Network structure and naive sequential learning

Krishna Dasaratha
Department of Economics, Harvard University

Kevin He
Division of Humanities and Social Sciences, California Institute of Technology and Department of
Economics, University of Pennsylvania

We study a sequential-learning model featuring a network of naive agents with Gaussian information structures. Agents apply a heuristic rule to aggregate predecessors' actions. They weigh these actions according the strengths of their social connections to different predecessors. We show this rule arises endogenously when agents wrongly believe others act solely on private information and thus neglect redundancies among observations. We provide a simple linear formula expressing agents' actions in terms of network paths and use this formula to characterize the set of networks where naive agents eventually learn correctly. This characterization implies that, on all networks where later agents observe more than one neighbor, there exist disproportionately influential early agents who can cause herding on incorrect actions. Going beyond existing social-learning results, we compute the probability of such mislearning exactly. This allows us to compare likelihoods of incorrect herding, and hence expected welfare losses, across network structures. The probability of mislearning increases when link densities are higher and when networks are more integrated. In partially segregated networks, divergent early signals can lead to persistent disagreement between groups.

Keywords. Network structure, sequential social learning, naive inference, mislearning, disagreeement.

JEL classification. D85, D83, D90.

1. Introduction

Consider an environment with a sequence of agents facing the same decision problem in turn, where each agent considers both her private information and the behavior of those who came before her in reaching a decision. When consumers choose between rival products, for instance, their decisions are often informed by the choices of early customers. When doctors decide on a treatment for their patients, they consult best practices established by other clinicians who came before them. Additionally, when a

new theory or rumor is introduced into a society, individuals are swayed by the discussions and opinions of those who have already taken a clear stance on the new idea.

A key feature of these examples is that agents observe only the behavior of a certain subset of their predecessors. For example, a consumer may know about her friends' recent purchases, but not the product choices of anyone outside her social circle. In general, each sequential social-learning problem has an *observation network* that determines which predecessors are observable to each agent. Observation networks associated with different learning problems may vary in density, extent of segregation, and other structural properties. Hence, our central research question: how does the structure of the observation network affect the probability of correct social learning in the long run?

To answer this question, our model must capture key behavioral patterns in how individuals process social information in these learning settings. Empirical research on social learning suggests that humans often exhibit *inferential naiveté*, failing to understand that their predecessors' actions reflect a combination of private information and the inference those predecessors have drawn from the behavior of still others (e.g., Chandrasekhar et al. 2020, Eyster et al. 2015). Returning to the examples, a consumer may mistake a herd on a product for evidence that everyone has positive private information about the product's quality. In an online community, a few early opinion-makers can make a rumor go viral, due to people not thinking through how the vast majority of a viral story's proponents are just following the herd and possess no private information about the rumor's veracity.

The present study examines the effect of the *observation network* on the extent of learning, in a setting where players suffer from *inferential naiveté*. We analyze the theoretical implications of a tractable log-linear learning rule that aggregates observations in a manner related to the DeGroot heuristic. Agents who fail to account for correlations in their observations (as in Eyster and Rabin 2010) choose actions according to this log-linear rule. We also introduce *weighted networks* and a weighted version of log-linear behavior, where agents place different weights on different neighbors' actions. We show that such decision weights arise when agents additionally misperceive the precisions of predecessors' signals according to the strengths of their links to said predecessors.

When combined with a Gaussian informational environment, our weighted log-linear learning rule lets us compute naive agents' exact probability of taking the correct action ("learning accuracy") on arbitrary weighted networks. In contrast, the existing literature on social learning has focused on whether long-run learning outcome exhibits certain properties (e.g., convergence to dogmatic beliefs in the wrong state) with positive probability, but not on how these probabilities vary across environments. In settings where learning is imperfect (i.e., society does not almost surely learn the correct state in the long run), we obtain a richer characterization of the learning outcome and compute comparative statics of learning accuracy with respect to network structures.

Under our weighted log-linear learning rule, actions take a simple form: given a continuous action space and a binary state space, we can express each agent's action as a log-linear function of her predecessors' private signal realizations with coefficients depending on the weighted network structure. We exploit this expression to develop

a necessary and sufficient condition for society to learn completely: no agent has too much "influence." Imperfect learning is the leading case: some agent in the network must have disproportionate influence whenever all but finitely many people observe more than one neighbor. Since this condition applies to a very broad class of networks, our analysis focuses on comparing differentially inefficient learning outcomes across networks. The detailed comparative statics we obtain from this approach are crucial: imperfect learning implies a wide range of welfare losses on different networks.

Introducing naiveté generates predictions matching empirical observations in several domains where the rational model's implications are unclear or counterfactual. We prove that increasing network density leads to more inaccurate social-learning outcomes in the long run. This prediction is supported by the experimental results in our companion paper (Dasaratha and He 2019), where we find human subjects' accuracy gain from social learning is twice as large on sparse networks compared to dense networks. As another example, disagreement among different communities is common in practice. In the domain of product adoption, respective subcommunities frequently insist upon the superiority of their preferred products. We prove that if agents' actions only coarsely reflect their beliefs and society is partially segregated, then two social subgroups can disagree forever.[1] Because of the limited information conveyed by actions, disagreement can persist even when agents observe the actions of unboundedly many individuals from another group. This presents a sharp contrast with Bayesian social-learning models (and other leading learning models such as DeGroot learning), where asymptotic agreement is a robust prediction.

1.1 *Related literature*

1.1.1 *Effect of network structure on learning*   Much of the literature on how network structure matters for learning outcomes has focused on networked agents repeatedly guessing a state while learning from the same set of neighbors each period (e.g., Bala and Goyal 1998).

The leading behavioral model here is the DeGroot heuristic, which forms a belief in each period by averaging the beliefs of neighbors in the previous period. A key prediction of DeGroot learning is that society converges to a long-run consensus (DeMarzo et al. 2003), which will be correct in large networks as long as the network is not too unbalanced (Golub and Jackson 2010). Much of the analysis focuses on how network structure (e.g., homophily) matters for the speed of convergence to correct consensus beliefs (Golub and Jackson 2012). We find that in a sequential setting, natural changes in network structure matter for asymptotic accuracy, not only for speed of learning. Changing network density, which has no effect on DeGroot learning in large networks, can substantially alter the probability that a society learns correctly. One intuition for this difference is that DeGroot agents assign weights adding up to 1 to their neighbors, but agents in our setting have increasing out-degrees with increasing network density and therefore can overweight their social information. Homophily also matters for this probability and even for whether consensus is ever reached.

---

[1]Segregation is extensively documented in social networks (see McPherson et al. 2001 for a survey).

While DeGroot proposes the averaging rule as an ad hoc heuristic, several recent papers have developed behavioral microfoundations for learning in the repeated-interaction setting (Molavi et al. 2018, Mueller-Frank and Neri 2019, Levy and Razin 2018). These models closely resemble ours at the level of individual behavior, but their predictions about society's long-run beliefs are more in line with DeGroot. As such, changes in network structure again have a limited scope for affecting learning outcomes in this literature.

1.1.2 *Sequential social learning*   We consider the same environment as the extensive literature on sequential social learning beginning with Banerjee (1992) and Bikhchandani et al. (1992). Acemoglu et al. (2011) and Lobel and Sadler (2015) characterize network features that lead to correct asymptotic learning for Bayesians who move sequentially. By providing a thorough understanding of rational learning in sequential settings, this literature provides a valuable benchmark as we study naive learning. We find that among network structures where Bayesian agents learn asymptotically, there is large variation in the probability of mislearning for naive agents.

Several authors look at sequential behavioral learning on a particular network structure, usually the complete network (Eyster and Rabin 2010, Bohren 2016, Bohren and Hauser 2018). We characterize several ways in which the choice of network structure matters for the distribution of long-run outcomes. Eyster and Rabin (2014) exhibit a general class of social-learning rules, which includes the weighted log-linear rule we study for certain values of the weights, where mislearning occurs with positive probability. We go beyond this general result by deriving expressions for the exact probabilities of mislearning on different networks, whose associated welfare losses cannot be compared using the binary classification of Eyster and Rabin (2014).

2. Model

2.1 *Sequential social learning on a weighted network*

There are two possible states of the world, $\omega \in \{0, 1\}$, both equally likely. There is an infinite sequence of agents indexed by $i \in \mathbb{N}$. Agents move in order, each acting once.

On her turn, agent i observes a private signal $s_i \in \mathbb{R}$. Private signals (s_i) are Gaussian and independent and identically distributed (i.i.d.) given the state. When $\omega = 1$, we have $s_i \sim \mathcal{N}(1, \sigma^2)$ for some conditional variance $\sigma^2 > 0$. When $\omega = 0$, we have $s_i \sim \mathcal{N}(-1, \sigma^2)$.

In addition to her private signal, agent i also observes the actions of previous agents. Then i chooses an action $a_i \in [0, 1]$. In our microfoundation for Definition 1, agent i chooses a_i to maximize the expectation of

$$u_i(a_i, \omega) := -(a_i - \omega)^2$$

given her belief about ω, which we describe later. So her chosen action corresponds to the probability she assigns to the event $\{\omega = 1\}$.

We find it convenient to work with the following change of variables.

Notation 1. We have $\tilde{s}_i := \ln(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]})$ and $\tilde{a}_i := \ln(\frac{a_i}{1-a_i})$.

In words, $\tilde{s}_i$ is the log-likelihood ratio of the events $\{\omega = 1\}$ and $\{\omega = 0\}$ given signal s_i, while it is easy to show that $\tilde{a}_i$ is the log-likelihood ratio of $\{\omega = 1\}$ and $\{\omega = 0\}$ corresponding to action a_i. That is to say, if a_i is optimal given i's beliefs, then $\tilde{a}_i$ is the log-likelihood ratio of $\{\omega = 1\}$ and $\{\omega = 0\}$ according to i's beliefs. Note that the transformations from s_i to $\tilde{s}_i$ and from a_i to $\tilde{a}_i$ are bijective, so no information is lost when we relabel variables.

Agents are linked to all of their predecessors on a weighted network, with a lower-triangular adjacency matrix M where all diagonal entries are equal to 0. For $i > j$, the weight of the link from i to j is given by $M_{i,j} \in [0, 1]$. The weights of the edges determine the relative importance agents place on others' actions in forming their beliefs. In Section 4, we derive comparative statics with respect to the network structure. Because studying continuous changes in the network is more tractable than discrete changes, we consider a model that allows interior network weights.

Throughout, we study *naive agents* who choose actions equal to a weighted sum of their observations according to the following log-linear updating rule.

Definition 1. Agents use the *weighted log-linear rule* if each agent i plays

$$\tilde{a}_i = \tilde{s}_i + \sum_{j<i} M_{i,j} \tilde{a}_j.$$

In words, each agent i's log action is a weighted sum of her predecessors' log actions and her own log signal. The network M exogenously determines the relative influences of different predecessors' behavior on i's play, with j's influence proportional to the strength of i's social connection to her. By contrast, a society of rational agents would put endogenous weights on others' actions that are not simply proportional to strengths of the network links between them. For example, if agents only observe the actions of linked neighbors, rational agents would play the unique perfect Bayesian equilibrium of the social-learning game, in which case some equilibrium decision weights may be negative.[2]

Remark 1. The formula in Definition 1 resembles the DeGroot updating rule. A key distinction is that we allow for agents to have any out-degree, while the DeGroot heuristic requires all agents' weights sum to 1. In an unweighted network, any agent with multiple observations has an out-degree greater than 1. This distinction is not just a normalization, but is, in fact, the source of redundancy under naive inference.

2.2 *Microfoundation for weighted log-linear rule*

In this subsection, we provide a psychological microfoundation for the weighted log-linear rule from Definition 1. We first show that on *unweighted networks* (i.e., when

---

[2]We omit the proof that actions are log linear at the perfect Bayesian equilibrium. This can be shown by induction, and the key step is a calculation similar to Lemma 3.

each $M_{i,j}$ is either 0 or 1), this rule follows from a primitive assumption about agents' inference.

A growing body of recent evidence in psychology and economics shows that people learning from peers are often not fully correct in their treatment of social structure (Chandrasekhar et al. 2020, Enke and Zimmermann 2017). Instead of calculating the optimal Bayesian behavior that fully takes into account all they know about the network and signal structure, agents often apply heuristic simplifications to their environment. When networks are complicated and/or uncertain, determining Bayesian behavior can be intractable (Hązła et al. forthcoming) and these heuristic learning rules become especially prevalent (Eyster et al. 2015). Motivated by this observation, we consider the following behavioral assumption.

ASSUMPTION 1. *Each agent wrongly believes that each predecessor chooses an action to maximize her expected payoff based only on her private signal and not on her observation of other agents.*

This inferential mistake can be equivalently described as agent i misperceiving $M_{j,k} = 0$ for all $k < j < i$. Under this interpretation, i acts as if her neighbors do not take into account their own predecessors' actions.

In the sequential-learning literature, Assumption 1 was first studied on the complete network by Eyster and Rabin (2010), who coined the term "best-response trailing naive inference" (BRTNI) to describe this behavior. The laboratory games in Eyster et al. (2015) and Mueller-Frank and Neri (2015) find evidence for this behavioral assumption.

Agents who make this inferential mistake use the log-linear rule on unweighted networks, providing a psychological microfoundation for the behavior we study.

LEMMA 1. *On an unweighted network where agents observe only the actions of linked predecessors, Assumption 1 implies agents use the weighted log-linear rule.*

Due to the inferential mistake, agent i wrongly infers that j's log action equals her log signal. (This inference is possible since the continuum action set is rich enough to exactly reveal beliefs of predecessors.) The action a_i is the product of the relevant likelihoods because an agent satisfying Assumption 1 thinks her observations are based on independent information, and therefore $\tilde{a}_i$ is the sum of the corresponding log-likelihood ratios.

REMARK 2. Inference under Assumption 1 is cognitively simple in that it does not rely on agents' knowledge about the network (beyond their own neighborhoods) or even knowledge about the order in which their predecessors moved. Our model therefore applies even to complex environments with random arrival of agents. In such environments, Assumption 1 may be more realistic than assuming full knowledge about the observation structure and move order.

Next we give a microfoundation for the same behavior on weighted networks. We provide an interpretation of network weights that formalizes the idea that agents place

more trust in neighbors with whom their connections are stronger: we suppose that agents underestimate the precision of others' private signals in a way that depends on $M_{i,j}$.

ASSUMPTION 2. *Given network weight $M_{i,j} \in [0, 1]$, agent i believes j's private signal has conditional variance $\sigma^2/M_{i,j}$ given the state.*

Weizsäcker's (2010) meta-analysis of sequential social-learning experiments finds that laboratory subjects underuse social information relative to their own private signals. Our Assumption 2 is consistent with this evidence, but also allows for different degrees of underuse for different predecessors. Weaker network connections formally correspond to predecessors whose signals are believed to be less informative about the state or less relevant. Conversely, if we know that i acts as if j's signal has conditional variance $V_{i,j} \geq \sigma^2$, then we can construct a weighted network with weights $M_{i,j} = \sigma^2/V_{i,j}$.

The next result shows the combination of the inferential mistake about others' social information and the underestimation of others' signal precisions (Assumptions 1 and 2) provides a microfoundation for the weighted log-linear rule.

LEMMA 2. *Agents who satisfy Assumptions 1 and 2 use the weighted log-linear rule.*

By a property of the Gaussian distribution, a log-transformed Gaussian variable is also Gaussian, which is the key to showing the above lemma.

2.3 *Complete learning and mislearning*

We define what it means for society to learn completely in terms of convergence of actions.

DEFINITION 2. Society *learns completely* if (a_n) converges to ω in probability.

Since a_n reflects agent n's belief in $\{\omega = 1\}$ in our microfoundation of weighted log-linear inference, this definition also describes a property about the convergence of beliefs.[3] In a setting where society learns completely, agent n becomes very likely to believe strongly in the true state of the world as n grows large.

One failure of complete learning is when society becomes fully convinced of the wrong state of the world with positive probability, an event we call mislearning.

DEFINITION 3. Society *mislearns* when $\lim_{n \to \infty} a_n = 0$ but $\omega = 1$ or when $\lim_{n \to \infty} a_n = 1$ but $\omega = 0$.

REMARK 3. Mislearning is not the only obstacle to complete learning. Consider a network where, for $i \geq 3$, we have $M_{i,1} = M_{i,2} = 1$ and $M_{i,j} = 0$ for all $j \neq 1, 2$. Clearly this society neither learns completely nor mislearns with positive probability. Instead, agents' beliefs almost surely do not converge.

---

[3]In general, we treat a_n as the belief of a naive agent who plays $\tilde{a}_n$.

3. Social influence and learning

In this section, we develop a necessary and sufficient condition on the network for society to mislearn. We argue that this condition is satisfied by a large class of networks of economic relevance.

3.1 *Path-counting interpretation of actions*

We now show that with naive agents, actions have a simple (log-) linear expression in terms of paths in the network. Unlike the expression in Definition 1, this next result expresses actions in terms of only signal realizations and the network structure, making no reference to predecessors' actions.

Let $M[n]$ refer to the $n \times n$ upper-left submatrix of M.

PROPOSITION 1. *Consider any weighted network M. For each n, the actions of the first n agents are determined by*

$$\begin{pmatrix} \tilde{a}_1 \\ \vdots \\ \tilde{a}_n \end{pmatrix} = \left(I - M[n]\right)^{-1} \cdot \begin{pmatrix} \tilde{s}_1 \\ \vdots \\ \tilde{s}_n \end{pmatrix}.$$

So $\tilde{a}_i$ is a linear combination of $(\tilde{s}_j)_{j=1}^{i}$, with coefficients given by the number of weighted paths from i to j in the network with adjacency matrix M.

From a combinatorial perspective, the formula says that the influence of j's signal on i's action depends on the number of weighted paths from i to j. In unweighted networks where all entries in M are 1 or 0, this is just the number of paths. In general, "weighted paths" means the path passing through agents $i_0, \ldots, i_K$ is counted with weight $\prod_{k=0}^{K-1} M_{i_k, i_{k+1}}$.

Our Proposition 1 resembles a formula for agents' actions in Levy and Razin (2018). In a setting of repeated interaction with a fixed set of neighbors instead of sequential social learning, Levy and Razin (2018) also find that the influence of i's private information on j's period t posterior belief depends on the number of length-t paths from i to j in the network.

3.2 *Condition for complete learning*

We now use the representation result of Proposition 1 to study which networks lead to mislearning by naive agents.

We define below a notion of network influence for the sequential social-learning environment, which plays a central role in determining whether society learns completely in the long run.

DEFINITION 4. Let $b_{i,j} := (I - M[i])_{i,j}^{-1}$ be the number of weighted paths from i to j in network M.

Because of Proposition 1, these path counts are important to our analysis.

DEFINITION 5. For $n > i$, the *influence* of i on n is $\mathbb{I}(i \rightarrow n) := b_{n,i}/\sum_{j=1}^{n} b_{n,j}$.

That is to say, the influence of i on n is the fraction of paths from n that end at i.

A different definition of influence appears in Golub and Jackson (2010), who study DeGroot learning in a network where agents act simultaneously each period. For them, the influence of an agent i is determined by the unit left eigenvector of the belief-updating matrix, which is proportional to i's degree in an undirected network with symmetric weights. Both definitions are related to the proportion of walks terminating at an agent, but because of the asymmetry between earlier and later agents in the sequential setting, the distribution of walks tends to be more unbalanced.

For Proposition 2 only, we consider networks that satisfy the following connectedness condition.

DEFINITION 6. Network M satisfies the *connectedness condition* if there is an integer N and constant $C > 0$ such that for all $i > N$, there exists $j < N$ with $b_{i,j} \geq C$.

Intuitively, this says that all sufficiently late agents are indirectly influenced by some early agent. An unweighted network satisfies the connectedness condition if and only if there are only finitely many agents who have no neighbors. If such a network violates the connectedness condition, then clearly the infinitely many agents without neighbors will prevent society from learning completely. All weighted networks studied in Section 4 also satisfy the connectedness condition.

PROPOSITION 2. *Consider any weighted network satisfying the connectedness condition. Society learns completely if and only if* $\lim_{n \rightarrow \infty} \mathbb{I}(i \rightarrow n) = 0$ *for all i.*

Proposition 2 says that beliefs always converge to the truth if and only if no agent has undue influence in the network. This is a recurring insight in research on social learning on networks, beginning with the "royal family" example and related results in Bala and Goyal (1998). Other examples where excessive influence hinders social learning in networks include Golub and Jackson (2010), Acemoglu et al. (2010), and Mossel et al. (2015). The main contribution of Proposition 2 is to identify the relevant measure of influence in our sequential-learning setting with naive agents. In our setting, unlike on large unordered networks as in Golub and Jackson (2010), the ordering of agents creates an asymmetry that prevents society from learning completely on most natural networks. Early movers influence many successors, an asymmetry unique to the sequential-learning setting. For instance, the results of Golub and Jackson (2010) imply that when agents move simultaneously and every agent weights every other agent equally, society converges to complete learning as the size of the network grows. But as we show in Section 4.2, society does not learn completely in the uniform weighted network where each agent is connected to every predecessor with the same weight.

The idea behind the proof is that if there were some i and $\epsilon > 0$ such that $\mathbb{I}(i \rightarrow n) > \epsilon$ for infinitely many n, then i exerts at least ϵ influence on all these future players. Since $\tilde{s}_i$

is unbounded, there is a rare but positive probability event where i gets such a strong but wrong private signal that any future player who puts ϵ weight on $\tilde{s}_i$ and $(1 - \epsilon)$ weight on other signals would come to believe in the wrong state of the world with high probability. But this would mean infinitely many players have a high probability of believing in the wrong state of the world, so society fails to learn completely. To gain an intuition for the converse, first observe that $\tilde{a}_n = \|\vec{b}_n\|_1 \sum_{i=1}^n \mathbb{I}(i \to n)\tilde{s}_i$. In the event that $\omega = 1$, the mean of $\tilde{a}_n$ converges to infinity with n. So, provided the variance of $\tilde{a}_n$ is small relative to its mean, $\tilde{a}_n$ will converge to infinity in probability and society will learn completely. Since the log signals ($\tilde{s}_i$) are i.i.d., the variance of $\tilde{a}_n$ is small relative to its mean precisely when all of the weights $\mathbb{I}(i \to n)$ in the summand are small; this is guaranteed by the condition on influence $\lim_{n\to\infty} \mathbb{I}(i \to n) = 0$.

We now argue, both analytically and through an example, that the condition for complete learning in Proposition 2 is violated by a large class of weighted networks. The *out-degree* of an agent i is $\sum_{j<i} M_{i,j}$, interpreted as the total number of neighbors who directly affect i's play. We first show that on any network where all but finitely many agents have out-degree at least $1 + \epsilon$ for some $\epsilon > 0$, complete learning fails.

PROPOSITION 3. *Suppose there exists $\epsilon > 0$ so that $\sum_{j<i} M_{i,j} \geq 1 + \epsilon$ for all except finitely many agents i. Then society does not learn completely.*

The proof establishes that such a network satisfies the connectedness condition, but the influence of at least one of the early agents does not converge to 0, so complete learning fails by Proposition 2. The intuition is that if an influential later agent has an out-degree greater than 1, then the earlier agents whose action indirectly affects her must have even more influence. We construct a correspondence between weighted paths ending at early agents and weighted paths ending at later agents.

The condition in Proposition 3 is satisfied by all of the weighted networks studied in Section 4, which all feature mislearning with positive probability. The network in Remark 3 also satisfies the condition in Proposition 3 and almost surely leads to non-convergence of beliefs.

As an additional example, consider a network where network weights decay exponentially in distance, so $M_{i,j} = \delta^{i-j}$ for some $\delta \geq 0$. When the rate of decay is strictly above the threshold of $\frac{1}{2}$, late enough agents have out-degree bounded away from 1, so society does not learn completely by Proposition 3. When the rate of decay is strictly below the same threshold, we can show the connectedness condition fails and agents' beliefs do not converge to ω due to lack of information. At the threshold value of $\delta = \frac{1}{2}$, private signals of all predecessors are given equal weight, so the law of large numbers implies complete learning. This highlights the fragility of complete learning in our model.

EXAMPLE 1. Suppose $M_{i,j} = \delta^{i-j}$ for some $\delta \geq 0$. Society learns completely if and only if $\delta = \frac{1}{2}$.   ◇

Details of the arguments are provided in the Appendix.

4. Probability of mislearning and network structure

In this section, we compare the probability of mislearning in networks where complete learning fails by Proposition 2. To do so, we first derive a formula for the probability of mislearning as a function of the observation network. Then applying this expression to several canonical network structures, we compute comparative statics of this probability with respect to network parameters.

The first network structure we consider assigns the same weight to each link. Next, we study a homophilic network structure with agents split into two groups, allowing different weights on links within groups and between groups.

4.1 *Probability of mislearning*

Due to the Gaussian signal structure, we can give explicit expressions for the distributions of agent actions in each period. We show that the probability that agent n is correct about the state is related to the ratio of ℓ_1 norm to ℓ_2 norm of the vector of weighted path counts to n's predecessors, $\vec{b}_n := (b_{n,1}, \ldots, b_{n,n})$. The ratio $\|\vec{b}_n\|_1 / \|\vec{b}_n\|_2$ can be viewed as a measure of distributional equality for the vector of weights $\vec{b}_n$.[4] Indeed, among positive n-dimensional vectors $\vec{b}_n$ with $\|\vec{b}_n\|_1 = 1$, the ℓ_1/ℓ_2 ratio is minimized by the vector $\vec{b}_n = (1, 0, \ldots, 0)$ and maximized by the vector $\vec{b}_n = (\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n})$.

We can express in terms of the network structure the ex ante probability that agent n puts more confidence in the state being $\omega = 1$ when this is, in fact, true. This gives the key result that later lets us compare the probabilities of mislearning on different networks.

PROPOSITION 4. *On any network, the probability that agent n thinks the correct state is more likely than the incorrect one is*

$$\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1] = \Phi\left(\frac{1}{\sigma} \cdot \frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}\right),$$

where Φ is the standard Gaussian distribution function.

As $\frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}$ increases, the probability of agent n playing higher actions in state $\omega = 1$ also increases. In other words, the agent is more likely to be correct about the state when the vector of path counts is more evenly distributed. This should make intuitive sense as she is more likely to be correct when her action is the average of many independent signals with roughly equal weights, and less likely to be correct when her action puts disproportionally heavy weights on a few signals.

The proof of Proposition 4 first expresses $\tilde{a}_n = \sum_{i=1}^n b_{n,i} \tilde{s}_i$ using Proposition 1 and then observes that (s_i) are distributed i.i.d. $\mathcal{N}(1, \sigma^2)$ conditional on $\omega = 1$. This means $(\tilde{s}_i)$ are also conditionally i.i.d. Gaussian random variables, since the proof of Lemma 2 establishes that $\tilde{s}_i = 2s_i/\sigma^2$. As a sum of conditionally i.i.d. Gaussian random variables,

---

[4]In fact, the ratio of ℓ_1 to ℓ_2 norm has been used in the applied mathematics literature as a measure of normalized sparsity.

the action $\tilde{a}_n$ is itself Gaussian. The result follows from calculating the mean and variance of this sum.

For the remainder of this section, we study specific weighted networks where the ratio $\|\vec{b}_n\|_1/\|\vec{b}_n\|_2$ can be expressed in terms of interpretable network parameters. Our basic technique is to count paths on a given network using an appropriate recurrence relation, and then to apply Proposition 4. This allows us to relate network parameters to the probability distribution over learning outcomes.

4.2 Uniform weights

The simplest network we consider assigns the same weight $q \in [0, 1]$ to each feasible link. By varying the value of q, we can ask how link density affects the probability of mislearning, which we now define.

PROPOSITION 5. *Consider the q-uniform weighted network. When $0 < q \le 1$, almost surely agents' actions a_n converge to 0 or 1. The probability that society mislearns is*

$$\Phi\left(-\frac{1}{\sigma} \cdot \sqrt{\frac{q+2}{q}}\right).$$

This probability is strictly increasing in q.

The first statement of the proposition tells us that agents eventually agree on the state of the world and that these beliefs are arbitrarily strong after some time. These consensus beliefs need not be correct, however. The probability of society converging to incorrect beliefs is nonzero for all positive q, and increases in q. When the observational network is more densely connected, society is more likely to be wrong, as in Figure 1.

When the observation network is sparse (i.e., q is low), early agents' actions convey a large amount of independent information because they do not influence each other
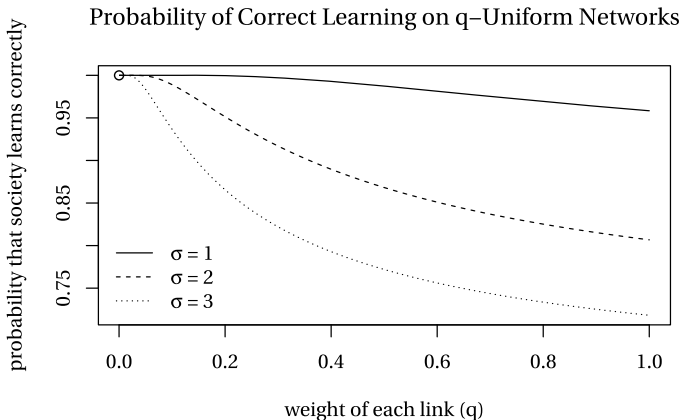

FIGURE 1. Probability that society learns correctly ($a_n \to \omega$) on a network where all feasible links have weight q. See Remark 5 for the discontinuity at $q = 0$.

too much. This facilitates later agents' learning. For high q, early agents' actions are highly correlated, so later naive agents cannot recover the true state as easily. A related intuition compares agents' beliefs about network structure to the actual network: as q grows, agents' beliefs about the network weights chosen by their neighbors differ more and more from the true weights. For small q, however, underweighting of social information partially mitigates the error due to Assumption 1. To complement this theoretical result, in a companion paper we conduct a sequential-learning experiment to evaluate a related comparative static (Dasaratha and He 2019). In line with the intuition above, we find that human subjects indeed exhibit lower long-run accuracy in the learning game when the density of the observation network increases.

The proof relies on the recurrence relation $b_{n,i} = (1 + q)b_{n-1,i}$. To see that this recurrence holds, let $\Psi[n \to i]$ be the set of all paths from n to i and let $\Psi[(n-1) \to i]$ be the set of all paths from $n-1$ to i. For each $\psi \in \Psi[(n-1) \to i]$ passing through agents $(n-1), j_1, j_2, \ldots, i$, we associate two paths $\psi', \psi'' \in \Psi[n \to i]$, with ψ' passing through $n, j_1, j_2, \ldots, i$, and ψ'' passing through $n, (n-1), j_1, j_2, \ldots, i$. This association exhaustively enumerates all paths in $\Psi[n \to i]$ as we consider all $\psi \in \Psi[(n-1) \to i]$. Path ψ' has the same weight as ψ since they have the same length, while path ψ'' has q fraction of the weight of ψ since it is longer by 1. This shows that the weight of all paths in $\Psi[n \to i]$ is equal to $1+q$ times the weight of all paths in $\Psi[(n-1) \to i]$; hence, $b_{n,i} = (1+q)b_{n-1,i}$.

REMARK 4. The case $q = 1$ is studied in Eyster and Rabin (2010), who use a slightly different signal structure. In their setting, Eyster and Rabin (2010) show that agents' beliefs converge to 0 or 1 almost surely and derive a nonzero lower bound on the probability of converging to the incorrect belief. By contrast, our result gives the exact probability of converging to the wrong belief for any $0 < q \leq 1$, under a Gaussian signal structure.

REMARK 5. There is a discontinuity at $q = 0$. As q approaches 0, the probability of society eventually learning correctly approaches 1. But when $q = 0$, each agent uses only her own private signal, so there is no social learning. This nonconvergence of actions means that society never learns correctly.

While we have focused on long-run learning accuracy, there is a trade-off between the speed of convergence and asymptotic accuracy for naive agents. The next proposition illustrates an extreme form of this trade-off. Start with a uniform-weights network with any link weight $0 < q^* \leq 1$. Sufficiently sparse uniform-weights networks will have worse accuracy than the q^*-uniform network for arbitrarily many early agents due to a lack of information aggregation. However, as implied by Proposition 5, late enough agents will have higher accuracy on these very sparse networks than on the q^*-uniform network.

PROPOSITION 6. *For any $0 < q^* \leq 1$ and $N \in \mathbb{N}$, there exists some $\bar{q} \in (0, q^*)$ so that $\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1]$ is strictly larger on the q^*-uniform weights network than the q-uniform weights network for all $2 \leq n \leq N$ and $q \in (0, \bar{q})$.*

FIGURE 2. First four agents in a two-groups network. Odd-numbered agents are in one group while even-numbered agents are in another group. Solid arrows have weight q_s and dashed arrows have weight q_d.

4.3 *Two groups*

We next consider a network with two groups and different weights for links within groups and between groups. By varying the link weights, we consider how homophily (i.e., segregation in communication) changes learning outcomes.

Odd-numbered agents are in one group and even-numbered agents are in a second group. Each feasible within-group link has weight q_s (s for same) and each feasible between-group link has weight q_d (d for different), so that for $i > j$, the link $M_{i,j} = q_s$ if $i \equiv j \pmod 2$ and $M_{i,j} = q_d$ otherwise. Figure 2 illustrates the first four agents in a two-group network.

We denote the probability of mislearning with weights q_s and q_d as $\xi(q_s, q_d)$.

PROPOSITION 7. *Consider the two-groups network with within-group link weight q_s and across-group link weight q_d. When $0 \le q_s \le 1$ and $0 < q_d \le 1$, almost surely agents' actions a_n converge to 0 or 1. The partial derivatives of the mislearning probability $\xi(q_s, q_d)$ satisfy*

$$\frac{\partial \xi}{\partial q_d} > \frac{\partial \xi}{\partial q_s} > 0,$$

i.e., the probability is increasing in q_s and q_d, but increasing q_d has a larger effect than increasing q_s.

The first statement again says that agents eventually agree on the state and eventually have arbitrarily strong beliefs. The fact that ξ is increasing in q_s and q_d is another example of higher link density implying more mislearning. The comparison $\partial \xi / \partial q_d > \partial \xi / \partial q_s$ tells us that more integrated (i.e., less homophilic) networks are more likely to herd on the wrong state of the world.

Convergence of beliefs is more subtle with two groups, as we might imagine the two homophilic groups holding different beliefs asymptotically. This does not happen because agents have continuous actions that allow them to precisely convey the strength of their beliefs. As such, eventually one group will develop sufficiently strong beliefs to convince the other given any arbitrarily weak connection $q_d > 0$ between groups. (In Section 5, however, we show that disagreement between two homophilic groups is possible with a coarser action space.)

To see that convergence must occur, observe that the belief of a later agent n depends mostly on the number of paths from that agent to early agents (and those agents' signal realizations). When n is large, most paths from agent n to an early agent pass between the two groups many times. So the number of paths does not depend substantially on agent n's group. Put another way, when $q_s \gg q_d > 0$ and n is large, agent n has many more length-1 paths to her own group than to the other group, but roughly the same total number of paths across all lengths to both groups. Therefore, agent n's belief does not depend substantially on whether n is in the odd group or the even group.[5]

Coleman's (1958) homophily index equals $(q_s - q_d)/(q_s + q_d)$ for this weighted network. To explore how homophily affects mislearning probability while holding fixed the average degree of each agent, we consider the total derivative $\frac{d}{d\Delta}\xi(q_s + \Delta, q_d - \Delta)$. To interpret, we are considering the marginal effect on mislearning of a Δ increase to all the within-group link weights, coupled with a Δ decrease to all the between-groups link weights. These two perturbations, applied simultaneously, leave each agent with roughly the same total degree and increases the homophily index by $(2\Delta)/(q_s + q_d)$. Using the chain rule and Proposition 7,

$$\frac{d}{d\Delta}\xi(q_s + \Delta, q_d - \Delta) = \frac{\partial\xi}{\partial q_s} - \frac{\partial\xi}{\partial q_d} < 0,$$

which means increasing the homophily index of the society and fixing average degrees always decreases the probability of mislearning. Note that this result holds regardless of whether society is currently homophilic ($q_s > q_d$) or heterophilic ($q_s < q_d$).

An important insight from the literature about social learning on networks is that beliefs converge more slowly on more segregated networks (Golub and Jackson 2012). In our model, faster convergence of beliefs tends to imply a higher probability of incorrect beliefs. When beliefs converge quickly, agents are putting far too much weight on early movers, while when beliefs converge more slowly, agents wait for more independent information. Since agents eventually agree, segregation helps society form strong beliefs more gradually.

5. Disagreement

In Section 4.3, we saw that even on partially segregated networks, agents eventually reach a consensus on the state of the world. This agreement relies crucially on the richness of the action space available to agents, which allows agents to communicate the strength of their beliefs. In this section, we modify our model so that the action space is binary and show that the two groups can disagree forever about the state of the world even when the number of connections across the groups is unbounded.

The contrasting results for the binary-actions model versus the continuum-actions model echo a similar contrast in the rational-herding literature, where society herds on the wrong action with positive probability when actions coarsely reflect beliefs (Banerjee

---

[5]Each path transitions between the two groups, and eventually the probability of ending in a given group is approximately independent of the starting group. This is analogous to a Markov chain approaching its stationary distribution.

1992, Bikhchandani et al. 1992), but almost surely converges to the correct action when the action set is rich enough (Lee 1993). Interestingly, while the rational-herding literature finds that an unboundedly informative signal structure prevents herding on the wrong action even when actions coarsely reflect beliefs (Smith and Sørensen 2000), we will show below that even with Gaussian signals, two groups may disagree with positive probability.

Suppose that the state of the world and the signal structure are the same as in Section 2, but agents now choose binary actions $a_i \in \{0, 1\}$. Agents still maximize the expectation of $u_i(a_i, \omega) := -(a_i - \omega)^2$ given their beliefs about ω, under the psychological errors given by Assumptions 1 and 2. This utility function now implies that an agent chooses the action corresponding to the state of world she believes is more likely. Agents live on the two-groups network from Section 4.3: for $i > j$, the link $M_{i,j} = q_s$ if $i \equiv j \pmod 2$ and $M_{i,j} = q_d$ otherwise. We assume $q_s > q_d > 0$, so that agents have stronger connections with predecessors from their own groups.

PROPOSITION 8. *Consider the two-groups network. Suppose $q_s > q_d > 0$ and agents play binary actions. Then there is a positive probability that all odd-numbered agents choose action 0 while all even-numbered agents choose action 1.*

Persistent disagreement is sustained even though agent n has approximately $nq_d/2$ weighted links to agents from the other group (when n is large) taking opposite actions.

Our result extends to two groups of unequal sizes as long as for all later agents, the total number of weighted links to their own group is larger than the total number of weighted links to the other group.

We also get the same result on a random-network analog of the two-groups model, where edges are unweighted and q_s is the probability of link formation within groups while q_d is the probability of link formation between groups. Agents observe only the actions of the predecessors to whom they are linked and wrongly believe all observed actions derive from private signals.[6] By contrast, with rational agents, Theorem 2 of Acemoglu et al. (2011) implies the groups almost surely agree asymptotically on this random network.

This result adds a new mechanism to the literature on disagreement in connected societies (Acemoglu et al. 2013, 2016, Sethi and Yildiz 2012). Bohren and Hauser (2018) also study disagreement in a binary sequential-learning setting with behavioral agents, but their results concern disagreement on a complete network among agents with different types of behavioral biases. By contrast, our Proposition 8 says that when all agents use the same naive heuristic, they can still disagree by virtue of belonging to two different homophilic social groups, even when there are many connections between those groups.

---

[6]Details of the statement and a proof are available in a previous draft at https://arxiv.org/pdf/1703.02105v5.pdf.

6. Conclusion

In this paper, we have explored the influence of network structures on learning outcomes when agents move sequentially and use a log-linear learning rule due to inferential naiveté. We have compared long-run welfare across networks by deriving the exact probabilities of mislearning on arbitrary networks.

We have studied the simplest possible social-learning environment to focus on the effect of network structure, but several extensions are straightforward. Analogs of our general results hold for finite state spaces with more than two elements, where we can define a log-likelihood ratio for each pair of states. We can also make the order of moves random and unknown, in which case naive behavior conditional on a given turn order is the same as when that order is certain.

We prove our comparative statics results for weighted networks as they are analytically more tractable than random graphs. For each weighted network with weights in $[0, 1]$, there corresponds an analogous (unweighted) random graph model where the $i \rightarrow j$ link exists with probability $M_{i,j}$. In numerical simulations, all comparative statics results proved for weighted networks in Section 4 continue to hold in the analogous random network models. The major obstacle to extending our proofs is that because our networks are directed and acyclic, the relevant adjacency matrices have no nonzero eigenvalues. As a consequence, most techniques from spectral random graph theory do not apply (but perhaps other methods would).

Appendix: Proofs

A.1 *Proof of Lemma 1*

The log-likelihood ratio of state $\omega = 1$ and state $\omega = 0$ conditional on the signal realizations of i's linked predecessors is

$$\ln\left(\frac{\mathbb{P}[\omega = 1|s_i, (s_j)_{j:M_{i,j}=1}]}{\mathbb{P}[\omega = 0|s_i, (s_j)_{j:M_{i,j}=1}]}\right)$$

$$= \ln\left(\frac{\mathbb{P}[s_i, (s_j)_{j:M_{i,j}=1}|\omega = 1]}{\mathbb{P}[s_i, (s_j)_{j:M_{i,j}=1}|\omega = 0]}\right) \quad \text{(two states equally likely)}$$

$$= \ln\left(\frac{\mathbb{P}[s_i|\omega = 1]}{\mathbb{P}[s_i|\omega = 0]} \cdot \prod_{j:M_{i,j}=1} \frac{\mathbb{P}[s_j|\omega = 1]}{\mathbb{P}[s_j|\omega = 0]}\right) \quad \text{(by independence)}$$

$$= \ln\left(\frac{\mathbb{P}[\omega = 1|s_i]}{\mathbb{P}[\omega = 0|s_i]}\right) + \sum_{j:M_{i,j}=1} \ln\left(\frac{\mathbb{P}[\omega = 1|s_j]}{\mathbb{P}[\omega = 0|s_j]}\right)$$

$$= \tilde{s}_i + \sum_{j:M_{i,j}=1} \tilde{s}_j$$

$$= \tilde{s}_i + \sum_{j<i} M_{i,j}\tilde{s}_j.$$

Due to Assumption 1, i thinks each predecessor j must have received signal s_j such that $\tilde{s}_j = \tilde{a}_j$. When i observes only the play of linked predecessors, her log-likelihood ratio of state $\omega = 1$ and state $\omega = 0$ given her social observations and private signal is therefore $\tilde{s}_i + \sum_{j<i} M_{i,j}\tilde{a}_j$. She maximizes her expected payoff by choosing an action a_i corresponding to her belief in state $\omega = 1$, which implies that $\tilde{a}_i$ is equal to this log-likelihood ratio. $\square$

A.2 *Proof of Lemma 2*

We first establish an auxiliary lemma.

LEMMA 3. *We have $\tilde{s}_i = 2s_i/\sigma^2$.*

PROOF. The log-likelihood ratio is

$$\ln\left(\frac{\mathbb{P}[\omega = 1|s_i]}{\mathbb{P}[\omega = 0|s_i]}\right) = \ln\left(\frac{\mathbb{P}[s_i|\omega = 1]}{\mathbb{P}[s_i|\omega = 0]}\right)$$

$$= \ln\left(\frac{\exp\left(\dfrac{-(s_i - 1)^2}{2\sigma^2}\right)}{\exp\left(\dfrac{-(s_i + 1)^2}{2\sigma^2}\right)}\right)$$

$$= \frac{-(s_i^2 - 2s_i + 1) + (s_i^2 + 2s_i + 1)}{2\sigma^2} = 2s_i/\sigma^2. \qquad \square$$

We now turn to the proof of Lemma 2.

Due to Assumptions 1, i thinks that j will choose a_j such that $\tilde{a}_j = 2s_j/\sigma^2$ by the result just established, since j thinks the conditional variance of her signal is σ^2. But since i believes j's signal has conditional variance $\sigma^2/M_{i,j}$ by Assumption 2, in i's view

$$\ln\left(\frac{\mathbb{P}[\omega = 1|s_j]}{\mathbb{P}[\omega = 0|s_j]}\right) = \frac{2s_j}{\sigma^2/M_{i,j}} = M_{i,j}\tilde{a}_j,$$

again applying the result above.

Omitting analogous algebraic arguments as in the proof of Lemma 1,

$$\ln\left(\frac{\mathbb{P}[\omega = 1|s_i, (s_j)_{j<i}]}{\mathbb{P}[\omega = 0|s_i, (s_j)_{j<i}]}\right) = \ln\left(\frac{\mathbb{P}[\omega = 1|s_i]}{\mathbb{P}[\omega = 0|s_i]}\right) + \sum_{j<i}\ln\left(\frac{\mathbb{P}[\omega = 1|s_j]}{\mathbb{P}[\omega = 0|s_j]}\right)$$

$$= \tilde{s}_i + \sum_{j<i} M_{i,j}\tilde{a}_j.$$

So $\tilde{s}_i + \sum_{j<i} M_{i,j}\tilde{a}_j$ is i's log-likelihood ratio of state $\omega = 1$ and state $\omega = 0$ given her social observations and private signal. $\square$

A.3 *Proof of Proposition 1*

By weighted log-linear inference, for each i we have $\tilde{a}_i = \tilde{s}_i + \sum_{j \in N_i} M_{i,j} \tilde{a}_j$. In vector notation, we therefore have

$$
\begin{pmatrix} \tilde{a}_1 \\ \vdots \\ \tilde{a}_n \end{pmatrix} = \begin{pmatrix} \tilde{s}_1 \\ \vdots \\ \tilde{s}_n \end{pmatrix} + M[n] \cdot \begin{pmatrix} \tilde{a}_1 \\ \vdots \\ \tilde{a}_n \end{pmatrix}.
$$

Algebra then yields the desired expression. Note that $(I - M[n])$ is invertible because $M[n]$ is lower triangular with all diagonal entries equal to 0.

To see the path-counting interpretation, write $(I - M[n])^{-1} = \sum_{k=0}^{\infty} M[n]^k$. Here, $(M[n]^k)_{i,j}$ counts the number of weighted paths of length k from i to j. $\qquad \square$

A.4 *Proof of Proposition 2*

Without loss of generality, assume $\omega = 1$. (The case of $\omega = 0$ is exactly analogous and is omitted.) Note that a_n converges in probability to 1 if and only if $\tilde{a}_n$ converges in probability to ∞.

First suppose that $\lim_{n \to \infty} \mathbb{I}(j \to n) \neq 0$ for some j. Then there exists $\epsilon > 0$ such that $\mathbb{I}(j \to n) > \epsilon$ for infinitely many n. For each such n, the probability that agent n chooses an action with $\tilde{a}_n < 0$ is equal to the probability that $\sum_{i=1}^{n} \mathbb{I}(i \to n) \tilde{s}_i$ is negative.

Because s is Gaussian, $\tilde{s}$ has finite variance, so we can find positive constants C and δ independent of n such that $\sum_{i \neq j} \mathbb{I}(i \to n) \tilde{s}_i < C$ with probability at least δ (for example, by applying Markov's inequality to $|\tilde{s}_i|$). Then agent n will be wrong if $\tilde{s}_j < -C/\epsilon$, which is a positive probability event since $\tilde{s}$ is unbounded. So the probability that an agent n such that $\mathbb{I}(j \to n) > \epsilon$ chooses $\tilde{a}_n < 0$ is bounded from below by a positive constant.

For the converse, suppose that $\lim_{n \to \infty} \mathbb{I}(i \to n) = 0$ for all i. By the independence of the log signals $\tilde{s}_i$, the log action $\tilde{a}_n = \sum_{i=1}^{n} \mathbb{I}(i \to n) \tilde{s}_i$ is a random variable with mean $\|\vec{b}_n\|_1$ and standard deviation $\|\vec{b}_n\|_2 \sigma$ when $\omega = 1$. We now use the connectedness condition to show that $\|\vec{b}_n\|_1 / \|\vec{b}_n\|_2 \to \infty$.

Find N and $C \leq 1$ as in the connectedness condition. For each $\epsilon > 0$, we can choose M_ϵ such that $\mathbb{I}(i \to n) < \epsilon$ whenever $i < N$ and $n > M_\epsilon$ by the hypothesis $\lim_{n \to \infty} \mathbb{I}(i \to n) = 0$ applied to the finitely many members $i < N$. But now for any $j \geq N$ and any $n > \max(j, M_\epsilon)$, concatenating a path from j to n with a path from i to j gives a path from i to n whose weight is the product of the weights of the two subpaths. This shows $b_{n,i} \geq b_{j,i} \cdot b_{n,j}$, which implies $\mathbb{I}(i \to n) \geq \mathbb{I}(j \to n) \cdot b_{j,i}$. We have $\mathbb{I}(j \to n) \leq \min_{i<N} \mathbb{I}(i \to n)/b_{j,i}$, where $b_{j,i} \geq C$ for at least one $i < N$ by the connectedness condition. This shows for any $j \in \mathbb{N}$ and for $n > M_\epsilon$, that we get $\mathbb{I}(j \to n) \leq \epsilon/C$.

We have for all $n > M_\epsilon$,

$$
\frac{\|\vec{b}_n\|_2}{\|\vec{b}_n\|_1} \leq \max_j \frac{\sqrt{\|\vec{b}_n\|_1 \cdot b_{n,j}}}{\|\vec{b}_n\|_1} = \max_{j<n} \sqrt{\mathbb{I}(j \to n)} < \sqrt{\epsilon/C}.
$$

Because $\epsilon > 0$ is arbitrary, $\|\vec{b}_n\|_1 / \|\vec{b}_n\|_2$ converges to infinity.

Let some $K > 0$ be given. We now show that $\mathbb{P}[\tilde{a}_n < K|\omega = 1] \to 0$, hence proving that $\tilde{a}_n$ converges to ∞ in probability. We compute

$$z_n := \frac{\mathbb{E}[\tilde{a}_n|\omega = 1] - K}{\text{Std}[\tilde{a}_n|\omega = 1]} = \frac{\|\vec{b}_n\|_1 \cdot \dfrac{2}{\sigma^2}}{\|\vec{b}_n\|_2 \cdot \dfrac{2}{\sigma}} - \frac{K}{\|\vec{b}_n\|_2 \cdot \dfrac{2}{\sigma}}.$$

Since $\|\vec{b}_n\|_1/\|\vec{b}_n\|_2 \to \infty$, the first term converges to infinity. By the connectedness condition, $\|\vec{b}_n\|_2 \geq C$ for all large enough n, so the second term is bounded. This implies $z_n \to \infty$. By Chebyshev's inequality, $\mathbb{P}[\tilde{a}_n < K|\omega = 1] \leq z_n^{-2}$. This shows $\mathbb{P}[\tilde{a}_n < K|\omega = 1] \to 0$.

We note that this shows convergence in probability, but does not characterize the joint distribution of actions, so these methods do not guarantee almost sure convergence (without further structure on the networks as in Section 4). □

A.5 *Proof of Proposition 3*

By the hypothesis of the proposition, there exists some $\epsilon > 0$ and $N \in \mathbb{N}$ so that for all $i > N$, $\sum_{j<i} M_{i,j} \geq 1 + \epsilon$. Modify the network to set all links originating from any of the first N agents to have weight 0, that is, $M_{i,j} = 0$ for all $i, j \leq N$.

We prove by induction that $\sum_{j\leq N} b_{i,j} \geq 1 + \epsilon$ for all $i \geq N + 1$ on the modified network. Consider agent $N + 1$. Since $\sum_{j<N+1} M_{N+1,j} \geq 1 + \epsilon$ and all of $(N+1)$'s out-degree comes from links to agents in position N or earlier, $\sum_{j\leq N} b_{N+1,j} \geq 1 + \epsilon$. By induction, suppose $\sum_{j\leq N} b_{N+k,j} \geq 1 + \epsilon$ holds for all $1 \leq k \leq K$. A lower bound on $\sum_{j\leq N} b_{N+K+1,j}$ is

$$\sum_{j\leq N} M_{N+K+1,j} + \sum_{N+1\leq i\leq N+K} M_{N+K+1,i} \cdot \left(\sum_{j'\leq N} b_{i,j'}\right)$$

$$\geq \sum_{j\leq N} M_{N+K+1,j} + \sum_{N+1\leq i\leq N+K} M_{N+K+1,i} \cdot (1 + \epsilon)$$

$$\geq \sum_{j\leq N+K} M_{N+K+1,j}$$

$$\geq 1 + \epsilon,$$

where in the first inequality we used the inductive hypothesis and in the last inequality we used the fact that $N + K + 1$ has an out-degree of at least $1 + \epsilon$. This establishes $\sum_{j\leq N} b_{N+K+1,j} \geq 1 + \epsilon$ and so by induction, $\sum_{j\leq N} b_{N+k,j} \geq 1 + \epsilon$ for all $k \geq 1$.

This result holds a fortiori on the original network with higher link weights. By the pigeonhole principle, the original network satisfies the connectedness condition with $C = \frac{1}{N}(1 + \epsilon)$.

Now return to the modified network (so M refers to the possibly modified network weights). We develop some notation for the rest of the proof and establish an intermediary lemma. For $j < i$, let $\Psi[i \to j]$ be the set of all paths from i to j. Let $\hat{\Psi}[i \to [N]]$ be the

set of paths from i to some agent $k \leq N$, such that the path contains no links between two different agents among the first N. Let $\hat{\Psi}[i \to [N] \mid j]$ be the subset of such paths that pass through j.

For a path ψ passing through agents $i_1, i_2, \ldots, i_L$, let $W(\psi) := \prod_{\ell=1}^{L-1} M_{i_{\ell+1}, i_\ell}$ denote its weight and let $D(\psi) := \prod_{\ell=1}^{L-1} (\sum_{j < i_\ell} M_{i_\ell, j})$ denote the product of out-degrees of all agents on the path except the last one.

LEMMA 4. *For $n > N$, $\sum_{\psi \in \hat{\Psi}[n \to [N]]} \frac{W(\psi)}{D(\psi)} = 1$.*

PROOF. We prove by induction on n. For $n = N + 1$, the set $\hat{\Psi}[n \to [N]]$ is the set of N paths each consisting of a link from $N + 1$ to some agent $j \leq N$. Each $\psi \in \hat{\Psi}[n \to [N]]$ therefore has $D(\psi) = \sum_{j < N+1} M_{N+1, j}$, and the path terminating at j has $W(\psi) = M_{N+1, j}$. So the claim holds for $n = N + 1$. By induction suppose it holds for all $n \leq N + K$ for some $K \geq 1$. For $n = N + K + 1$, partition $\hat{\Psi}[n \to [N]]$ into $K + 1$ groups. For $1 \leq k \leq K$, each path $\psi \in \Psi(k)$ in the kth group consists of the link $n \to (N + k)$ concatenated in front of a path $\psi' \in \hat{\Psi}[(N + k) \to [N]]$, so $\psi = ((n, N + k), \psi')$. The final $(K + 1)$th group consists of paths where n links directly to an agent among the first N. We have

$$
\sum_{\psi \in \hat{\Psi}[n \to [N]]} \frac{W(\psi)}{D(\psi)} = \sum_{k=1}^{K} \left(\sum_{\psi' \in \hat{\Psi}[(N+k) \to [N]]} \frac{W((n, N+k), \psi')}{D((n, N+k), \psi')} \right) + \sum_{j=1}^{N} \frac{M_{n,j}}{\sum_{h<n} M_{n,h}}
$$

$$
= \sum_{k=1}^{K} \left(\sum_{\psi' \in \hat{\Psi}[(N+k) \to [N]]} \frac{M_{n,N+k} \cdot W(\psi')}{\sum_{h<n} M_{n,h} \cdot D(\psi')} \right) + \sum_{j=1}^{N} \frac{M_{n,j}}{\sum_{h<n} M_{n,h}}
$$

$$
= \sum_{k=1}^{K} \left(\frac{M_{n,N+k}}{\sum_{h<n} M_{n,h}} \cdot 1 \right) + \sum_{j=1}^{N} \frac{M_{n,j}}{\sum_{h<n} M_{n,h}} \quad \text{(by inductive hypothesis)}
$$

$$
= \frac{\sum_{h<n} M_{n,h}}{\sum_{h<n} M_{n,h}} = 1.
$$

So by induction, this claim holds for all $n > N$.   □

We now return to the proof of Proposition 3. For $N < i < n$,

$$
b_{n,i} = \sum_{\psi \in \Psi[n \to i]} W(\psi) = \sum_{\psi \in \Psi[n \to i]} \left[W(\psi) \cdot \left(\sum_{\hat{\psi} \in \hat{\Psi}[i \to [N]]} \frac{W(\hat{\psi})}{D(\hat{\psi})} \right) \right],
$$

where the second equality follows because Lemma 4 implies the term in the inner parentheses is 1. For a path ψ passing through i, let $\psi[i]$ denote the subpath starting with i. So

the above says

$$b_{n,i} = \sum_{\psi \in \hat{\Psi}[n \to [N]|i]} \frac{W(\psi)}{D(\psi[i])}.$$

Summing across i, we may re-index the sum by paths in $\hat{\Psi}[n \to [N]]$. To be more precise, for $\psi \in \hat{\Psi}[n \to [N]]$, write $A(\psi) \subseteq \{N+1, \dots, n-1\}$ to be the set of agents that ψ passes through. For each $j \in A(\psi)$, we have $\psi \in \hat{\Psi}[n \to [N] \mid j]$, so it contributes $W(\psi)/D(\psi[j])$ to the overall sum,

$$\sum_{i=N+1}^{n-1} b_{n,i} = \sum_{i=N+1}^{n-1} \sum_{\psi \in \hat{\Psi}[n \to [N]|i]} \frac{W(\psi)}{D(\psi[i])}$$

$$= \sum_{\psi \in \hat{\Psi}[n \to [N]]} \sum_{j \in A(\psi)} \frac{W(\psi)}{D(\psi[j])}$$

$$\leq \sum_{\psi \in \hat{\Psi}[n \to [N]]} W(\psi) \cdot \sum_{j \in A(\psi)} \frac{1}{(1+\epsilon)^{|\psi[j]|-1}}$$

$$\leq \sum_{\psi \in \hat{\Psi}[n \to [N]]} W(\psi) \cdot \frac{1}{\epsilon},$$

where $|\psi[j]|$ denotes the number of agents in the subpath $\psi[j]$. On the third line, we used the fact that all agents on ψ except the last one must have out-degree at least $1+\epsilon$, so $D(\psi[j]) \geq (1+\epsilon)^{|\psi[j]|-1}$. The result $\sum_{i=N+1}^{n-1} b_{n,i} \leq \sum_{\psi \in \hat{\Psi}[n \to [N]]} W(\psi) \cdot \frac{1}{\epsilon}$ also holds for the original network, since we have not modified the subnetwork among agents $N+1, \dots, n$.

We also have $\sum_{i=1}^{N} b_{n,i} = \sum_{\psi \in \hat{\Psi}[n \to [N]]} W(\psi)$. On the original network, we have higher link weights among the first N agents, so we in fact have

$$\sum_{i=1}^{N} b_{n,i} \geq \sum_{\psi \in \hat{\Psi}[n \to [N]]} W(\psi).$$

So, on the original network,

$$\sum_{i=1}^{N} \mathbb{I}(i \to n) \geq \frac{1}{1 + 1/\epsilon}.$$

This inequality holds for every n, so it cannot be the case that $\lim_{n \to \infty} \mathbb{I}(i \to n) = 0$ for all $1 \leq i \leq N$. □

A.6 *Proof of Example 1*

The coefficients $b_{i,n}$ satisfy the recurrence relation $b_{n,i} = 2\delta b_{n-1,i}$ whenever $n - i > 1$.

When $\delta = \frac{1}{2}$, from the recurrence relation, all predecessors' signals are given equal weight, so by the law of large numbers, actions converge to ω almost surely.

When $\delta > \frac{1}{2}$, $\sum_{k=0}^{\infty} \delta^k > 1$, so $\sum_{j<i} M_{i,j}$ is bounded above 1 for n large enough. So by Proposition 3, society does not learn correctly.

The final case is $\delta < \frac{1}{2}$. We show that $\mathbb{P}[a_n \le \frac{1}{2} \mid \omega = 1]$ is bounded away from 0 for all $n \ge 1$, so (a_n) cannot converge in probability to ω.

Without loss of generality, normalize to $\sigma = 1$. From the recurrence relation for the coefficients $b_{n,i}$, it is easy to check that $\tilde{a}_n = 2\delta\tilde{a}_{n-1} + \tilde{s}_n - \delta\tilde{s}_{n-1}$ for each n. Evidently $b_{i+1,i} = \delta$, so from recursion, $b_{n,i} = (2\delta)^{n-i-1}\delta$ for $i \le n-1$, $b_{n,n} = 1$. So $\tilde{a}_n = \tilde{s}_n + \delta \sum_{j=0}^{n-2}(2\delta)^j \cdot \tilde{s}_{n-1-j}$, meaning

$$\tilde{a}_n \mid (\omega = 1) \sim \mathcal{N}\left(1 + \delta\frac{1 - (2\delta)^{n-1}}{1 - 2\delta}, 1 + \delta^2\frac{1 - (4\delta^2)^{n-1}}{1 - 4\delta^2}\right).$$

So

$$\mathbb{P}[\tilde{a}_n \le 0 \mid \omega = 1] \ge \Phi\left(-1 - \frac{\delta}{1 - 2\delta}\right)$$

for all n, which implies for all n,

$$\mathbb{P}\left[a_n \le \frac{1}{2} \,\middle|\, \omega = 1\right] \ge \Phi\left(-1 - \frac{\delta}{1 - 2\delta}\right). \qquad \square$$

A.7 Proof of Proposition 4

We first state and prove a lemma that gives the ex ante distribution of agent n's log action.

LEMMA 5. *When $\omega = 1$, the log action of agent n on any weighted network is distributed as*

$$\tilde{a}_n \sim \mathcal{N}\left(\frac{2}{\sigma^2}\|\vec{b}_n\|_1, \frac{4}{\sigma^2}\|\vec{b}_n\|_2^2\right).$$

PROOF. By Proposition 1, $\tilde{a}_n = \sum_{i=1}^{n} b_{n,i}\tilde{s}_i$. This is equal to $\sum_{i=1}^{n} 2b_{n,i}s_i/\sigma^2$ according to Lemma 3. Conditional on $\omega = 1$, (s_i) are i.i.d. $\mathcal{N}(1, \sigma^2)$ random variables, so

$$\sum_{i=1}^{n} \frac{2b_{n,i}}{\sigma^2}s_i \sim \mathcal{N}\left(\frac{2}{\sigma^2}\sum_{i=1}^{n} b_{n,i}, \frac{4}{\sigma^2}\sum_{i=1}^{n} b_{n,i}^2\right) = \mathcal{N}\left(\frac{2}{\sigma^2}\|\vec{b}_n\|_1, \frac{4}{\sigma^2}\|\vec{b}_n\|_2^2\right). \qquad \square$$

Now we give the proof of Proposition 4.

By Lemma 5, $\tilde{a}_n|(\omega = 1) \sim \mathcal{N}(\frac{2}{\sigma^2}\|\vec{b}_n\|_1, \frac{4}{\sigma^2}\|\vec{b}_n\|_2^2)$. So using properties of the Gaussian distribution,

$$\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1] = \Phi\left(\frac{\frac{2}{\sigma^2}\|\vec{b}_n\|_1}{\frac{2}{\sigma}\|\vec{b}_n\|_2}\right) = \Phi\left(\frac{1}{\sigma} \cdot \frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}\right). \qquad \square$$

A.8 *Proof of Proposition 5*

The numbers of paths from various agents to agent i satisfy the recurrence relation $b_{n,i} = (1+q)b_{n-1,i}$ when $n-i > 1$. By a simple computation, we find that

$$\tilde{a}_n = \sum_{i=1}^{n-1} q(1+q)^{n-i-1}\tilde{s}_i + \tilde{s}_n.$$

Since $\tilde{s}_i$ are independent Gaussian random variables, our argument uses the fact that for n large, $\tilde{a}_n$ has the same sign as another Gaussian random variable, whose mean and variance we can compute.

We first show that $\tilde{a}_n$ converges to $-\infty$ or ∞ almost surely. Consider the random variable

$$X_n(\tilde{s}) := \frac{1}{2}\sum_{i=1}^{n-1}(1+q)^{-i}\tilde{s}_i,$$

where $\tilde{s} := (s_i)_{i=1}^{\infty}$ is the profile of private signal realizations. By a standard result, $X_n(\tilde{s})$ converges almost surely to a random variable $Y(\tilde{s})$ such that the conditional distribution of Y in each state of the world is Gaussian. For each n, $\tilde{a}_n(\tilde{s}) = 2q(1+q)^{n-1} \cdot X_n(\tilde{s}) + \tilde{s}_n$. Since $\sum_{n=1}^{\infty}\mathbb{P}[\tilde{s}_n > n] < \infty$, by the Borel–Cantelli lemma, $\mathbb{P}[\tilde{s}_n > n$ infinitely often$] = 0$. So almost surely, $\lim_{n\to\infty}\tilde{a}_n(\tilde{s}) = \lim_{n\to\infty} 2q(1+q)^{n-1} \cdot Y(\tilde{s}) + \tilde{s}_n \in \{-\infty, \infty\}$. This in turn shows that a_n converges to 0 or 1 almost surely.

Now we show $\mathbb{P}[a_n \to 0 \mid \omega = 1] = \Phi(-\sigma^{-1}\sqrt{(q+2)/q})$, which is the same probability as $\mathbb{P}[\tilde{a}_n \to -\infty \mid \omega = 1]$. The random variable $Y(\tilde{s})$ that $X_n(\tilde{s})$ converges to a.s. has the distribution $\mathcal{N}(1/(\sigma^2 q), 1/(\sigma^2 q(q+2)))$ when $\omega = 1$, and $\tilde{a}_n$ has the same sign as $X_n(\tilde{s})$ with probability converging to 1 for n large. The distribution $\mathcal{N}(1/(\sigma^2 q), 1/(\sigma^2 q(q+2)))$ assigns $\Phi(-\sigma^{-1}\sqrt{(q+2)/q})$ probability to the positive region. The symmetric argument holds for $\omega = 0$. $\qquad\square$

A.9 *Proof of Proposition 6*

First, we derive a closed-form expression for the probability that the nth agent thinks the correct state is more likely in the uniform weights network, conditional on $\omega = 1$.

LEMMA 6. *In the q-uniform weights network,*

$$\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1] = \Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2n-2}}}\right).$$

This probability is strictly increasing in n when $0 < q \le 1$.

PROOF. From the proof of Proposition 5, we have

$$\tilde{a}_n = \sum_{i=1}^{n-1} q(1+q)^{n-i-1}\tilde{s}_i + \tilde{s}_n,$$

where the different $\tilde{s}_i$s are conditionally independent given $\omega = 1$, with $\tilde{s}_i \mid (\omega = 1) \sim \mathcal{N}(2/\sigma^2, 4/\sigma^2)$ from Lemma 3. Thus, the sum $\tilde{a}_n$ is conditionally Gaussian with a mean of

$$\frac{2}{\sigma^2} \cdot \left[1 + \sum_{i=1}^{n-1} q(1+q)^{n-i-1} \right] = \frac{2}{\sigma^2} \cdot \left[1 + q \cdot \frac{(1+q)^{n-1} - 1}{(1+q) - 1} \right]$$

$$= \frac{2}{\sigma^2} \cdot (1+q)^{n-1}$$

and a variance of

$$\frac{4}{\sigma^2} \cdot \left[1 + \sum_{i=1}^{n-1} q^2(1+q)^{2n-2i-2} \right] = \frac{4}{\sigma^2} \cdot \left[1 + q^2 \cdot \frac{(1+q)^{2n-2} - 1}{(1+q)^2 - 1} \right]$$

$$= \frac{4}{\sigma^2} \cdot \frac{2 + (1+q)^{2n-2}q}{2 + q}.$$

Thus, 0 is

$$\frac{\dfrac{2}{\sigma^2} \cdot (1+q)^{n-1}}{\sqrt{\dfrac{4}{\sigma^2} \cdot \dfrac{2 + (1+q)^{2n-2}q}{2 + q}}} = \frac{1}{\sigma} \cdot \frac{(1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2n-2}}}$$

standard deviations below the mean in the distribution of $\tilde{a}_n \mid (\omega = 1)$, so

$$\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1] = \Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2n-2}}} \right).$$

To see that this expression is strictly increasing, let $n \geq 1$. Then

$$\mathbb{P}[\tilde{a}_{n+1} > 0 \mid \omega = 1] = \Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q) \cdot (1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + (1+q)^2 \cdot q(1+q)^{2n-2}}} \right)$$

$$> \Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q) \cdot (1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{(1+q)^2 \cdot 2 + (1+q)^2 \cdot q(1+q)^{2n-2}}} \right)$$

$$= \Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2n-2}}} \right)$$

$$= \mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1]$$

as desired.   □

Now we give the proof of Proposition 6.

Let $\mathbb{P}[\tilde{a}_2 > 0 \mid \omega = 1]$ on the q^*-uniform weights network be denoted by p. Lemma 6 implies $\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1]$ is strictly increasing in n on the q^*-uniform weights network, so $p > \Phi(1/\sigma)$ and, furthermore, $\mathbb{P}[\tilde{a}_n > 0 \mid \omega = 1] \geq p$ for all $n \geq 2$ on the same network.

The function

$$q \mapsto \Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q)^{N-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2N-2}}}\right)$$

is continuous and equals $\Phi(1/\sigma)$ when $q = 0$. So we may find a small enough $\bar{q} \in (0, q^*)$ so that whenever $0 < q < \bar{q}$,

$$\Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q)^{N-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2N-2}}}\right) < p.$$

From the monotonicity result of Lemma 6, this also implies

$$\Phi\left(\frac{1}{\sigma} \cdot \frac{(1+q)^{n-1} \cdot \sqrt{(2+q)}}{\sqrt{2 + q(1+q)^{2n-2}}}\right) < p$$

for all $2 \leq n \leq N$. $\qquad\square$

A.10 *Proof of Proposition 7*

Suppose we have two groups, and agents observe predecessors in the same group with weight q_s and predecessors in the other group with weight q_d. Then the coefficients $b_{n,i}$ satisfy the recurrence relation

$$b_{n,i} = q_d b_{n-1,i} + (1 + q_s) b_{n-2,i}$$

when $n - i > 2$. Since the network is translation invariant, $b_{n,i}$ only depends on $n - i$. By a standard algebraic fact, there exist constants c_+, c_-, ζ_+, ζ_- (only depending on $n - i$) so that

$$b_{n,i} = c_+ \zeta_+^{n-i} + c_- \zeta_-^{n-i},$$

where $\zeta_\pm$ are the solutions to the polynomial $x^2 - q_d x - (1 + q_s) = 0$ and c_+, c_- are constants that we can determine from $b_{2,1}$ and $b_{3,1}$. We compute

$$\zeta_\pm = \frac{q_d \pm \sqrt{4q_s + q_d + 4}}{2},$$

where $\zeta_+ > 1$ and $\zeta_- < 0$. By arguments analogous to those in the proof of Proposition 5, we may again establish that a_n converges to 0 or 1 almost surely. We now analyze the probability of mislearning.

Since $\zeta_+ > |\zeta_-|$, the exponential term with base ζ_+ dominates as n grows large. This shows $c_+ > 0$, since $b_{n,i}$ counts the number of weighted paths in a network so it must be

a positive number. This also shows that $\mathbb{P}[\tilde{a}_n < 0 \mid \omega = 1] \to \mathbb{P}[\sum_{i=0}^{\infty}(\zeta_+)^{-i}\tilde{s}_i < 0 \mid \omega = 1]$ as $n \to \infty$. Conditional on $\omega = 1$, the sum $\sum_{i=0}^{\infty}(\zeta_+)^{-i}\tilde{s}_i$ has the distribution

$$\mathcal{N}\left(\frac{2}{\sigma^2(\zeta_+ - 1)}, \frac{4}{\sigma^2(\zeta_+ - 1)(\zeta_+ + 1)}\right),$$

so it is easy to show that the probability assigned to the negative region is increasing in ζ_+.

Having shown that the probability of mislearning is monotonically increasing in ζ_+, we can take comparative statics:

$$\frac{\partial \zeta_+}{\partial q_d} = \frac{q_d}{2\sqrt{4q_s + q_d + 4}} + \frac{1}{2} \quad \text{and} \quad \frac{\partial \zeta_+}{\partial q_s} = \frac{1}{\sqrt{4q_s + q_d + 4}}.$$

It is easy to see that $\partial \zeta_+/\partial q_d > \partial \zeta_+/\partial q_s > 0$ for all $q_s \geq 0$ and $q_d > 0$. □

A.11 *Proof of Proposition 8*

PROOF. Define κ_q to be a naive agent i's log-likelihood ratio of state $\omega = 1$ versus state $\omega = 0$ upon observing one neighbor j who picks action 1 with weight q. Then we have

$$\kappa_q := \ln\left(\frac{\mathbb{P}[\omega = 1 \mid s_j \geq 0]}{\mathbb{P}[\omega = 0 \mid s_j < 0]}\right) > 0,$$

where $\mathbb{P}$ is taken under i's beliefs about the conditional distributions of s_j under Assumption 2, that is, $s_j \mid (\omega = 1) \sim \mathcal{N}(1, \sigma^2/q)$ and $s_j \mid (\omega = 0) \sim \mathcal{N}(-1, \sigma^2/q)$. In particular, this log likelihood κ_q is decreasing in q and so $\kappa_{q_s} - \kappa_{q_d} > 0$ for $q_s > q_d$.

By symmetry of the Gaussian distribution, the log-likelihood ratio after observing one neighbor who chooses action 0 with weight q is $-\kappa_q$.

Suppose after $2n$ agents have moved, the actions taken so far involve every odd-numbered agent playing 1 and every even-numbered agent playing 0. Then agent $2n + 1$ has a log-likelihood ratio of $n(\kappa_{q_s} - \kappa_{q_d})$ from her social observations. The probability that private signal s_{2n+1} is so strongly in favor of $\omega = 0$ as to make $2n + 1$ play 0 is

$$\epsilon_n := \mathbb{P}\left[s_i \in \mathbb{R} : \ln\left(\frac{\mathbb{P}[\omega = 1 \mid s_i]}{\mathbb{P}[\omega = 0 \mid s_i]}\right) < -n(\kappa_{q_s} - \kappa_{q_d}) \mid \omega = 1\right].$$

For the Gaussian distribution, $\ln(\mathbb{P}[\omega = 1 \mid s_i]/\mathbb{P}[\omega = 0 \mid s_i]) = 2s_i/\sigma^2$, so

$$\sum_{n=1}^{\infty} \epsilon_n = \sum_{n=1}^{\infty} \Phi\left(-\frac{\sigma^2}{2}n(\kappa_{q_s} - \kappa_{q_d}); 1, \sigma^2\right) < \infty$$

because the Gaussian distribution function tends to 0 faster than geometrically. This shows that there is a positive probability that every odd-numbered agent plays 1.

By an analogous argument, there is also a positive probability that every even-numbered agent plays 0. In that argument we would use the fact that

$$\sum_{n=1}^{\infty}\left[1 - \Phi\left(\frac{\sigma^2}{2}n(\kappa_{q_s} - \kappa_{q_d}); 1, \sigma^2\right)\right] < \infty.$$

□

References

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz (2016), "Fragility of asymptotic agreement under Bayesian learning." *Theoretical Economics*, 11, 187–225. [430]

Acemoglu, Daron, Giacomo Como, Fabio Fagnani, and Asuman Ozdaglar (2013), "Opinion fluctuations and disagreement in social networks." *Mathematics of Operations Research*, 38, 1–27. [430]

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), "Bayesian learning in social networks." *Review of Economic Studies*, 78, 1201–1236. [418, 430]

Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi (2010), "Spread of (mis) information in social networks." *Games and Economic Behavior*, 70, 194–227. [423]

Bala, Venkatesh and Sanjeev Goyal (1998), "Learning from neighbours." *Review of Economic Studies*, 65, 595–621. [417, 423]

Banerjee, Abhijit V. (1992), "A simple model of herd behavior." *Quarterly Journal of Economics*, 107, 797–817. [418, 429, 430]

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy*, 100, 992–1026. [418, 430]

Bohren, J. Aislinn (2016), "Informational herding with model misspecification." *Journal of Economic Theory*, 163, 222–247. [418]

Bohren, J. Aislinn and Daniel Hauser (2018), "Bounded rationality and learning: A framework and a robustness result." Unpublished paper, SSRN 3236842. [418, 430]

Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri (2020), "Testing models of social learning on networks: Evidence from two experiments." *Econometrica*, 88, 1–32. [416, 420]

Coleman, James (1958), "Relational analysis: The study of social organizations with survey methods." *Human Organization*, 17, 28–36. [429]

Dasaratha, Krishna and Kevin He (2019), "An experiment on network density and sequential learning." Unpublished paper. Available at arXiv:1909.02220. [417, 427]

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003), "Persuasion bias, social influence, and uni-dimensional opinions." *Quarterly Journal of Economics*, 118, 909–968. [417]

Enke, Benjamin and Florian Zimmermann (2017), "Correlation neglect in belief formation." *Review of Economic Studies*, 86, 313–332. [420]

Eyster, Erik and Matthew Rabin (2010), "Naive herding in rich-information settings." *American Economic Journal: Microeconomics*, 2, 221–243. [416, 418, 420, 427]

Eyster, Erik and Matthew Rabin (2014), "Extensive imitation is irrational and harmful." *Quarterly Journal of Economics*, 129, 1861–1898. [418]

Eyster, Erik, Matthew Rabin, and Georg Weizsacker (2015), "An experiment on social mislearning." Unpublished paper, SSRN 2704746. [416, 420]

Golub, Benjamin and Matthew O. Jackson (2010), "Naive learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2, 112–149. [417, 423]

Golub, Benjamin and Matthew O. Jackson (2012), "How homophily affects the speed of learning and best-response dynamics." *Quarterly Journal of Economics*, 127, 1287–1338. [417, 429]

Hązła, Jan, Ali Jadbabaie, Elchanan Mossel, and M. Amin Rahimian (forthcoming), "Bayesian decision making in groups is hard." *Operations Research*. Available at arXiv:1705.04770. [420]

Lee, In Ho (1993), "On the convergence of informational cascades." *Journal of Economic Theory*, 61, 395–411. [430]

Levy, Gilat and Ronny Razin (2018), "Information diffusion in networks with the Bayesian peer influence heuristic." *Games and Economic Behavior*, 109, 262–270. [418, 422]

Lobel, Ilan and Evan Sadler (2015), "Information diffusion in networks through social learning." *Theoretical Economics*, 10, 807–851. [418]

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook (2001), "Birds of a feather: Homophily in social networks." *Annual Review of Sociology*, 27, 415–444. [417]

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018), "A theory of non-Bayesian social learning." *Econometrica*, 86, 445–490. [418]

Mossel, Elchanan, Allan Sly, and Omer Tamuz (2015), "Strategic learning and the topology of social networks." *Econometrica*, 83, 1755–1794. [423]

Mueller-Frank, Manuel and Claudia Neri (2015), "A general model of boundedly rational observational learning: Theory and evidence." Unpublished paper, SSRN 2566210. [420]

Mueller-Frank, Manuel and Claudia Neri (2019), "A general analysis of boundedly rational learning in social networks." Unpublished paper, SSRN 2933411. [418]

Sethi, Rajiv and Muhamet Yildiz (2012), "Public disagreement." *American Economic Journal: Microeconomics*, 4, 57–95. [430]

Smith, Lones and Peter Sørensen (2000), "Pathological outcomes of observational learning." *Econometrica*, 68, 371–398. [430]

Weizsäcker, Georg (2010), "Do we follow others when we should? A simple test of rational expectations." *American Economic Review*, 100, 2340–2360. [421]