}_0(R^y_t) - y\mathbb{G}_1(R^y_t)\right]\mathrm{d}\mathbb{H}(t) < 0$$

for any $y \in [y^*, 1)$.

To complete the proof, we show that for a sufficiently small choice of ϵ, there exists $y_\epsilon < 1$ such that

$$\int_{1-\epsilon}^1\left[(1-y)\mathbb{G}_0(R^y_t) - y\mathbb{G}_1(R^y_t)\right]\mathrm{d}\mathbb{H}(t) < 0$$

for all $y \in [y_\epsilon, 1)$. Thus, for $y \in [\max(y_\epsilon, y^*), 1)$, the entire improvement term is negative. Again using Lemma 1 and integration by parts we have

$$\begin{aligned}&(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\\ &\quad= \int_0^{R_t^y} (1-y)\,\mathrm{d}\mathbb{G}_0(r) - y\,\mathrm{d}\mathbb{G}_1(r) = \int_0^{R_t^y} \frac{1-y-r}{r}\,\mathrm{d}\mathbb{G}_1(r)\\ &\quad= \frac{y(1-2t)}{t}\mathbb{G}_1(R_t^y) + \int_0^{R_t^y} \frac{(1-y)\mathbb{G}_1(r)}{r^2}\,\mathrm{d}r.\end{aligned}$$

Since $\mathbb{G}_1$ is increasing and bounded, there exist constants $0 < \underline{c} < \bar{c}$, such that

$$\underline{c}r^K \leqslant \mathbb{G}_1(r) \leqslant \bar{c}r^K$$

for all $r \in [0, 1]$. Therefore, for any $t > \frac{1}{2}$ we have

$$\begin{aligned}&\frac{y(1-2t)}{t}\mathbb{G}_1(R_t^y) + \int_0^{R_t^y} \frac{(1-y)\mathbb{G}_1(r)}{r^2}\,\mathrm{d}r\\ &\quad\leqslant \frac{\underline{c}y(1-2t)}{t}[R_t^y]^K + \frac{\bar{c}(1-y)}{K-1}[R_t^y]^{K-1}.\end{aligned}$$

The right-hand side is negative whenever $t > y(\underline{c} \cdot (K-1) + \bar{c})/(2y\underline{c}(K-1) + (2y-1)\bar{c})$. This threshold is decreasing in y, with a limiting value that is strictly less than 1. Therefore, for any $\epsilon < 1 - \underline{c}(K-1) + \bar{c}/(2\underline{c} \cdot (K-1) + \bar{c})$, we can find y_ϵ such that

$$(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y) < 0$$

for all $t \geqslant 1-\epsilon$ and all $y \in [y_\epsilon, 1)$, completing the proof. □

Proof of Proposition 1

Fixing an $\epsilon > 0$, we can bound the improvement term as

$$\begin{aligned}&\int_0^1 \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\Big]\,\mathrm{d}\mathbb{H}(t)\\ &\quad\geqslant \int_0^{1-\epsilon} \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\Big]\,\mathrm{d}\mathbb{H}(t)\\ &\qquad - \int_{1-\epsilon}^1 y\mathbb{G}_1(R_t^y)\,\mathrm{d}\mathbb{H}(t).\end{aligned}$$

Carrying out a similar exercise as in the previous proof, we can show that for any ϵ', we can find y^* such that

$$\begin{aligned}&\int_0^{1-\epsilon} \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\Big]\,\mathrm{d}\mathbb{H}(t)\\ &\quad\geqslant \frac{c}{K}(1-y)^K \left(\int_0^1 \frac{t^{K-1}(K-(2K-1)t)}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t) - \epsilon'\right)\end{aligned}$$

for all $y \in [y^*, 1)$. There exists a constant $\bar{c}$ such that $\mathbb{G}_1(r) \leqslant \bar{c}r^K$ for all r, so we have

$$\begin{aligned}&\int_{1-\epsilon}^1 y\mathbb{G}_1(R_t^y)\,\mathrm{d}\mathbb{H}(t)\\ &\quad\leqslant \bar{c}(1-y)^K \int_{1-\epsilon}^1 y\frac{t^K}{(t(1-y)+y(1-t))^K}\,\mathrm{d}\mathbb{H}(t)\\ &\quad\leqslant \bar{c}(1-y)^K \int_{1-\epsilon}^1 \frac{t^K}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t)\end{aligned}$$

for all $\epsilon < \frac{1}{2}$. Since $\int_0^1 t^{K-1}(K-(2K-1)t)/(1-t)^K\,\mathrm{d}\mathbb{H}(t)$ is positive, we know that $\int_0^1 t^K/(1-t)^K\,\mathrm{d}\mathbb{H}(t)$ is finite. Hence, for any ϵ', we can choose ϵ such that $(K\bar{c}/c) \cdot \int_{1-\epsilon}^1 t^K/(1-t)^K\,\mathrm{d}\mathbb{H}(t) < \epsilon'$.

Take $\epsilon' < \frac{1}{3}\int_0^1 t^{K-1}(K-(2K-1)t)/(1-t)^K\,\mathrm{d}\mathbb{H}(t)$, choose ϵ so that $(K\bar{c}/c)\int_{1-\epsilon}^1 t^K/(1-t)^K\,\mathrm{d}\mathbb{H}(t) < \epsilon'$, and find y^* such that

$$\begin{aligned}&\int_0^{1-\epsilon} \Big[(1-y)\mathbb{G}_0(R_t^y) - y\mathbb{G}_1(R_t^y)\Big]\,\mathrm{d}\mathbb{H}(t)\\ &\quad\geqslant \frac{c}{K}(1-y)^K \left(\int_0^1 \frac{t^{K-1}(K-(2K-1)t)}{(1-t)^K}\,\mathrm{d}\mathbb{H}(t) - \epsilon'\right)\end{aligned}$$

for all $y \in [y^*, 1)$. Then for all $y \in [y^*, 1)$ there is a constant C such that $\mathscr{Z}(y) > y + C(1-y)^K$. The analysis of Lobel et al. (2009) now implies the result. □

We use the following version of the improvement principle to establish positive learning results.

Lemma 4 (Improvement Principle). *Let Assumption 3 hold. Suppose there exists a sequence of neighbor choice functions* $\{\gamma_n\}_{n\in\mathbb{N}}$ *and a continuous, increasing function* $\mathscr{Z}\colon [1, 3/2] \to [1, 3/2]$ *with the following properties*:

(a) *The chosen neighbor topology* $\mathbb{Q}_\gamma$ *features expanding subnetworks.*

(b) *We have* $\mathscr{Z}(y) > y$ *for any* $y < \frac{3}{2}$.

(c) *For any* $\epsilon > 0$, *we have*

$$\begin{aligned}\lim_{n\to\infty} \mathbb{P}_\sigma\{&\mathbb{E}_\sigma[u(t_n, x_n, \theta) \mid \gamma_n(B(n))]\\ &< \mathscr{Z}(\mathbb{E}_\sigma[u(t_{\gamma_n(B(n))}, x_{\gamma_n(B(n))}, \theta)]) - \epsilon\} = 0.\end{aligned}$$

Asymptotic learning obtains.

Proof. Note that asymptotic learning is equivalent to

$$\lim_{n\to\infty} \mathbb{E}_\sigma[u(t_n, x_n, \theta)] = \frac{3}{2}.$$

We construct two sequences $\{\eta_k\}$ and $\{\phi_k\}$ with the property that for all $k \geqslant 1$ and $n \geqslant \eta_k$, $\mathbb{E}_\sigma[u(t_n, x_n, \theta)] \geqslant \phi_k$. Upon showing that $\lim_{k\to\infty}\phi_k = \frac{3}{2}$, we shall have our result. Using our assumptions, for any integer K and $\epsilon > 0$, we can find a positive integer $N(K, \epsilon)$ such that

$$\mathbb{Q}(\gamma_n(B(n)) = \{m\},\, m < K) < \frac{\epsilon}{2}, \quad \text{and}$$

$$\begin{aligned}\mathbb{P}_\sigma\{&\mathbb{E}_\sigma[u(t_n, x_n, \theta) \mid \gamma_n(B(n))]\\ &< \mathscr{Z}(\mathbb{E}_\sigma[u(t_{\gamma_n(B(n))}, x_{\gamma_n(B(n))}, \theta)]) - \epsilon\} < \frac{\epsilon}{2}\end{aligned}$$

for all $n \geqslant N(K, \epsilon)$. Set $\eta_1 = 1$ and $\phi_1 = 1$, and let $\epsilon_k \equiv \frac{1}{2}(1 + \mathscr{Z}(\phi_k) - \sqrt{1 + 2\phi_k + \mathscr{Z}(\phi_k)^2})$; we define the rest of the sequences recursively by

$$\eta_{k+1} = N(\eta_k, \epsilon_k), \qquad \phi_{k+1} = \frac{\mathscr{Z}(\phi_k) + \phi_k}{2}.$$

Given the assumptions on $\mathscr{Z}(\phi_k)$, these sequences are well defined.

Proceed by induction to show that $\mathbb{E}_\sigma[u(t_n, x_n, \theta)] \geqslant \phi_k$ for all $n \geqslant \eta_k$. The base case $k = 1$ trivially holds because an agent may always choose the action preferred a priori according to her type, obtaining expected utility at least 1. Considering $n \geqslant \eta_{k+1}$ we have

$$\begin{aligned}&\mathbb{E}_\sigma[u(t_n, x_n, \theta)]\\ &\quad\geqslant \sum_{m<n} \mathbb{E}_\sigma[u(t_n, x_n, \theta) \mid \gamma_n(B(n)) = m]\mathbb{Q}(\gamma_n(B(n)) = m)\\ &\quad\geqslant (1-\epsilon_k)(\mathcal{Z}(\phi_k) - \epsilon_k) = \phi_{k+1}.\end{aligned}$$

To see that ϕ_k converges to $\frac{3}{2}$, note the definition implies $\{\phi_k\}$ is a nondecreasing, bounded sequence, so it has a limit ϕ^*. Since $\mathcal{Z}$ is continuous, the limit must be a fixed point of $\mathcal{Z}$. The only fixed point is $\frac{3}{2}$, finishing the proof. □

Proof of Theorem 3

Our goal is to construct a function $\mathcal{Z}$ on which to apply Lemma 4. We begin by characterizing the decision $\tilde{x}_n$, defined as

$$\tilde{x}_n = \operatorname*{arg\,max}_{y \in \{0,1\}} \mathbb{E}_\sigma[u(t_n, y, \theta) \mid \gamma_n(B(n)),\ x_{\gamma_n(B(n))}].$$

The decision $\tilde{x}_n$ is a decision based on a coarser information set than what agent n actually has access to. Therefore, the utility derived from this decision provides a lower bound on agent n's utility in equilibrium.

Suppose $\gamma_n(B(n)) = m$. To simplify notation, we define $P_{it} = \mathbb{P}_\sigma(x_m = i \mid \theta = i,\ t_m = t)$ and $P_i = \int_0^1 P_{it}\, d\mathbb{H}_{t_m \mid t_n}(t)$, where $\mathbb{H}_{t_m \mid t_n}$ denotes the distribution of t_m conditioned on the realized t_n and the event $\gamma_n(B(n)) = m$. We further define E_{it} as the expected utility of agent m given that $t_m = t$ and E_i analogously to P_i. These quantities are related by

$$E_{0t} = 2tP_{0t} + 1 - t, \qquad E_{1t} = 2(1-t)P_{1t} + t. \tag{7}$$

Note that P_{it} and E_{it} are constants, independent of the realization of t_n, while P_i and E_i are random variables as functions of t_n through the conditional distribution $\mathbb{H}_{t_m \mid t_n}$. Define the thresholds

$$L_{t_n} = \frac{t_n(1-P_0)}{t_n(1-P_0) + (1-t_n)P_1}, \qquad U_{t_n} = \frac{t_n P_0}{t_n P_0 + (1-t_n)(1-P_1)}.$$

For the remainder of the proof we suppress the subscript t_n to simplify notation. An application of Bayes' rule shows that

$$\tilde{x}_n = \begin{cases} 0 & \text{if } p_n < L \\ x_m & \text{if } p_n \in (L, U) \\ 1 & \text{if } p_n > U. \end{cases} \tag{8}$$

Fixing t_n, the expected payoff from the decision $\tilde{x}_n$ is easily computed as

$$\begin{aligned}&\frac{1}{2}[\mathbb{G}_0(L)(1+t_n) + (\mathbb{G}_0(U) - \mathbb{G}_0(L))(2t_nP_0 + 1 - t_n)\\ &\quad + (1-\mathbb{G}_0(U))(1-t_n) + (1-\mathbb{G}_1(U))(2-t_n)\\ &\quad + (\mathbb{G}_1(U) - \mathbb{G}_1(L))(2(1-t_n)P_1 + t_n) + \mathbb{G}_1(L)t_n],\end{aligned} \tag{9}$$

with the first line corresponding to the payoff when $\theta = 0$ and the second to the payoff when $\theta = 1$. An application of Lemma 1 provides the following inequalities.

$$\begin{aligned}&\mathbb{G}_0(L) \geqslant \frac{1-L}{L}\mathbb{G}_1(L) + \frac{L}{4}\mathbb{G}_1\left(\frac{L}{2}\right),\\ &1 - \mathbb{G}_1(U) \geqslant \frac{U}{1-U}(1-\mathbb{G}_0(U)) + \frac{1-U}{4}\left(1 - \mathbb{G}_1\left(\frac{1+U}{2}\right)\right).\end{aligned}$$

Substituting into Equation (9), we find

$$\begin{aligned}&\mathbb{E}_\sigma[u(t_n, \tilde{x}_n, \theta) \mid t_n,\ \gamma_n(B(n)) = m]\\ &\quad\geqslant \frac{1}{2} + t_nP_0 + (1-t_n)P_1 + t_n(1-P_0)\frac{L}{4}\mathbb{G}_1\left(\frac{L}{2}\right)\\ &\qquad + (1-t_n)(1-P_1)\frac{1-U}{4}\left(1 - \mathbb{G}_0\left(\frac{1+U}{2}\right)\right).\end{aligned} \tag{10}$$

We collectively refer to the first three terms above as the "base" terms, and the last two as the "improvement" terms. We focus first on the base terms. Using Equation (7) the base terms can be written as

$$\begin{aligned}&\frac{E_0 + E_1}{2} + \int_0^1 (t_n - t)(P_{0t} - P_{1t})\, d\mathbb{H}_{t_m \mid t_n}(t)\\ &\quad = \mathbb{E}_\sigma[u(t_m, x_m, \theta) \mid \gamma_n(B(n)) = m,\ t_n]\\ &\qquad + \int_0^1 (t_n - t)(P_{0t} - P_{1t})\, d\mathbb{H}_{t_m \mid t_n}(t).\end{aligned} \tag{11}$$

For a given $\epsilon > 0$, define $p_\epsilon = \mathbb{P}(|t_m - t_n| > \epsilon \mid \gamma_n(B(n)) = m)$. Integrating over t_n and using Assumption 3, for any $\epsilon > 0$ we can bound the integrated base terms from below by

$$\begin{aligned}&\mathbb{E}_\sigma[u(t_m, x_m, \theta)] + \int_0^1\int_0^1 (s-t)(P_{0t} - P_{1t})\, d\mathbb{H}_{t_m \mid t_n}(t)\, d\mathbb{H}_{t_n}(s)\\ &\quad \geqslant \mathbb{E}_\sigma[u(t_m, x_m, \theta)] - \epsilon - p_\epsilon.\end{aligned} \tag{12}$$

Moving to the improvement terms, note that

$$P_0 = \int_0^1 P_{0t}\, d\mathbb{H}_{t_m \mid t_n}(t) = \int_0^1 \frac{E_{0t} - (1-t)}{2t}\, d\mathbb{H}_{t_m \mid t_n}(t).$$

The last integrand is Lipschitz continuous in t for t bounded away from zero. Therefore, for any ϵ with $0 < 2\epsilon < t_n$, we can find a constant c such that

$$\begin{aligned}&\int_0^1 \frac{E_{0t} - (1-t_n)}{2t_n}\, d\mathbb{H}_{t_m \mid t_n}(t) - c\epsilon - \frac{p_\epsilon}{t_n}\\ &\quad \leqslant \int_0^1 \frac{E_{0t} - (1-t)}{2t}\, d\mathbb{H}_{t_m \mid t_n}(t)\\ &\quad \leqslant \int_0^1 \frac{E_{0t} - (1-t_n)}{2t_n}\, d\mathbb{H}_{t_m \mid t_n}(t) + c\epsilon + \frac{p_\epsilon}{t_n}.\end{aligned}$$

This is equivalent to

$$\frac{E_0 - (1-t_n)}{2t_n} - c\epsilon - \frac{p_\epsilon}{t_n} \leqslant P_0 \leqslant \frac{E_0 - (1-t_n)}{2t_n} + c\epsilon + \frac{p_\epsilon}{t_n}.$$

Similarly, we can find a constant c such that

$$\frac{E_1-t_n}{2(1-t_n)} - c\epsilon - \frac{p_\epsilon}{1-t_n} \leqslant P_1 \leqslant \frac{E_1-t_n}{2(1-t_n)} + c\epsilon + \frac{p_\epsilon}{1-t_n}.$$

Consider a modification of the improvement terms in Equation (10) where we replace P_0 by $(E_0-(1-t_n))/(2t_n)$ and P_1 by $(E_1-t_n)/(2(1-t_n))$, including in the definitions of L and U. Our work above, together with the continuity of the belief distributions, implies that the modified terms differ from the original improvement terms by no more than some function $\delta(\epsilon, p_\epsilon, t_n)$, where δ converges to zero as ϵ and p_ϵ approach zero together, and the convergence is uniform in t_n for t_n bounded away from 0 and 1. We can then bound the improvement terms by

$$t_n(1-P_0)\frac{L}{4}\mathbb{G}_1\left(\frac{L}{2}\right) \geqslant \frac{1}{8}\frac{(1+t_n-E_0)^2}{1-E_0+E_1}\mathbb{G}_1\left(\frac{1+t_n-E_0}{2(1-E_0+E_1)}\right) - \delta(\epsilon, p_\epsilon, t_n), \qquad (13)$$

$$(1-t_n)(1-P_1)\frac{1-U}{4}\left(1-\mathbb{G}_0\left(\frac{1+U}{2}\right)\right) \geqslant \frac{1}{8}\frac{(2-t_n-E_1)^2}{1+E_0-E_1}\left(1-\mathbb{G}_0\left(1-\frac{2-t_n-E_1}{2(1+E_0-E_1)}\right)\right) - \delta(\epsilon, p_\epsilon, t_n). \qquad (14)$$

Let $y^* = (E_0+E_1)/2$; we must have either $E_0 \leqslant \frac{2}{3}y^* + t_n$ or $E_1 \leqslant \frac{2}{3}y^* + 1 - t_n$ since $y^* \leqslant \frac{3}{2}$. The first term on the right-hand side of Equation (13) can be rewritten as

$$\frac{1}{8}\frac{(1+t_n-E_0)^2}{1-2E_0+2y^*}\mathbb{G}_1\left(\frac{1+t_n-E_0}{2(1-2E_0+2y^*)}\right),$$

which we note is decreasing in E_0 for $E_0 \leqslant \frac{2}{3}y^* + t_n$ (we have used that $y^* \geqslant \frac{1}{2}$). Therefore, if $E_0 \leqslant \frac{2}{3}y^* + t_n$, then Equation (13) is bounded below by

$$\frac{1}{16}\left(1-\frac{2}{3}y^*\right)^2\mathbb{G}_1\left(\frac{1-(2/3)y^*}{4}\right) - \delta(\epsilon, p_\epsilon, t_n). \qquad (15)$$

Similarly, if $E_1 \leqslant \frac{2}{3}y^* + 1 - t_n$, then Equation (14) is bounded below by

$$\frac{1}{16}\left(1-\frac{2}{3}y^*\right)^2\left(1-\mathbb{G}_0\left(1-\frac{1-(2/3)y^*}{4}\right)\right) - \delta(\epsilon, p_\epsilon, t_n). \qquad (16)$$

To simplify notation, define

$$Z(y^*) = \frac{1}{16}\left(1-\frac{2}{3}y^*\right)^2\min\left[\mathbb{G}_1\left(\frac{1-(2/3)y^*}{4}\right), \left(1-\mathbb{G}_0\left(1-\frac{1-(2/3)y^*}{4}\right)\right)\right].$$

Recall that $y^* = y^*(t_n)$ is a function of t_n through E_0 and E_1. Using Equations (15) and (16), we can integrate over t_n to bound the contribution of the improvement terms to agent n's utility. Since the improvement terms are nonnegative, we can choose any $\epsilon' > 0$ and restrict the range of integration to obtain a lower bound of

$$\int_{\epsilon'}^{1-\epsilon'}\Big[Z(y^*(t)) - 2\delta(\epsilon, p_\epsilon, t)\Big]\,\mathrm{d}\mathbb{H}(t).$$

Now define $y = \mathbb{E}_\sigma[u(t_m, x_m, \theta)]$ and note that $y = \int_0^1 y^*(t)\,\mathrm{d}\mathbb{H}(t)$. Observe since $y^* \geqslant \frac{1}{2}$, this implies $\mathbb{P}_\sigma(y^* \leqslant \frac{3}{4} + y/2) \geqslant (3-2y)/(1+2y)$. Choosing ϵ' sufficiently small we can bound the improvement terms below by

$$\frac{3-2y}{2(1+2y)}Z(y) - \int_{\epsilon'}^{1-\epsilon'} 2\delta(\epsilon, p_\epsilon, t)\,\mathrm{d}\mathbb{H}(t). \qquad (17)$$

Finally, define $\mathscr{Z}(y) = y + ((3-2y)/2(1+2y))Z(\frac{3}{4} + y/2)$. Combining Equations (10), (12), and (17), and using that Z is nonincreasing, we have

$$\mathbb{E}_\sigma[u(t_n, x_n, \theta) \mid \gamma_n(B(n)) = m] \geqslant \mathscr{Z}(y) - \epsilon - p_\epsilon - 2\int_{\epsilon'}^{1-\epsilon'} 2\delta(\epsilon, p_\epsilon, t)\,\mathrm{d}\mathbb{H}(t)$$

for any $\epsilon > 0$ and some fixed $\epsilon' > 0$. The hypothesis of the theorem implies that for any $\epsilon > 0$, p_ϵ approaches zero as n grows without bound. Using the uniform convergence to 0 of δ for $t_n \in [\epsilon', 1-\epsilon']$, we see that $\mathscr{Z}$ satisfies the hypothesis of Lemma 4, completing the proof. □

Proof of Proposition 3

Since $\mathbb{Q}'$ is strongly homophilous, there exist neighbor choice functions $\{\gamma_n\}_{n\in\mathbb{N}}$ for which $\lim_{n\to\infty}\mathbb{Q}'(|t_n - t_{\gamma_n(B(n))}| > \epsilon) = 0$, and the improvement principle of Lemma 4 applies to establish learning along the chains of observations in $\mathbb{Q}'_\gamma$. We shall see the same neighbor choice functions can be used to establish learning in $\mathbb{Q}^{(\lambda)}$.

Considering an arbitrary agent n with $\gamma_n(B(n)) = m$, recall the decision thresholds L_{t_n} and U_{t_n} defined in the proof of Theorem 3. Here we let L_{t_n} and U_{t_n} denote the corresponding thresholds for the network $\mathbb{Q}$, and let L'_{t_n} and U'_{t_n} denote the thresholds for the network $\mathbb{Q}'$. As in the proof of Theorem 3, we suppress the subscript t_n from here on. A lower bound on the improvement agent n makes in network $\mathbb{Q}'$ can be obtained by using the suboptimal decision thresholds L and U. Since the last of the base terms (see Equation (11)) does not depend on the decision thresholds, the proof of Theorem 3 may be followed with essentially no changes to obtain an improvement function $\mathscr{Z}'$ satisfying the assumptions of Lemma 4. Crucially, this improvement function applies for decisions made via the suboptimal decision thresholds that are appropriate to the network $\mathbb{Q}$ as opposed to $\mathbb{Q}'$.

By assumption, we have an improvement function $\mathscr{Z}$ for the network $\mathbb{Q}$. Considering an arbitrary λ-mixture of $\mathbb{Q}$

and $\mathbb{Q}'$, we suppose that n uses the decision thresholds L and U, instead of the optimal thresholds for $\mathbb{Q}^{(\lambda)}$. It is immediately clear that

$$\mathscr{Z}^{(\lambda)} = \min_{\lambda \in [0,1]} \{\lambda \mathscr{Z} + (1-\lambda)\mathscr{Z}'\}$$

is an improvement function for the suboptimal decision rule. This provides a lower bound on the true improvement, so $\mathscr{Z}^{(\lambda)}$ is an improvement function for $\mathbb{Q}_{\gamma}^{(\lambda)}$. Lemma 4 applies. ☐

The following lemmas provide the essential machinery for the proof of Theorem 4. The first lemma shows that if a subsequence of decisions in a cluster provides information that converges to full knowledge of the state, the entire cluster must learn. The essence of the proof is that the expected social belief of agents in an identified cluster, conditional on the decisions of this subsequence, must also converge on the state θ. Since social beliefs are bounded between 0 and 1, the realized social belief of each agent in the cluster deviates from this expectation only with very small probability.

The second lemma establishes a lower bound on the amount of new information provided by the decision of an agent within a uniformly diverse cluster. Given a subset of the information available to that agent, we bound how much information the agent's decision conveys on top of that. Under the assumption of uniform diversity, we can show that either the agent acts on her own signal with positive probability, or she has access to significant additional information not contained in the subset we considered. This is a key technical innovation that allows us to apply martingale convergence arguments even when agents within a cluster do not have nested information sets.

LEMMA 5. *Let Assumption 4 hold; suppose* $\{\alpha_i\}_{i\in\mathbb{N}}$ *is a uniformly diverse, identified cluster. If there exists a subsequence* $\{\alpha_{i(j)}\}_{j\in\mathbb{N}}$ *such that the social beliefs* $\hat{q}_k = \mathbb{P}_\sigma(\theta = 1 \mid x_{\alpha_{i(j)}},\ j \leqslant k)$ *converge to* θ *almost surely, then asymptotic learning obtains within the cluster.*

PROOF. Consider the case where $\theta = 0$; the case $\theta = 1$ is analogous. Fixing any $\epsilon > 0$, by assumption there exists some random integer K_ϵ, which is finite with probability 1, such that $\hat{q}_k \leqslant \epsilon/3$ for all $k \geqslant K_\epsilon$. Fix a k large enough so that $\mathbb{P}(K_\epsilon > k) \leqslant \epsilon/3$. Since the cluster is identified, for large enough n there are neighbor choice functions $\gamma_{\alpha_n}^{(1)}, \ldots, \gamma_{\alpha_n}^{(k)}$ such that

$$\mathbb{P}_\sigma(\exists\, i \leqslant k \mid \gamma_{\alpha_n}^{(i)}(B(\alpha_n)) \neq \alpha_i) \leqslant \frac{\epsilon}{3}.$$

We then have

$$\begin{aligned}
&\mathbb{E}_\sigma[q_{\alpha_n} \mid x_{\gamma_{\alpha_n}^{(i)}(B(\alpha_n))},\ i \leqslant k,\ \theta = 0] \\
&\quad\leqslant \left(1 - \frac{\epsilon}{3}\right)\mathbb{E}_\sigma[\hat{q}_k \mid \theta = 0] + \frac{\epsilon}{3} \\
&\quad\leqslant \left(1 - \frac{\epsilon}{3}\right)^2 \mathbb{E}_\sigma[\hat{q}_k \mid k \geqslant K_\epsilon,\ \theta = 0] + \frac{2\epsilon}{3} \leqslant \epsilon.
\end{aligned}$$

Since the social belief q_{α_n} is bounded below by 0, we necessarily have

$$\mathbb{P}_\sigma(q_{\alpha_n} > \sqrt{\epsilon} \mid \theta = 0) \leqslant \sqrt{\epsilon}.$$

Therefore,

$$\begin{aligned}
&\mathbb{P}_\sigma(x_{\alpha_n} = 0 \mid \theta = 0) \\
&\geqslant (1 - \sqrt{\epsilon}) \int_0^1 \mathbb{P}_\sigma(x_{\alpha_n} = 0 \mid \theta = 0,\ q_{\alpha_n} = \sqrt{\epsilon},\ t_{\alpha_n} = t)\, \mathrm{d}\mathbb{H}_{\alpha_n}(t) \\
&= (1 - \sqrt{\epsilon}) \int_0^1 \mathbb{G}_0(R_t^{\sqrt{\epsilon}})\, \mathrm{d}\mathbb{H}_{\alpha_n}(t).
\end{aligned}$$

Condition (b) of uniform diversity ensures this last integral approaches 1 as ϵ approaches zero, completing the proof. ☐

LEMMA 6. *Let Assumption 4 hold. Let* α *be a stopping time within a cluster, and suppose* $\mathscr{I}$ *is a random variable taking values contained in agent* α*'s information set with probability* 1. *For a given realization* I *of* $\mathscr{I}$, *define* $\hat{q}_\alpha = \mathbb{P}_\sigma(\theta = 1 \mid \mathscr{I} = I)$, *and suppose there exist* $p, d > 0$ *such that*

$$\mathbb{H}_\alpha(R_{\bar{\beta}}^{1-\hat{q}_\alpha} - d) - \mathbb{H}_\alpha(R_{\underline{\beta}}^{1-\hat{q}_\alpha} + d) \geqslant p.$$

If for some $\epsilon > 0$ *we have* $\hat{q}_\alpha \geqslant \epsilon$, *then there exists some* $\delta > 0$ *such that*

$$\frac{\mathbb{P}_\sigma(x_\alpha = 0 \mid \mathscr{I} = I,\ \theta = 0)}{\mathbb{P}_\sigma(x_\alpha = 0 \mid \mathscr{I} = I,\ \theta = 1)} \geqslant 1 + \delta.$$

Similarly, if $\hat{q}_\alpha \leqslant 1 - \epsilon$, *then we can find* $\delta > 0$ *such that*

$$\frac{\mathbb{P}_\sigma(x_\alpha = 1 \mid \mathscr{I} = I,\ \theta = 1)}{\mathbb{P}_\sigma(x_\alpha = 1 \mid \mathscr{I} = I,\ \theta = 0)} \geqslant 1 + \delta.$$

PROOF. Fix a realization I of $\mathscr{I}$, and define

$$P_{qi} = \mathrm{d}\mathbb{P}_\sigma(q_\alpha = q \mid I,\ \theta = i).$$

We can rewrite the ratio

$$\begin{aligned}
&\frac{\mathbb{P}_\sigma(x_\alpha = 0 \mid I, \theta = 0)}{\mathbb{P}_\sigma(x_\alpha = 0 \mid I, \theta = 1)} \\
&\quad= \frac{\int_0^1 \mathbb{P}_\sigma(x_\alpha = 0 \mid I, \theta = 0, t_\alpha = t)\, \mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1 \mathbb{P}_\sigma(x_\alpha = 0 \mid I, \theta = 1, t_\alpha = t)\, \mathrm{d}\mathbb{H}_\alpha(t)} \\
&\quad= \frac{\int_0^1 \int_0^1 P_{q0} \mathbb{G}_0(R_t^q)\, \mathrm{d}q\, \mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1 \int_0^1 P_{q1} \mathbb{G}_1(R_t^q)\, \mathrm{d}q\, \mathrm{d}\mathbb{H}_\alpha(t)} \\
&\quad= \frac{\int_0^1 \int_0^1 P_{q1} \mathbb{G}_0(R_t^q) + (P_{q0} - P_{q1}) \mathbb{G}_0(R_t^q)\, \mathrm{d}q\, \mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1 \int_0^1 P_{q1} \mathbb{G}_1(R_t^q)\, \mathrm{d}q\, \mathrm{d}\mathbb{H}_\alpha(t)} \\
&\quad= \bigg(\int_0^1 \int_0^1 P_{q1} \mathbb{G}_0(R_t^q) + (P_{q0} - P_{q1})(\mathbb{G}_0(R_t^q) \\
&\qquad\quad - \mathbb{G}_0(R_t^{\hat{q}_\alpha}))\, \mathrm{d}q\, \mathrm{d}\mathbb{H}_\alpha(t)\bigg) \\
&\qquad \cdot \left(\int_0^1 \int_0^1 P_{q1} \mathbb{G}_1(R_t^q)\, \mathrm{d}q\, \mathrm{d}\mathbb{H}_\alpha(t)\right)^{-1}.
\end{aligned} \tag{18}$$

It follows from the definition of the social belief and an application of Bayes' rule that for $\mathbb{P}_\sigma$ almost all q we have

$$q = \left[1 + \left(\frac{1}{\hat{q}_\alpha} - 1\right)\frac{P_{q0}}{P_{q1}}\right]^{-1}.$$

In particular, if $q < \hat{q}_\alpha$, then $P_{q0} > P_{q1}$, and the same holds with the inequalities reversed. Now for any given $\epsilon' > 0$, define $q_{\epsilon'}^- < q_{\epsilon'}^+$ as the values given above when we take P_{q1}/P_{q0} equal to $1-\epsilon'$ and $1+\epsilon'$, respectively.

We shall consider three cases. First, suppose $\mathbb{P}_\sigma(q_{\epsilon'}^- \geqslant q_\alpha \mid I) \geqslant 1/3$. For $q \leqslant q_{\epsilon'}^-$ we have

$$\begin{aligned}(P_{q0} - P_{q1})&(\mathbb{G}_0(R_t^q) - \mathbb{G}_0(R_t^{\hat{q}_\alpha}))\\ &= P_{q0}\left(1 - \frac{P_{q1}}{P_{q0}}\right)(\mathbb{G}_0(R_t^q) - \mathbb{G}_0(R_t^{\hat{q}_\alpha}))\\ &\geqslant \epsilon' P_{q0}(\mathbb{G}_0(R_t^{q_{\epsilon'}^-}) - \mathbb{G}_0(R_t^{\hat{q}_\alpha})).\end{aligned}$$

Integrating over all such q and choosing ϵ' small enough we have

$$\begin{aligned}\int_0^1\int_0^1 &(P_{q0} - P_{q1})\mathbb{G}_0(R_t^q)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}_\alpha(t)\\ &\geqslant \frac{\epsilon'}{6}\int_0^1 (\mathbb{G}_0(R_t^{q_{\epsilon'}^-}) - \mathbb{G}_0(R_t^{\hat{q}_\alpha}))\,\mathrm{d}\mathbb{H}_\alpha(t)\\ &\geqslant \frac{p\epsilon'}{6}\min_{t\in[R_{\underline{\beta}}^{1-\hat{q}_\alpha}+d,\, R_{\bar{\beta}}^{1-\hat{q}_\alpha}-d]}(\mathbb{G}_0(R_t^{q_{\epsilon'}^-}) - \mathbb{G}_0(R_t^{\hat{q}_\alpha})) > 0,\end{aligned}$$

where the last inequality follows because $R_t^{\hat{q}_\alpha}$ is bounded strictly between $\underline{\beta}$ and $\bar{\beta}$ for the given range, and $\mathbb{G}_0$ has full support in $(\underline{\beta}, \bar{\beta})$. Since the denominator in Equation (18) is bounded by 1, and the first term of the numerator is at least as large by Lemma 1 part (d), the inequality above gives us our desired δ. For the second case, if $\mathbb{P}_\sigma(q_{\epsilon'}^+ \leqslant q_\alpha \mid \mathcal{I} = I) \geqslant 1/3$, then a similar exercise gives

$$\begin{aligned}\int_0^1\int_0^1 &(P_{q0} - P_{q1})\mathbb{G}_0(R_t^q)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}_\alpha(t)\\ &\geqslant \frac{\epsilon'}{6}\int_0^1 (\mathbb{G}_0(R_t^{\hat{q}_\alpha}) - \mathbb{G}_0(R_t^{q_{\epsilon'}^+}))\,\mathrm{d}\mathbb{H}_\alpha(t).\end{aligned}$$

For small ϵ', the required δ is forthcoming by the same argument.

Finally, we assume that $\mathbb{P}_\sigma(q_{\epsilon'}^- \leqslant q_\alpha \leqslant q_{\epsilon'}^+ \mid \mathcal{I} = I) \geqslant 1/3$. We can then bound Equation (18) below by

$$\begin{aligned}&\frac{\int_0^1\int_0^1 P_{q1}\mathbb{G}_0(R_t^q)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1\int_0^1 P_{q1}\mathbb{G}_1(R_t^q)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}_\alpha(t)}\\ &\qquad\geqslant \frac{\int_0^1\int_{q_{\epsilon'}^-}^{q_{\epsilon'}^+} P_{q1}\mathbb{G}_0(R_t^q)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}_\alpha(t)+1}{\int_0^1\int_{q_{\epsilon'}^-}^{q_{\epsilon'}^+} P_{q1}\mathbb{G}_1(R_t^q)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}_\alpha(t)+1}. \qquad (19)\end{aligned}$$

For q in the given range, Lemma 1 part (d) gives

$$\begin{aligned}&\frac{\int_0^1\mathbb{G}_0(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}\\ &= \frac{\int_0^q\mathbb{G}_0(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)+\int_q^1\mathbb{G}_0(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}\\ &\geqslant \frac{(\mathbb{G}_0(1/2)/\mathbb{G}_1(1/2))\int_0^q\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)+\int_q^1\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}\\ &= 1+\left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}-1\right)\frac{\int_0^q\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}\\ &\geqslant 1+\left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}-1\right)\frac{\int_0^{q_{\epsilon'}^-}\mathbb{G}_1(R_t^{q_{\epsilon'}^+})\,\mathrm{d}\mathbb{H}_\alpha(t)}{\int_0^1\mathbb{G}_1(R_t^q)\,\mathrm{d}\mathbb{H}_\alpha(t)}. \qquad (20)\end{aligned}$$

The denominator of the last expression is at most 1, and the numerator is nonnegative. Define $\underline{b} = (2\underline{\beta}+1)/4$. Observe that for any $q \in (0, 1)$, we have $q > R_q^{\underline{b}}$, so we may fix ϵ' small enough so that $q_{\epsilon'}^- > R_{q_{\epsilon'}^+}^{\underline{b}}$. Restricting the range of integration, Equation (19) is bounded below by

$$\begin{aligned}1 &+ \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}-1\right)\int_{R_{q_{\epsilon'}^+}^{\underline{b}}}^{q_{\epsilon'}^-}\mathbb{G}_1(R_t^{q_{\epsilon'}^+})\,\mathrm{d}\mathbb{H}_\alpha(t)\\ &\geqslant 1+\left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}-1\right)\int_{R_{q_{\epsilon'}^+}^{\underline{b}}}^{q_{\epsilon'}^-}\mathbb{G}_1(\underline{b})\,\mathrm{d}\mathbb{H}_\alpha(t)\\ &= 1+\left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}-1\right)\mathbb{G}_1(\underline{b})(\mathbb{H}_\alpha(q_{\epsilon'}^-)-\mathbb{H}_\alpha(R_{q_{\epsilon'}^+}^{\underline{b}})) > 1\end{aligned}$$

for small enough ϵ'. Choosing ϵ' small enough for all three cases to satisfy the corresponding inequalities finishes the proof of the first statement; the second statement is analogous. ☐

Proof of Theorem 4

We first construct a subsequence on which we can apply Lemma 5. Define the indices $i(j)$ recursively, beginning with $i(1) = 1$. For $j > 1$, let

$$\begin{aligned}\epsilon_j = \frac{1}{2}\min_{\mathbf{x}}\min\{&\mathbb{P}_\sigma(\theta = 1 \mid x_{\alpha_{i(l)}},\ l < j),\\ &1 - \mathbb{P}_\sigma(\theta = 1 \mid x_{\alpha_{i(l)}},\ l < j)\}.\end{aligned}$$

Since the cluster is identified, we can choose $i(j)$ large enough so that

$$\mathbb{P}_\sigma(\exists\, l < j \mid \gamma_{\alpha_{i(j)}}^{i(l)}(B(\alpha_{i(j)})) \neq \alpha_{i(l)}) \leqslant \epsilon_j.$$

Since the cluster is uniformly diverse, we can choose the indices so that the resulting subsequence is also uniformly diverse.

For the subsequence $\{\alpha_{i(j)}\}_{j\in\mathbb{N}}$, define $\hat{q}_k = \mathbb{P}_\sigma(\theta = 1 \mid x_{\alpha_{i(j)}},\ j \leqslant k)$. The sequence $\hat{q}_k$ is clearly a bounded martingale, so the sequence converges almost surely to a random variable q^* by the martingale convergence theorem. Conditional on $\theta=1$, the likelihood ratio $(1-\hat{q}_k)/\hat{q}_k$ is also a nonnegative martingale (Doob 1953, Eq. (7.12)). Therefore, conditional on $\theta=1$, the ratio $(1-\hat{q}_k)/\hat{q}_k$ converges with probability 1 to a limiting random variable. In particular,

$$\mathbb{E}_\sigma\left[\frac{1-q^*}{q^*}\,\middle|\,\theta=1\right]<\infty,$$

(Breiman 1968, Theorem 5.14), and therefore, $q^*>0$ with probability 1 when $\theta=1$. Similarly, $\hat{q}_k/(1-\hat{q}_k)$ is a nonnegative martingale conditional on $\theta=0$, and by a similar argument we have $q^*<1$ with probability 1 when $\theta=0$.

We shall see that the random variable q^* equals θ with probability 1. Let $x^{(k)}$ denote the vector comprised of the actions $\{x_{\alpha_{i(j)}}\}_{j\leqslant k}$, and suppose without loss of generality that $x_{\alpha_{i(k+1)}}=0$ for infinitely many k. Using Bayes' rule twice,

$$\begin{aligned}\hat{q}_{k+1}&=\mathbb{P}(\theta=1\mid x_{\alpha_{i(k+1)}}=0,\,x^{(k)})\\&=\left[1+\frac{\mathbb{P}(x^{(k)}\mid\theta=0)\mathbb{P}(x_{\alpha_{i(k+1)}}=0\mid x^{(k)},\,\theta=0)}{\mathbb{P}(x^{(k)}\mid\theta=1)\mathbb{P}(x_{\alpha_{i(k+1)}}=0\mid x^{(k)},\,\theta=1)}\right]^{-1}\\&=\left[1+\left(\frac{1}{\hat{q}_k}-1\right)\frac{\mathbb{P}(x_{\alpha_{i(k+1)}}=0\mid x^{(k)},\,\theta=0)}{\mathbb{P}(x_{\alpha_{i(k+1)}}=0\mid x^{(k)},\,\theta=1)}\right]^{-1}.\qquad(21)\end{aligned}$$

To simplify notation, let

$$f(x^{(k)})=\frac{\mathbb{P}(x_{\alpha_{i(k+1)}}=0\mid x^{(k)},\,\theta=0)}{\mathbb{P}(x_{\alpha_{i(k+1)}}=0\mid x^{(k)},\,\theta=1)}.$$

Thus,

$$\hat{q}_{k+1}=\left[1+\left(\frac{1}{\hat{q}_k}-1\right)f(x^{(k)})\right]^{-1}.$$

Suppose $\hat{q}_k\in[\epsilon,1-\epsilon]$ for all sufficiently large k and some $\epsilon>0$. From our construction of the subsequence $\alpha_{i(j)}$, this implies

$$\begin{aligned}\tilde{q}_k&\equiv\mathbb{P}_\sigma(\theta=1\mid x_m,\,m=\gamma^{i(l)}_{\alpha_{i(k+1)}}(B(\alpha_{i(k+1)})),\,l\leqslant k)\\&\quad\in\left[\frac{\epsilon}{2},1-\frac{\epsilon}{2}\right]\end{aligned}$$

for all such k. By Lemma 6 together with the uniform diversity property, there exists $\delta_\epsilon>0$ such that $f(x^{(k)})\geqslant 1+\delta_\epsilon$. Given this δ_ϵ, we can bound the difference $\hat{q}_k-\hat{q}_{k+1}$. If $x_{\alpha_{i(k+1)}}=0$, we have

$$\begin{aligned}\hat{q}_k-\hat{q}_{k+1}&\geqslant\hat{q}_k-\left[1+\left(\frac{1}{\hat{q}_k}-1\right)(1+\delta_\epsilon)\right]^{-1}\\&=\frac{\hat{q}_k(1-\hat{q}_k)\delta_\epsilon}{1+\delta_\epsilon(1-\hat{q}_k)}\geqslant\frac{\epsilon(1-\epsilon)\delta_\epsilon}{1+(1-\epsilon)\delta_\epsilon}\\&\equiv K_0(\epsilon)>0.\qquad(22)\end{aligned}$$

If $\hat{q}_k\in[\epsilon,1-\epsilon]$ for all sufficiently large k, this implies the sequence $\{\hat{q}_k\}$ is not Cauchy. This contradicts the almost sure convergence of the sequence, so the support of q^* cannot contain $[\epsilon,1-\epsilon]$ for any $\epsilon>0$. Lemma 5 completes the proof. ☐

Proof of Proposition 4

The negative implication follows similar reasoning as Example 4. If one cluster fails to observe the other beyond a certain point, then once social beliefs become strong enough in favor of the cluster's preferred state, the cluster will herd. Since this happens with positive probability, the other cluster obtains bounded information from this cluster, and herds with positive probability as well. We now prove the positive implication.

Define $\hat{q}^C_k=\mathbb{P}_\sigma(\theta=1\mid x_{c_i},\,i\leqslant k)$ and $\hat{q}^D_k=\mathbb{P}_\sigma(\theta=1\mid x_{d_i},\,i\leqslant k)$. Following the same argument as in the proof of Theorem 4, the sequences $\hat{q}^C_k$ and $\hat{q}^D_k$ converge almost surely to random variables q^C and q^D, respectively. We further have for each $i\in\{0,1\}$ that

$$\mathbb{P}_\sigma(q^C=1-i\mid\theta=i)=\mathbb{P}_\sigma(q^D=1-i\mid\theta=i)=0.$$

Without loss of generality, assume $t_{c_1}\in[0,1/2)$; following the argument from Theorem 4, and applying Lemma 6, we immediately find that the support of q^C is contained in $\{0\}\cup[1-\underline{\beta},1]$ and the support of q^D is contained in $[0,1-\bar{\beta}]\cup\{1\}$.

Define the constants δ^C and δ^D by

$$\begin{aligned}&\delta^C=\inf\{\delta\mid\mathbb{P}_\sigma(q^C\in\{0\}\cup[1-\delta,1])=1\}\quad\text{and}\\&\delta^D=\inf\{\delta\mid\mathbb{P}_\sigma(q^D\in[0,\delta]\cup\{1\})=1\}.\end{aligned}$$

Our first task is to show that if $q^C\in[1-\delta^C,1)$, then only finitely many agents in C select action 0; similarly if $q^D\in(0,\delta^D]$, only finitely many agents in D selection action 1. We will analyze only the cluster C since the other result is analogous. The argument should be familiar by now; we establish a contradiction with almost sure convergence of the martingale.

Suppose $x_{c_{k+1}}=0$ with positive probability, and recall the computation leading to Equation (18) in the proof of Lemma 6. Consider the corresponding result with $\mathcal{I}=\{x_{c_i}\}_{i\leqslant k}$,

$$\begin{aligned}&\frac{\mathbb{P}_\sigma(x_{c_{k+1}}=0\mid x_{c_i},\,i\leqslant k,\,\theta=0)}{\mathbb{P}_\sigma(x_{c_{k+1}}=0\mid x_{c_i},\,i\leqslant k,\,\theta=1)}\\&=\frac{\int_0^1\int_0^1 P_{q1}\mathbb{G}_0(R^q_t)+(P_{q0}-P_{q1})(\mathbb{G}_0(R^q_t)-\mathbb{G}_0(R^{\hat{q}_\alpha}_t))\,\mathrm{d}q\,\mathrm{d}\mathbb{H}\_(t)}{\int_0^1\int_0^1 P_{q1}\mathbb{G}_1(R^q_t)\,\mathrm{d}q\,\mathrm{d}\mathbb{H}\_(t)}.\end{aligned}$$

If $q^C\in[1-\delta^C,1-\epsilon]$ for some fixed $\epsilon>0$, then by almost sure convergence, for all k large enough we must have $\hat{q}^C_k\geqslant 1-\delta^C-\epsilon\geqslant 1-\underline{\beta}-\epsilon$. Since $t_{c_{k+1}}\leqslant\frac{1}{2}$ with probability 1, we have $\mathbb{G}_0(R^q_t)=\mathbb{G}_1(R^q_t)=0$ for any $q>1-\underline{\beta}$, so these values of q do not contribute to the integrals; from here on we consider the densities P_{qi} above to be conditional on this event. We analyze two cases: First suppose that conditional on $q\leqslant 1-\underline{\beta}$, we have $q\leqslant 1-\underline{\beta}-2\epsilon$ with probability at least $1/2$. Following the argument of Lemma 6, this gives us a lower bound on the second term in the numerator, and therefore bounds the ratio strictly above one.

Alternatively, suppose that conditional on $q \leqslant 1-\underline{\beta}$ we have $q \in [1-\underline{\beta}-2\epsilon, 1-\underline{\beta}]$ with probability at least 1/2. For any t sufficiently close to 1/2, we have

$$\frac{\mathbb{G}_0(R_t^q)}{\mathbb{G}_1(R_t^q)} \geqslant \frac{\mathbb{G}_0(\underline{\beta}+2\epsilon)}{\mathbb{G}_1(\underline{\beta}+2\epsilon)} > 1.$$

Since $\mathbb{H}\_$ is supported in any interval $[1/2-\delta, 1/2]$, positive measure is assigned to these values of t, and together with Lemma 1 part (d) this again allows us to bound the ratio strictly above one. Following the argument of Theorem 4, infinitely many agents in C selecting action 0 would contradict almost sure convergence of $\hat{q}_k^C$ if $q^C \in [1-\delta^C, 1-\epsilon]$. Since ϵ was arbitrary, the result follows.

Our next task is to show for a fixed $q \in [1-\delta^C, 1)$ we have

$$\mathbb{P}_\sigma(q^D \in (0, \delta^D] \mid q^C = q,\ \theta = 1) \geqslant \frac{(1-\underline{\beta})(1-q)}{\underline{\beta} q}. \tag{23}$$

Similarly, for a fixed $q \in (0, \delta^D]$ we have

$$\mathbb{P}_\sigma(q^C \in [\delta^C, 1) \mid q^C = q,\ \theta = 0) \geqslant \frac{\bar{\beta} q}{(1-\bar{\beta})(1-q)}. \tag{24}$$

Let N^C and N^D denote the smallest integer valued random variables such that $x_{c_k} = 1$ for all $k \geqslant N^C$ and $x_{d_k} = 0$ for all $k \geqslant N^D$. From the above work, we have that N^C is finite with probability 1 conditional on $q^C \in [1-\delta^C, 1)$, and N^D is finite with probability 1 if $q^D \in (0, \delta^D]$. Consider a particular sequence of decisions $\{x_{c_k}\}_{k\in\mathbb{N}}$ such that $q^C = q \in [1-\delta^C, 1)$. This immediately implies

$$\begin{aligned}\lim_{k\to\infty} &\frac{\mathbb{P}_\sigma(x_{d_k} = 0 \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 1)}{\mathbb{P}_\sigma(x_{d_k} = 0 \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 0)} \\ &= \frac{\mathbb{P}_\sigma(q^D \in (0, \delta^D] \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 1)}{\mathbb{P}_\sigma(q^D \in [0, \delta^D] \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 0)} \\ &= \mathbb{P}_\sigma(q^D \in (0, \delta^D] \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 1).\end{aligned}$$

Given an $\epsilon > 0$ we can choose N_ϵ large enough so that $\mathbb{P}_\sigma(N_C \leqslant N_\epsilon \mid \{x_{c_k}\}_{k\leqslant N_\epsilon}) \geqslant 1-\epsilon$. If further we have N_ϵ larger than the realized N^C for our fixed sequence, we have the bounds

$$\begin{aligned}&\frac{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 1)}{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 0)} \\ &\qquad \leqslant \frac{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\leqslant n},\ \theta = 1)+\epsilon}{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\leqslant n},\ \theta = 0)-\epsilon}, \quad \text{and} \\ &\frac{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 1)}{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\in\mathbb{N}},\ \theta = 0)} \\ &\qquad \geqslant \frac{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\leqslant n},\ \theta = 1)-\epsilon}{\mathbb{P}_\sigma(x_{d_j} = 0 \mid \{x_{c_k}\}_{k\leqslant n},\ \theta = 0)+\epsilon}\end{aligned}$$

for any $n \geqslant N_\epsilon$ and all j.

Suppose $\mathbb{P}_\sigma(q^D \in (0, \delta^D] \mid q^C = q,\ \theta = 1) < (1-\underline{\beta})(1-q)/(\underline{\beta} q)$. From the above work, we can find a collection of sequences $\{x_{c_k}\}_{k\in\mathbb{N}}$ with positive measure conditional on $q^C = q$, an $\epsilon > 0$, and an integer N'_ϵ such that

$$\frac{\mathbb{P}_\sigma(x_{d_m} = 0 \mid \{x_{c_k}\}_{k\leqslant n},\ \theta = 1)}{\mathbb{P}_\sigma(x_{d_m} = 0 \mid \{x_{c_k}\}_{k\leqslant n},\ \theta = 0)} \leqslant \frac{(1-\underline{\beta})(1-q)}{\underline{\beta} q} - \epsilon$$

for all $n, m \geqslant N'_\epsilon$ when one of the sequences $\{x_{c_k}\}_{k\in\mathbb{N}}$ is realized. For all sufficiently large n, $\hat{q}_n^C$ is within $\epsilon/2$ of q; this in turn implies that for all such n and $m \geqslant N'_\epsilon$, we have an $\epsilon' > 0$ such that

$$\mathbb{P}_\sigma(\theta = 1 \mid x_{d_m} = 0,\ x_{c_k},\ k \leqslant n) \leqslant 1-\underline{\beta}-\epsilon'.$$

Note that for any subsequence of the cluster D, since types are conditionally independent and arbitrarily close to 1, infinitely many members of the subsequence will select action 0. If $d_m \in B(c_n)$ for n, m as above, we have $\mathbb{E}_\sigma[q_{c_n} \mid x_{d_m} = 0,\ x_{c_k},\ k \leqslant n] \leqslant 1-\underline{\beta}-\epsilon'$. Since social beliefs are bounded between 0 and 1, there is a positive lower bound on the probability that $q_{c_n} \leqslant 1-\underline{\beta}-\epsilon'/2$ conditional on these observations. Since infinitely many agents c_n observe such d_m and have types arbitrarily close to 1/2, with probability 1 infinitely many agents in C choose action 0. We conclude that either $\delta^C = 0$ or

$$\begin{aligned}\lim_{n\to\infty} &\mathbb{P}_\sigma(q^D \in (0, \delta^D] \mid x_{c_k},\ k \leqslant n,\ \theta = 1) \\ &= \mathbb{P}_\sigma(q^D \in (0, \delta^D] \mid q^C = q,\ \theta = 1) \geqslant \frac{(1-\underline{\beta})(1-q)}{\underline{\beta} q}\end{aligned}$$

as claimed. In the former case, decisions of agents in C become fully informative, and asymptotic learning clearly obtains. On the other hand, if Equations (23) and (24) hold,

$$\begin{aligned}&\mathbb{P}_\sigma(q^C \in [1-\delta^C, 1] \mid \theta = 0) \\ &\quad = \int_0^{\delta^D} \mathbb{P}_\sigma(q^C \in [1-\delta^C, 1] \mid q^D = q, \theta = 0) \\ &\qquad \cdot \mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 0) \\ &\quad \geqslant \int_0^{\delta^D} \frac{\bar{\beta} q}{(1-\bar{\beta})(1-q)} \mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 0) \\ &\quad = \int_0^{\delta^D} \frac{\bar{\beta} q}{(1-\bar{\beta})(1-q)} \\ &\qquad \cdot \frac{\mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 0)}{\mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 1)} \mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 1).\end{aligned} \tag{25}$$

Now, observe the definition of q^D implies

$$\frac{\mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 0)}{\mathrm{d}\mathbb{P}_\sigma(q^D = q \mid \theta = 1)} = \frac{\mathbb{P}_\sigma(\theta = 0 \mid q^D = q)}{\mathbb{P}_\sigma(\theta = 1 \mid q^D = q)} = \frac{1-q}{q}.$$

Substituting into Equation (25) gives

$$\begin{aligned}\mathbb{P}_\sigma(q^C \in [1-\delta, 1] \mid \theta=0) &\geqslant \frac{\bar{\beta}}{1-\bar{\beta}} \int_0^{\delta^D} d\mathbb{P}_\sigma(q^D = q \mid \theta=1)\\ &= \frac{\bar{\beta}}{1-\bar{\beta}} \mathbb{P}_\sigma(q^D \in [0, \delta^D] \mid \theta=1).\end{aligned}$$

A similar calculation for q^D gives

$$\mathbb{P}_\sigma(q^D \in [0, \delta^D] \mid \theta=1) \geqslant \frac{1-\underline{\beta}}{\underline{\beta}} \mathbb{P}_\sigma(q^C \in [1-\delta^C, 1] \mid \theta=0).$$

Combining the two gives

$$\begin{aligned}&\mathbb{P}_\sigma(q^C \in [1-\delta^C, 1] \mid \theta=0)\\ &\quad \geqslant \frac{\bar{\beta}(1-\underline{\beta})}{(1-\bar{\beta})\underline{\beta}} \mathbb{P}_\sigma(q^C \in [1-\delta^C, 1] \mid \theta=0).\end{aligned}$$

Since the ratio on the right-hand side is strictly larger than one, this leads to a contradiction if the probability is larger than zero. We conclude that both probabilities are equal to zero, and consequently, asymptotic learning obtains. ☐

Endnotes

1. See Marsden (1988), McPherson et al. (2001), and Currarini et al. (2009, 2010).
2. This includes the foundational papers by Banerjee (1992) and Bikhchandani et al. (1992) as well as more recent work by Bala and Goyal (1998), Gale and Kariv (2003), Çelen and Kariv (2004), Guarino and Ianni (2010), Acemoglu et al. (2011, 2014), Mossel et al. (2012), Mueller-Frank (2013), and Lobel and Sadler (2015).
3. See Smith and Sorensen (2000), Acemoglu et al. (2011), and Lobel and Sadler (2015).
4. The range $[1, 3/2]$ of the function $\mathcal{Z}$ in this definition corresponds to the possible range of an agent's ex ante expected utility, not the possible range of an agent's probability of matching the state of the world.

References

Acemoglu D, Bimpikis K, Ozdaglar A (2014) Dynamics of information exchange in endogenous social networks. *Theoret. Econom.* 9(1):41–97.

Acemoglu D, Dahleh M, Lobel I, Ozdaglar A (2011) Bayesian learning in social networks. *Rev. Econom. Stud.* 78(4):1201–1236.

Aral S, Van Alstyne M (2011) The diversity-bandwidth tradeoff. *Amer. J. Sociol.* 117(1):90–171.

Aral S, Walker D (2014) Tie strength, embeddedness, and social influence: Evidence from a large scale networked experiment. *Management Sci.* 60(6):1352–1370.

Bakshy E, Rosenn I, Marlow C, Adamic L (2012) The role of social networks in information diffusion. *Proc. 21st Internat. Conf. World Wide Web* (ACM, New York), 519–528.

Bala V, Goyal S (1998) Learning from neighbours. *Rev. Econom. Stud.* 65(3):595–621.

Banerjee A (1992) A simple model of herd behavior. *Quart. J. Econom.* 107:797–817.

Banerjee A, Fudenberg D (2004) Word-of-mouth learning. *Games Econom. Behav.* 46:1–22.

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as information cascades. *J. Political Econom.* 100:992–1026.

Breiman L (1968) *Probability* (Addison-Wellesley, Reading, MA).

Burt R (2004) Structural holes and good ideas. *Amer. J. Sociol.* 110(2):349–399.

Çelen B, Kariv S (2004) Observational learning under imperfect information. *Games Econom. Behav.* 47(1):72–86.

Centola D (2011) An experimental study of homophily in the adoption of health behavior. *Science* 334(6060):1269–1272.

Centola D, Macy M (2007) Complex contagions and the weakness of long ties. *Amer. J. Sociol.* 113(3):702–734.

Cipriani M, Guarino A (2008) Herd behavior and contagion in financial markets. *The B.E. J. Theoret. Econom.* 8(1):1–54.

Conley T, Udry C (2010) Learning about a new technology: Pineapple in Ghana. *Amer. Econom. Rev.* 100:35–69.

Currarini S, Jackson M, Pin P (2009) An economic model of friendship: Homophily, minorities, and segregation. *Econometrica* 77(4):1003–1045.

Currarini S, Jackson M, Pin P (2010) Identifying the roles of race-based choice and chance in high school friendship network formation. *Proc. Natl. Acad. Sci.* 107(11):4857–4861.

Doob J (1953) *Stochastic Processes* (John Wiley & Sons, New York).

Gale D, Kariv S (2003) Bayesian learning in social networks. *Games Econom. Behav.* 45(2):329–346.

Goeree J, Palfrey T, Rogers B (2006) Social learning with private and common values. *Econom. Theory* 28(2):245–264.

Golub B, Jackson M (2012) How homophily affects the speed of learning and best-response dynamics. *Quart. J. Econom.* 127(3):1287–1338.

Granovetter M (1973) The strength of weak ties. *Amer. J. Sociol.* 78(6):1360–1380.

Guarino A, Ianni A (2010) Bayesian social learning with local interactions. *Games* 1(4):438–458.

Hansen M (1999) The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Admin. Sci. Quart.* 44(1):82–111.

Lobel I, Sadler E (2015) Information diffusion in networks through social learning. *Theoret. Econom.* Forthcoming.

Lobel I, Acemoglu D, Dahleh M, Ozdaglar A (2009) Lower bounds on the rate of learning in social networks. *IEEE Proc. Amer. Control Conf., AAC '09* (IEEE, Piscataway, NJ), 2825–2830.

Marsden P (1988) Homogeneity in confiding relations. *Soc. Networks* 10(1):57–76.

McPherson M, Smith-Lovin L, Cook J (2001) Birds of a feather: Homophily in social networks. *Annual Rev. Sociol.* 27:415–444.

Mossel E, Sly A, Tamuz O (2012) On agreement and learning. Working paper, University of California, Berkeley, Berkeley, CA.

Mueller-Frank M (2013) A general framework for rational learning in social networks. *Theoret. Econom.* 8(1):1–40.

Smith L, Sorensen P (2000) Pathological outcomes of observational learning. *Econometrica* 68:371–398.

Sorensen A (2006) Social learning and health plan choice. *RAND J. Econom.* 37:929–945.

Vives X (1993) How fast do rational agents learn? *Rev. Econom. Stud.* 60(2):329–347.

Vives X (1995) The speed of information revelation in a financial market mechanism. *J. Econom. Theory* 67(1):178–204.

Ilan Lobel is an assistant professor of information, operations, and management sciences at the Stern School of Business at New York University. He is interested in questions of pricing, learning, and contract design in dynamic and networked markets.

Evan Sadler is a doctoral student in information, operations, and management sciences at the Stern School of Business at New York University. He is interested in the ways people influence each other's behavior and how the structure of a social network affects behavior in aggregate.