Network Structure and Naive Sequential Learning*

Krishna Dasaratha† Kevin He‡

First version: January 17, 2017
This version: August 3, 2018

Abstract

We study a sequential learning model featuring naive agents on a network. Agents wrongly believe their predecessors act solely on private information, so they neglect redundancies among observed actions. We provide a simple linear formula expressing agents' actions in terms of network paths and use this formula to completely characterize the set of networks guaranteeing eventual correct learning. This characterization shows that on almost all networks, disproportionately influential early agents can cause herding on incorrect actions. Going beyond existing social-learning results, we compute the probability of such mislearning exactly. This lets us 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. We conduct an experiment and find that the accuracy gain from social learning is twice as large on sparser networks, which is consistent with naive inference but inconsistent with the rational-learning model.

*We thank Daron Acemoglu, J. Aislinn Bohren, Jetlir Duraj, Ben Enke, Erik Eyster, Drew Fudenberg, Ben Golub, David Laibson, Jonathan Libgober, Margaret Meyer, Pooya Molavi, Xiaosheng Mu, Matthew Rabin, Tomasz Strzalecki, Omer Tamuz, Linh T. Tô, and Muhamet Yildiz for useful comments. Financial support from the Eric M. Mindich Research Fund for the Foundations of Human Behavior is gratefully acknowledged.

†Department of Economics, Harvard University. Email: krishnadasaratha@gmail.com
‡Department of Economics, Harvard University. Email: hesichao@gmail.com

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 her decision. For instance, when consumers choose between rival products, their decision is 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. And when a new theory or rumor is introduced into a society, individuals are influenced by the discussions of those who have already taken positions on the new idea.

A key feature of these examples is that agents only observe 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 of her social circle. In general, each sequential social learning problem has an *observation network* that determines which predecessors are observable to each agent. The observation networks associated with different learning problems can vary in density, extent of segregation, and other structural properties. Our main question is: how does the structure of the observation network affect the probability of correct social learning in the long-run?

To answer this question, one must first realistically capture key aspects of how individuals process social information in such learning settings. Empirical research on social learning, including the field experiment of Chandrasekhar, Larreguy, and Xandri (2015) and the laboratory evidence of Eyster, Rabin, and Weizsacker (2015), suggests that humans often exhibit *inferential naiveté*, failing to understand that their predecessors' actions reflect a combination of their private information and the inference those predecessors have drawn from the behavior of still others. 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 the viral story's proponents are just following the herd and possess no private information as to the rumor's veracity.

The present work studies 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 the most tractable model of such naiveté (as in Eyster and Rabin, 2010), which lets us exactly compute naive agents' probability of taking the correct action ("learning accuracy") on arbitrary networks. In contrast, the existing literature on social learning has focused on the binary classification of learning situations into those where correct learning occurs with probability one and those where learning is less than perfect with positive probability. We obtain a richer characterization of imperfect learning, so we can compute comparative statics of learning accuracy with respect to the network structure. We then test these comparative statics in an experiment, which gives a new form of evidence for inferential naiveté.

Our analysis focuses on comparisons among the network structures that imply imperfect learning. We show this is the leading case: the only networks leading to asymptotically correct learning for naive agents are those where every agent's "influence" on group opinion is negligible, a condition violated by almost all common models of social networks. Outside of this class, we show early agents taking wrong actions can cause herding on the same incorrect action and express the exact probability that society ends up fully convinced of the wrong state of the world ("mislearning") in terms of the network structure. This allows us to study how changes in network parameters influence agents' welfare. Since the set of networks that preclude naive agents from always learning asymptotically is very broad, the detailed comparative statics we can offer are crucial: imperfect learning may imply a wide range of welfare losses. In addition, the prediction that naive learning should be differentially inefficient on various networks motivates our empirical test of inferential naiveté, which experimentally compares learning accuracy on sparse and dense networks.

Introducing naiveté reverses several predictions of rational social-learning models which are at odds with empirical observations. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011)'s Bayesian learning model predicts that as long as agents can observe recent predecessors and private signals can have unbounded informativeness, agents' beliefs about the state of the world will eventually settle down on a common consensus and their actions will converge to the objectively optimal one. This prediction holds for Bayesian agents regardless of observation network density or the extent of network segregation. By contrast, we prove that increasing network density leads to worse social-learning outcomes. Our experiment reported in Section 6 confirms this prediction: the gain in accuracy from social learning is twice as large on sparse networks compared to dense networks. This finding is consistent with redundancy neglect, as denser networks introduce more opportunities for naive agents to fail to account for redundancies. As another example, disagreement among different communities is common in practice. In the product-adoption domain, different subcommunities frequently insist upon the superiority of their preferred products. With partially segregated observation networks, we prove disagreement between different subgroups can persist even when there are unboundedly many connections across subgroups.[1] This presents a sharp contrast with Bayesian social-learning models (and leading naive-learning models such as DeGroot), where asymptotic agreement is a robust prediction.

Finally, studying the impact of the network structure on learning outcomes can help us understand major societal changes. As social and technological changes restructure our networks, social learning dynamics change as well. Most notably, the rise of the Internet has dramatically altered the transmission of information about current events, new products, etc. By clarifying how network structures affect learning outcomes for a society of naifs, we aim to provide a framework for understanding such changes.

[1] Segregation is extensively documented in social networks (see McPherson, Smith-Lovin, and Cook (2001) for a survey).

1.1 Preview of Results

Actions in our model take a simple form: assuming a continuous action space and a binary state space, we can express each agent's action as a linear function of her predecessors' private signal realizations with coefficients depending on the network structure. We exploit this expression to develop a necessary and sufficient condition for a society of naifs to learn correctly: no agent has too much "influence." We identify the relevant notion of influence in a sequential learning setting with naive agents, which turns out to be related to the number of network paths terminating at each agent. This differs from other concepts of influence identified so far in the literature because in sequential learning settings early movers are observed by many successors, an asymmetry which is absent in other social-learning models.

We then derive a formula for the exact probability of asymptotic social mislearning as a function of the network structure. We apply this formula to a number of commonly used network models and find two new effects distinguishing naive sequential learning from rational and DeGroot learning. In a model with uniform link strength between agents, society is more likely to mislearn when there are more links. When society is composed of two homophilic (i.e. segregated) subgroups, mislearning is more likely when segregation between the subgroups decreases. A theme is that changes in network structure that would increase the speed of learning in rational and DeGroot models tend to make mislearning more likely in our model. This is because faster convergence of beliefs creates more (neglected) redundancy under inferential naiveté, whereas slower convergence of beliefs tends to imply more independent sources of information get incorporated into the social consensus.

Next, we show that if agents' actions only coarsely reflect their beliefs and society is partially segregated, then two homophilic social subgroups can disagree forever. Because of agents' naiveté and the limited information conveyed by actions, disagreement can persist even when agents observe the actions of unboundedly many individuals from the other group. We thus contribute to an active literature on disagreement (Acemoglu, Como, Fagnani, and Ozdaglar, 2013; Acemoglu, Chernozhukov, and Yildiz, 2016; Sethi and Yildiz, 2012) by demonstrating that network homophily can generate divergence in beliefs.

Finally, we conduct an experiment to distinguish the naive-learning model from the rational-learning model in terms of their comparative statics predictions. We place subjects in sequential networks of 40 agents, where each agent has a 25% chance of observing each predecessor in the sparse treatment and a 75% chance in the dense treatment. The subjects know the network-generating process and must guess the binary state using their private signal and their social observations about the guesses made by their predecessors, with incentives for accuracy. Prior to data collection, we pre-registered a measure of the long-run accuracy for each network: the fraction of the final 8 agents who correctly guess the state. Comparing this pre-registered measure on 130 sparse networks versus 130 dense networks, we find denser networks lead to worse learning accuracy. In denser networks, the last fifth of the agents guess correctly 5.7% more frequently than agents who have no social observations, but this improvement is 12.6% in sparse networks. Thus, we find the accuracy gains from social learning for these agents are twice as large in the sparse treatment as in the dense treatment (p-value 0.0239). This is consistent with models of inferential naiveté, including the functional form developed in Section 2 and other variants featuring a mixture of naive agents and autarkic agents (who ignore all social information). On the other hand, we show the doubling of accuracy gain is inconsistent with the rational model (for agents 33 through 40 in finite networks matching our experimental setup). We also find network density has no statistically significant effect on overall guess accuracy averaged across all forty agents, because early agents are more accurate on the dense network—as predicted by inferential naiveté.

The rest of the paper is organized as follows: the remainder of Section 1 reviews related literature. Section 2 introduces the model of naive sequential learning on a network. Section 3 contains a path-counting interpretation of actions and a necessary and sufficient condition for correct social learning based on a suitable notion of "influence" in the network. Section 4 begins by deriving a formula for the probability of mislearning on a general network, then applies this expression to a number of canonical network models and obtains comparative statics on the probability of mislearning with respect to network parameters. Section 5 contains results about disagreement in homophilic societies when actions only coarsely reflect beliefs. We present our experimental design and evidence in Section 6. Section 7 concludes.

1.2 Related Literature

1.2.1 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, Vayanos, and Zwiebel, 2003). Moreover, this prediction will be correct in large networks as long as the network is not too unbalanced (Golub and Jackson, 2010). So the asymptotic learning outcomes under DeGroot and rational learning are qualititatively similar: DeMarzo, Vayanos, and Zwiebel (2003) show that in this setting, rational agents learn the true state of the world in finite time. In our sequential learning model, by contrast, naiveté implies a wider scope for incorrect learning and often even disagreement, because early agents exert undue influence on the eventual social consensus. So, it is easy to distinguish naive and Bayesian learning.

Because the DeGroot heuristic is linear, how the network structure affects learning is well understood. Much of this analysis focuses on how quickly beliefs converge to the consensus, and Golub and Jackson (2012) show that the speed of convergence is determined by the degree of homophily in the network. We find that in a sequential setting, natural changes in network structure matter for asymptotic beliefs, not only for the speed of learning. Changing network density, which has no effect on DeGroot learning in large networks, can substantially alter the probability society learns correctly. Homophily now matters for this probability and even for whether consensus is ever reached.

While DeGroot proposes the averaging rule as an ad-hoc heuristic, several recent papers have developed behavioral microfoundations for naive learning in the repeated-interaction setting. Molavi, Tahbaz-Salehi, and Jadbabaie (2018) axiomatize "log-linear learning", where the log-odds ratio of each player's belief in period $t+1$ equals to a linear combination of the log-odds ratios of her neighbors' beliefs from period t. This characterization resembles our Lemma 1 and Lemma 2, which show that log-linear learning also appears in a sequential learning model under our inferential naiveté assumption. Mueller-Frank and Neri (2015) discuss conditions on utility functions and signal structures guaranteeing that naive agents will learn correctly on *all* networks. Levy and Razin (2018) define the Bayesian Peer Influence heuristic, which is based on the same error as the behavioral assumption we study. As in DeGroot, their model leads to correct learning under reasonable assumptions on the network structures. 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.

1.2.2 Sequential Social Learning

We consider the same environment as the extensive literature on sequential social learning. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015) characterize network features that lead to correct asymptotic learning for Bayesians who move sequentially. We consider instead a version of the problem where the sequential learners suffer from naive inference. By providing a thorough understanding of rational learning in sequential settings, this literature provides a valuable benchmark. We find that among network structures where Bayesian agents learn asymptotically, there is large variation in the probability of mislearning and thus the extent of welfare loss for naive agents.

Several authors look at sequential naive learning on a particular network structure, usually the complete network (Eyster and Rabin, 2010; Bohren, 2016; Bohren and Hauser, 2017). 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 naive learning with network observation, 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 Eyster and Rabin (2014)'s binary classification. This lets us develop richer understanding of how the network matters for the extent of mislearning.

1.2.3 Experimental Evidence for Naive Inference

Our experimental results add to a growing body of evidence that humans do not properly account for correlations in social-learning settings. In a field experiment where agents interact repeatedly with the same set of neighbors, Chandrasekhar, Larreguy, and Xandri (2015) find agents fail to account for redundancies. Enke and Zimmermann (2018) show that correlation neglect is prevalent even in simple environments where the observed information sources are mechanically correlated.

Most closely related to the present work, the laboratory games in Eyster, Rabin, and Weizsacker (2015) and Mueller-Frank and Neri (2015) directly evaluate behavioral assumptions matching ours. Eyster, Rabin, and Weizsacker (2015) find that on the complete network agents' behavior is closer to the rational model than the naive model. On a more complex network the naive model matches more observations than the rational model, and there is little anti-imitation (which would be required for correct Bayesian inference).[2] Mueller-Frank and Neri (2015) find most observations are consistent with the behavioral assumption we study (which they call quasi-Bayesian updating) in a setting where agents have limited information about the network. These experiments suggest naiveté may be more likely in settings where agents either have a limited knowledge of the network or the network is known but very complicated. In these settings, the correct Bayesian belief given one's observations can be far from obvious, so agents are more likely to resort to behavioral heuristics.

Unlike this previous work, our experiment tests the comparative statics predictions of naive and rational learning. This allows us to cleanly test redundancy neglect against rational updating. Our approach allows us to focus on long-term learning outcomes—which are the welfare-relevant outcome as we vary the learning environment—instead of solely on measuring individual behavior.

2 Model

2.1 Sequential Social Learning on a 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 from signal space S, as well as the actions of some previous agents. Then, i chooses an action $a_i \in [0, 1]$ to maximize the expectation of

$$u_i(a_i, \omega) \coloneqq -(a_i - \omega)^2$$

given her belief about ω, so her chosen action corresponds to the probability she assigns to the event $\{\omega = 1\}$.

Private signals (s_i) are i.i.d. conditional on the state of the world. To simplify notation, we assume that every $s \in S$ satisfies $0 < \mathbb{P}[\omega = 1 | s_i = s] < 1$, so that no

[2] In the complex network, four agents move in each period after observing predecessors from previous periods.

signal is fully revealing about the state of the world. To avoid trivialities, assume $\mathbb{E}[\mathbb{P}[\omega = 1|s_i = s]|\omega = 1] > \frac{1}{2}$, so that the signal structure is not fully uninformative.

In anticipation of the naive inference assumption we will make, we will now assume that the signal structure has unbounded informativeness:

Assumption 1 (Unbounded Signal Informativeness)**.** *For every $p \in (0, 1)$, there exists some $s \in S$ such that $\mathbb{P}[\omega = 1|s_i = s] = p$.*

Assumption 1 implies private beliefs have unbounded support: the distribution of posterior belief in state $\{\omega = 1\}$ after observing only one's own private signal has support $[0, 1]$.

Our initial results in Sections 2 and 3 will not require any further assumptions on the signal structure. For our results in Section 4, we will assume signals are Gaussian and symmetric across the two states:

Assumption 2 (Gaussian Signals)**.** *When $\omega = 1$, we have $s_i \sim \mathcal{N}(1, \sigma^2)$. When $\omega = 0$, we have $s_i \sim \mathcal{N}(-1, \sigma^2)$.*

Write $N_i \subseteq \{1, 2, ..., i-1\}$ for the subset of i's predecessors whose actions i can observe. Members of N_i are called *neighbors* of i. The sets (N_i) define a directed network. The adjacency matrix M of the network is defined by $M_{i,j} = 1$ if $j \in N_i$ and $M_{i,j} = 0$ otherwise. Because of the sequential observation structure, M is lower-triangular with diagonal entries all equal to zero.

If $M_{i,j} = 1$ whenever $j < i$, then M is associated with the complete network (up to the constraint that earlier agents cannot observe the behavior of later agents). In that case, we recover the commonly used social-learning model where every agent observes the actions of all predecessors.

2.2 Naive Inference Assumption

A growing body of recent evidence in psychology and economics shows that agents learning from peers are often not fully correct in their treatment of social structure.

Instead of calculating the optimal Bayesian behavior that fully takes into account all that they know about the network and signal structure, agents often apply heuristic simplifications to their environment and then act according to the solution of this simpler problem. When networks are complicated and/or uncertain, determining Bayesian behavior can be very difficult and these naive learning rules become especially prevalent. Motivated by this observation, we make the following behavioral assumption, which will be maintained throughout.

Assumption 3 (Naive Inference Assumption). *Each agent wrongly believes that each predecessor chooses an action to maximize her expected payoff based on only her private signal, and not on her observation of other agents.*

The naive inference assumption can be equivalently described as misperceiving $N_j = \emptyset$ for all $j < i$. Under this interpretation, agents act as if their neighbors do not observe anyone.

Besides the error in Assumption 3, agents are otherwise correctly specified and Bayesian. They correctly optimize their expected utility given their mistaken beliefs.[3]

Remark 1. The behavior given by Assumption 3 is cognitively simple in that it does not rely on agent's knowledge about the network or even knowledge about the order in which the predecessors moved. So our model applies even to complex environments with random arrival of agents and/or random networks. In such environments, the naiveté assumption may be more realistic than assuming full knowledge about the observation structure and the move order.

In the literature, Assumption 3 was first considered in a sequential learning setting by Eyster and Rabin (2010), who coined the term "best-response trailing naive inference" (BRTNI) to describe this behavior. They find that even in an "information-rich" setting where private signals have full support and actions perfectly reveal be-

[3]Thanks to Unbounded Signal Informativeness (Assumption 1), every action that a predecessor might take can be rationalized using only their private signal. This means a naive agent acting according to Assumption 3 will never be "surprised" by a predecessor's action that she deems impossible. The social-learning model with Assumption 3 is therefore well-defined.

liefs, BRTNI agents on a complete network have a positive probability of converging to playing the wrong action.

Lemma 1 characterizes the behavior of a naive agent who sees private signal s_i and the profile of neighbor actions $(a_j)_{j \in N_i}$. To state the lemma and our subsequent results, it is convenient to first perform the following change of variables:

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

In words, $\tilde{s}_i$ is the log-likelihood ratio of the events $\{\omega = 1\}$ and $\{\omega = 0\}$ given signal s_i, while $\tilde{a}_i$ 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.[4]

Lemma 1. *A naive agent who observes private signal s_i and the profile of neighbor actions $(a_j)_{j \in N_i}$ plays a_i such that $\tilde{a}_i = \tilde{s}_i + \sum_{j \in N_i} \tilde{a}_j$.*

The key step is to see that, due to naiveté, agent i who observes action a_j wrongly infers that j's private log signal must satisfy $\tilde{s}_j = \tilde{a}_j$. (This inference is possible since the continuum action set is rich enough to exactly reveal beliefs of predecessors.) This means i's belief about the state of the world can be expressed simply in terms of the $(\tilde{a}_j)_{j \in N_i}$ that she observes. The action a_i is the product of the relevant likelihoods because a naive agent assumes observed actions are independent, and therefore $\tilde{a}_i$ is the sum of the corresponding log-likelihood ratios.

This lemma characterizes agent behavior by expressing actions in terms of observed information. Moreover, the relevant formula is simple and linear. In Proposition 1, we will use Lemma 1 to express actions in terms of only signals and the network structure.

We assume that agents are naive with respect to social structure, but are able to correctly interpret private signals. This partial sophistication is not central to our analysis: it is crucial that agents treat their social observations as independent and

[4]Note 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.

exchangeable, but misspecified beliefs about signal structure would not substantially alter our findings.

2.3 Weighted Network Edges

In Section 4, we will derive comparative statics with respect to the network structure. Because studying continuous changes in the network is more tractable than discrete changes, we extend our network model to allow non-integer weights on network edges.

Henceforth our adjacency matrix entries can be real numbers $M_{i,j}$, which we usually take to lie in the interval $[0, 1]$. Below we give two interpretations of non-integer weights implying that naive agents act according to the formula

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

which is a version of Lemma 1 where observed neighbor actions $(\tilde{a}_j)_{j \in N_i}$ are assigned weights proportional to the corresponding link weights $(M_{i,j})_{j \in N_i}$. The underlying idea is that agents place more trust in or more emphasis on neighbors with whom their connections are stronger.

Noisy Observations Interpretation: This interpretation applies when signals are Gaussian as in Assumption 2. Suppose that instead of observing $\tilde{a}_j$, agent i observes $\tilde{a}_j + \epsilon_{i,j}$, where $\epsilon_{i,j} \sim \mathcal{N}\left(0, \frac{4}{\sigma^2}\left(\frac{1}{M_{i,j}} - 1\right)\right)$ is a Gaussian random variable with mean zero and variance decreasing in $M_{i,j}$, the weight of the network edge from i to j. Note that $M_{i,j} = 0$ corresponds to i observing a signal about $\tilde{a}_j$ with "infinite variance", i.e. a purely uninformative signal. On the other hand, $M_{i,j} = 1$ means i observes $\tilde{a}_j$ perfectly. We then consider actions in the case where observations are correct (or equivalently, actions of agents who mistakenly perceive observation as noisy).

Generations Interpretation: Replace each agent i with a continuum of agents $[0, 1]$, interpreted as the i^{th} generation. Each member of generation i has $M_{i,j}$ chance of observing a uniformly random member of generation j, where the randomization over observation or no observation is independent across members of generation i. Everyone in generation i receives the same private signal s_i. This setup induces a distribution over actions in each generation, and we re-interpret $\tilde{a}_i$ as the average log-action among generation i.[5]

By arguments analogous to Lemma 1, it is easy to show that:

Lemma 2. *Under either the Noisy Observations Interpretation or the Generations Interpretation,* $\tilde{a}_i = \tilde{s}_i + \sum_{j \in N_i} M_{i,j} \tilde{a}_j$.

Intuitively, a weighted network with all link weights in the interval $[0, 1]$ is a deterministic version of the random network model where link (i, j) is formed with probability $M_{i,j}$.[6] We prove our comparative statics results for weighted networks as they are analytically more tractable than random graphs.[7]

Remark 2. One distinction between our model and DeGroot is that we allow for agents to have any out-degree, while the DeGroot heuristic requires all agents' weights sum to one. When edges are not weighted, so that all $M_{i,j} \in \{0, 1\}$, any agent with multiple observations has out-degree greater than one. This distinction is not just a normalization, but is in fact the source of redundancy and naive inference.[8]

[5]A model where each generation contains a continuum of agents is central to Gagnon-Bartsch and Rabin (2016)'s study of naive social learning. We stipulate that generation i players make i.i.d. observations of generation j's behavior, just as members of the same generation receive i.i.d. private signals in Gagnon-Bartsch and Rabin (2016).

[6]In simulations, all comparative statics results proved for weighted networks in Section 4 continue to hold in the analogous random network models.

[7]The major obstacle to extending proofs is that because our networks are directed and acyclic, the relevant adjacency matrices have no non-zero eigenvalues. As a consequence, most techniques from spectral random graph theory do not apply.

[8]But we can embed DeGroot learning with k agents in our model (with a slight modification so that only the first k agents receive signals). Consider our sequential model with agents grouped into blocks of k, where agents $tk+1, tk+2, ..., (t+1)k$ correspond to the period t actions of the k agents in the DeGroot model. Agents $1, ..., k$ receive private signals and have no observations. For $t \geq 1$, agents $tk+1, tk+2, ..., (t+1)k$ do not have private signals and are connected only to the agents in the previous block, with edge weights equal to the DeGroot weights. Under this setup, the action of agent $tk+j$ in the sequential model is equal to the period t action of agent j in the DeGroot model.

3 Social Influence and Learning

3.1 Path-Counting Interpretation of Actions

If players were Bayesians, then the vector of their actions (a_i) would be a complicated function of their signal realizations, (s_i). We now show that with naive agents, actions actually have a simple (log-)linear expression in terms of paths in the network. This result requires no further assumptions about the signal structure.

Below we abuse notation by using M to also refer to the $n \times n$ upper-left submatrix of M.

Proposition 1. *For each n, the actions of the first n naive agents are determined by*

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

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.[9]

Proposition 1 expresses agent behavior in terms of only the network structure and signal realizations. Actions are a linear combination of signals, where the coefficients have a simple formula in terms of the adjacency matrix.

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. "Weighted paths" means the path passing through agents $i_0, ..., i_K$ 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

[9] Recall that M is lower-triangular due to sequential learning imposing the restriction that agent i does not observe any agent j for $j > i$. Therefore, the set of *paths* between two agents is the same as the set of *walks* between them.

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 Correct Learning

We will now use the representation result of Proposition 1 to study which networks lead to mislearning by naive agents. We first define what it means for society to learn correctly in terms of convergence of actions.

Definition 1. Society *learns correctly in the long-run* if (a_n) converges to ω in probability.

In a setting where society learns correctly in the long-run, agent n becomes very likely to believe strongly in the true state of the world as n grows large.

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

Definition 2. Let $b_{i,j} \coloneqq (I - M)^{-1}_{ij}$ be the number of weighted paths from i to j in network M.

Without network weights, $b_{i,j}$ is equal to the number of paths from i to j. With network weights, each path is counted with multiplicity equal to the product of the weights of the edges included in the path. Because of Proposition 1, these path counts are important to our analysis.

Definition 3. For $n > i$, the *influence* of i on n is $\mathbb{I}(i \to n) \coloneqq b_{n,i} / \sum_{j=1}^{n} b_{n,j}$.

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

An alternate definition of influence appears in Golub and Jackson (2010), who study naive social 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 updating matrix, which is proportional to i's degree in an undirected network with symmetric weights.[10] 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.

We will make the following connectedness assumption on M.

Assumption 4 (Connectedness Assumption)**.** *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 indirectly observe some early agent. When all links have weight 0 or 1, this assumption holds if and only if there are only finitely many agents who have no neighbors. If such a network fails the connectedness assumption, then clearly the infinitely many agents without neighbors will prevent society from learning correctly.

Proposition 2. *Suppose the connectedness assumption holds and $\tilde{s}_i$ has finite variance. Society learns correctly in the long-run if and only if $\lim_{n\to\infty} \mathbb{I}(i \to n) = 0$ 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).[11] The main contribution of Proposition 2 is to identify the relevant measure of influence in a sequential learning setting with naive agents. In our sequential setting, unlike on large unordered networks like those in Golub and Jackson (2010), the ordering of agents creates an asymmetry that prevents society from learning correctly on most natural networks.

[10] When we embed the DeGroot model in our sequential model as in Footnote 8, the eigenvector-based definition of influence for the i^{th} DeGroot agent is equal to the limit $\lim_{n\to\infty} \mathbb{I}(i \to n)$ of i's influence under our definition.

[11] As discussed above, our condition for learning correctly has a similar structure to the condition in Golub and Jackson (2010). The idea that a few influential agents can drive mislearning is also central to Acemoglu, Ozdaglar, and ParandehGheibi (2010). There, influence depends on whether agents are "forceful" and not only on their network position.

The idea behind the proof is that if there were some i and $\epsilon > 0$ such that $\mathbb{I}(i \to n) > \epsilon$ 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 so 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 correctly. 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 correctly. 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—but this is guaranteed by the condition on influence $\lim_{n\to\infty} \mathbb{I}(i \to n) = 0$.

4 Probability of Mislearning and Network Structure

In this section, we compare the probability of mislearning across networks where learning is imperfect. To do so, we first derive a formula for the probability of asymptotic mislearning as a function of the observation network. Proposition 2 implies this probability will be non-zero on many network structures of interest. 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

In Section 4, we assume signals are Gaussian and symmetric across the two states as in Assumption 2 in order to get tractable closed-form results.[12]

With Gaussian signals, we can give explicit expressions for the distributions of agent actions in each period. To this end, we first state a convenient fact about models with binary states and Gaussian signals: the associated log-likelihood ratios also happen to have a Gaussian distribution:

Lemma 3. *Under Assumption 2,* $\tilde{s}_i = 2s_i/\sigma^2$*.*

We now show 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, $\vec{b}_n := (b_{n,1}, ..., b_{n,n})$. The ratio $\frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}$ can be viewed as a measure of distributional equality for the vector of weights $\vec{b}_n$.[13] 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 $\vec{b}_n = (1, 0, ..., 0)$ and maximized by the vector $\vec{b}_n = (\frac{1}{n}, \frac{1}{n}, ..., \frac{1}{n})$.

The next lemma gives the ex-ante distribution of agent n's log-action.

Lemma 4. *When* $\omega = 1$*, the log-action of agent* n *is distributed as* $\tilde{a}_n \sim \mathcal{N}\left(\frac{2}{\sigma^2}\|\vec{b}_n\|_1, \frac{4}{\sigma^2}\|\vec{b}_n\|_2^2\right)$.

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 Lemma 4 first expresses $\tilde{a}_n = \sum_{i=1}^{n} b_{n,i}\tilde{s}_i$ using Proposition 1, then observes (s_i) are distributed i.i.d. $\mathcal{N}(1, \sigma^2)$ conditional on $\omega = 1$. This means $(\tilde{s}_i)$ are

[12] From simulations, the comparative statics results also hold under the triangular signal structure of Eyster and Rabin (2010).

[13] 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.

also conditionally i.i.d. Gaussian random variables, due to Lemma 3. As a sum of conditionally i.i.d. Gaussian random variables, the action $\tilde{a}_n$ is itself Gaussian. The result follows from calculating the mean and variance of this sum.

As a consequence, we can express in terms of the network structure the ex-ante probability that agent n puts more confidence in the state being $\omega = 0$ when in fact $\omega = 1$. This gives the key result that lets us compare the extent of imperfect learning on different networks.

Theorem 1. *The probability that agent n thinks the incorrect state is more likely than the correct one is* $\mathbb{P}[\tilde{a}_n < 0 | \omega = 1] = \Phi\left(-\frac{1}{\sigma} \cdot \frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}\right)$.

For the remainder of this section, we will study specific weighted networks where the ratio $\frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}$ 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 Lemma 4 and Theorem 1. This will allow us to relate network parameters to learning outcomes.

For the remainder of this section, let the signal variance $\sigma^2 = 1$. Note that in the limit as signals become uninformative ($\sigma^2 \to \infty$), the probability of correct learning converges to $\frac{1}{2}$.

4.2 Uniform Weights

The simplest network we consider assigns the same weight $q \in [0, 1]$ to each feasible link. This network is analogous to an Erdős–Rényi random network, where each feasible link is present with a fixed probability q. By varying the value of q, we can ask how link density affects the probability of mislearning, which we now define.

Definition 4. Society *mislearns* when $\lim_{n\to\infty} a_n = 0$ but $\omega = 1$, or when $\lim_{n\to\infty} a_n = 1$ but $\omega = 0$.

Society mislearns when all late enough agents are very confident but incorrect.

Proposition 3. *When $q \in (0, 1]$, agents' actions a_n converge to 0 or 1 almost surely. The probability that society mislearns is $\Phi\left(-\sqrt{\frac{q+2}{q}}\right)$ where Φ is the standard Gaussian distribution function. 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 non-zero for all positive q, and increases in q.[14] When the observational network is more densely connected, society is more likely to be wrong, as in Figure 1.

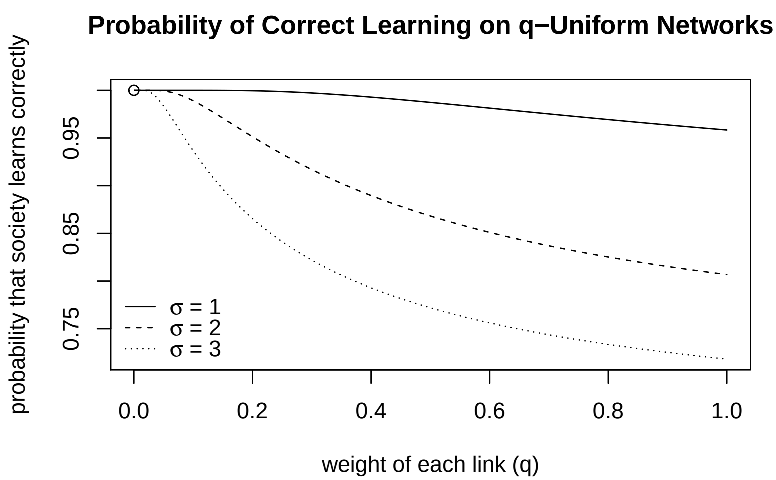

Figure 1: Probability that society learns correctly ($a_n \to \omega$) on a network where all links have weight q. (Proposition 3 shows this probability is $1 - \Phi\left(-\sqrt{\frac{q+2}{q}}\right)$ under the normalization $\sigma = 1$, but one can similarly show this probability is $1 - \Phi\left(-\frac{1}{\sigma} \cdot \sqrt{\frac{q+2}{q}}\right)$ more generally.) See Remark 4 for the discontinuity at $q = 0$.

When the observation network is sparse (i.e. q is low), early agents' actions convey

[14] We can extend the uniform weights network to incorporate recency effect: the weight of the link between i and j is δ^{i-j} for some $0 < \delta < 1$, so that more recent predecessors get more weight. Intuitively, this models a situation where agents are better connected to or place more trust in peers who are closer to them in the network. The conclusion that denser network causes more mislearning extends naturally with similar proof: when $\delta > \frac{1}{2}$, the probability of mislearning is increasing in δ. (When $\delta < \frac{1}{2}$, the actions (a_n) almost surely do not converge. Society learns correctly in the long-run on the boundary $\delta = \frac{1}{2}$.)

a large amount of independent information because they do not influence each other 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.[15] A related intuition compares agents' beliefs about network structure to the actual network: as q grows, agents' beliefs about the network structure differ more and more from the true network. To complement this theoretical result, in Section 6 we conduct a sequential learning experiment to evaluate the comparative static. In line with the intuition above, we find that human subjects indeed do worse 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 this recurrence holds, let Ψ_n be the set of all paths from n to i and Ψ_{n-1} be the set of all paths from $n-1$ to i. For each $\psi \in \Psi_{n-1}$, $\psi = (n-1) \to j_1 \to j_2 \to ... \to i$, we associate two paths $\psi', \psi'' \in \Psi_n$, with $\psi' = n \to j_1 \to j_2 \to ... \to i$ and $\psi'' = n \to (n-1) \to j_1 \to j_2 \to ... \to i$. This association exhaustively enumerates all paths in Ψ_n as we consider all $\psi \in \Psi_{n-1}$. 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 one. This shows the weight of all paths in Ψ_n is equal to $1+q$ times the weight of all paths in Ψ_{n-1}, hence $b_{n,i} = (1+q)b_{n-1,i}$.

Remark 3. 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 non-zero 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 $q \in (0, 1]$, under a Gaussian signal structure.

Remark 4. There is a discontinuity at $q = 0$. As q approaches 0, the probability of

[15] Sadler (2017) discusses a related intuition in the context of experimentation with social learning. In denser networks there is less experimentation, and thus agents are more likely to be wrong about the state.

society eventually learning correctly approaches 1.[16] But when $q = 0$, each agent only observes her own private signal. In this case there is no social learning since all agents choose actions equal to their signals. This non-convergence of actions means that society never learns correctly.

4.3 Two Groups

We next consider a network with two groups and different weights for links within groups and between groups. This is a deterministic version of a simple stochastic block model, which will be discussed further in Section 5. By varying the link weights, we will 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 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-groups network.

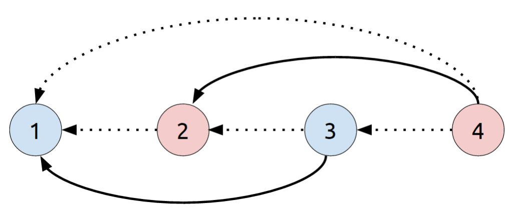

Figure 2: First four agents in a two-groups network. Odd-numbered agents are in one group (blue) while even-numbered agents are in another group (red). Solid arrows have weight q_s and dashed arrows have weight q_d.

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

Proposition 4. *When $q_s \in [0, 1]$ and $q_d \in (0, 1]$, agents' actions a_n converge almost surely to 0 or 1. The partial derivatives of the mislearning probability $\xi(q_s, q_d)$ satisfy*

[16] Although the speed of learning slows.

$\frac{\partial \xi}{\partial q_d} > \frac{\partial \xi}{\partial q_s} > 0$, *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 $\frac{\partial \xi}{\partial q_d} > \frac{\partial \xi}{\partial q_s}$ tells us that more integrated (i.e. less homophilic) networks are more likely to herd on the wrong state of the world.[17]

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. (In Section 5, however, we will see 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. Therefore, agent n's belief does not depend substantially on whether n is in the odd group or the even group.[18]

Coleman (1958)'s homophily index equals to $\frac{q_s - q_d}{q_s + q_d}$ 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-

[17]We can also extend the two-groups network to model recency effect: for some $0 < \delta < 1$, the weight of the link between i and j (with $i > j$) is $\delta^{i-j} q_s$ when i and j are in the same group and $\delta^{i-j} q_d$ otherwise. The conclusions of Proposition 4 remain robust under this extension with a similar proof. Namely, there exists a threshold $\delta_0(q_s, q_d)$ such that when $\delta > \delta_0(q_s, q_d)$, the probability of mislearning is increasing in each of δ, q_s, q_d, but increasing q_d has a larger effect than increasing q_s. If $\delta < \delta_0(q_s, q_d)$, then almost surely the actions (a_n) do not converge.

[18]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.

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 $\frac{2\Delta}{q_s+q_d}$. Using the chain rule and Proposition 4,

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

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 segregated networks agents eventually agree 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. With binary actions, two groups can disagree about the state of the world even when the number of connections across the groups is unbounded.

The contrasting results for 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, 1992; Bikhchandani, Hirshleifer, and Welch, 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 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 $u_i(a_i, \omega) \coloneqq -(a_i - \omega)^2$ given their beliefs about ω. This utility function now implies that an agent chooses the action corresponding to the state of world she believes is more likely.

In this section, we consider the random network analog of the two-groups model from Section 4.3, which was studied there in a weighted networks setting (the conclusions about disagreement are the same on random networks and weighted networks). Odd- and even-numbered agents are in different groups. Let q_s be the probability of link formation within groups and q_d be the probability of link formation between groups. So when $i > j$, agent i observes agent j with probability q_s if $i \equiv j \pmod 2$, and with probability q_d otherwise. This is a simple directed case of stochastic block networks, which have been used in economics to consider the effect of homophily on speed of learning (Golub and Jackson, 2012) and are widely studied in computer science and related fields (e.g. Airoldi, Blei, Feinberg, and Xing, 2008). Our result extends to two groups of unequal sizes as long as for all later agents, the expected number of observations within their group is larger than the expected number of observations from the other group.

We will state our disagreement result for symmetric binary signals[19] and for Gaussian signals. These functional forms can be relaxed without altering the result, and

[19] With the action space changed, we no longer need unbounded signal informativeness (Assumption 1).

the key assumptions on the signal structure are stated in the proof.[20] Most importantly, signal structures with fat tails need to be ruled out so that the probability of an agent seeing an extremely informative private signal that overturns the information contained in her neighbor's actions vanishes to zero fast enough along the sequence.

Assumption 5 (Binary Signals)**.** *$s_i \in \{0, 1\}$, and $s_i = \omega$ with probability $p > \frac{1}{2}$ in either state.*

This says that signals are binary and symmetric. Note that since all agents choose binary actions, the naive inference assumption still makes sense even with binary signals.

Theorem 2. *Consider a random network with agents partitioned into odd and even groups. Suppose $q_s > q_d > 0$, agents play binary actions, and the environment either has Gaussian signals from Assumption 2 or binary signals from Assumption 5. Then there is a positive probability (with respect to the joint distribution over network structures and signal realizations) that all odd-numbered agents choose action 0 while all even-numbered agents choose action 1.*

The two groups can disagree even though agent n observes approximately $nq_d/2$ agents from the other group (when n is large) taking opposite actions. Note that in this random network model, an agent counts each observed predecessor from her own social group with the same weight as an observed predecessor from the opposite social group—the disagreement result will only be strengthened if we assume agents discount observations from the opposite social group. By contrast, with rational agents and Gaussian signals, Theorem 2 of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) implies the groups agree asymptotically.

This result adds a new mechanism to the literature on public disagreement. Acemoglu, Como, Fagnani, and Ozdaglar (2013) also study a model generating disagree-

[20] This result also holds under any signal structure satisfying the three technical assumptions of *No Uninformative Signals*, *Symmetry*, and *Thin Tails* defined in the proof of Theorem 2.

ment via network learning, but their disagreement takes the form of persistent fluctuations in beliefs whereas our model predicts that two subgroups will converge to opposite extreme beliefs with positive probability. In Acemoglu, Como, Fagnani, and Ozdaglar (2013), disagreement arises from "stubborn" agents who hold different initial opinions and never change their mind. Here, the key forces driving disagreement are the segregated network structure as well as the tendency of naive agents to herd. Bohren and Hauser (2017) 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 Theorem 2 says when all agents use the same naive heuristic, they can still disagree by virtue of belonging to two different homophilic social groups.

6 Experimental Evidence

6.1 Theoretical Motivation

Proposition 3 suggests an empirical test for the naive inference assumption: in the context of sequential learning on random networks, does increasing the link-formation probability q cause more mislearning? In this section, we experimentally test this comparative static in networks of 40 agents by comparing learning outcomes when $q = \frac{1}{4}$ and $q = \frac{3}{4}$. For the sake of experimental clarity we assume that the action space is binary.[21] Our experimental setup features Gaussian signals, binary actions, and random networks as in Section 5.

The naive-learning model and the rational-learning model make competing predictions about this comparative static. The intuition for naive learning comes from Proposition 3, which suggests that in this case later agents are less likely to correctly guess the state on denser networks. We do not expect human subjects to behave

[21]We felt subjects would have difficulty interpreting continuous actions of neighbors, and that their interpretation would be sensitive to the framing of the experiment.

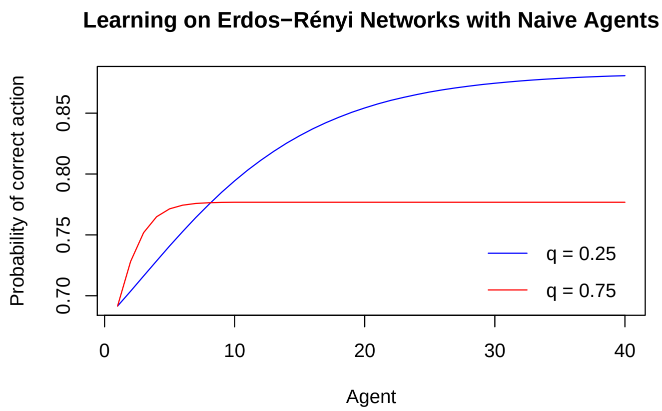

Figure 3: Learning on Erdős–Rényi random networks with 40 naive agents, binary actions and $\sigma^2 = 4$. Networks with link probabilities $q = \frac{1}{4}$ (blue curve) and $q = \frac{3}{4}$ (red curve).

exactly as in the naive model—for example, the meta-analysis of Weizsäcker (2010) reports that laboratory subjects in sequential learning games suffer from autarky bias, underweighting their social observations relative to the payoff-maximizing strategy.[22]

We compute the probabilities of correct actions with 40 naive agents and the network and signal structures corresponding to our experimental design, which we describe in Section 6.2. Because naive agents' actions only depend on the number of their predecessors choosing L and R and not the order of these actions, recursively calculating the distributions of actions is computationally feasible (see Appendix B.2 for details). As shown in Figure 3, early naive agents do worse under $q = \frac{1}{4}$ because there is very little social information, but the comparison quickly switches as we examine later naive agents due to the same logic as in Proposition 3.

On the other hand, the rational learning model predicts that later agents will have either similar or greater accuracy on the dense $q = \frac{3}{4}$ network compared to the sparse $q = \frac{1}{4}$ network. Asymptotically, Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) show that in an environment matching our experimental setup, rational agents will learn

[22]However, the comparative static prediction of the naive model remains robust even after introducing any fraction of autarkic agents who follow only private signals. The naive agents eventually agree, and the probability these agents settle on the correct consensus decreases in q.

the true state in the long-run, regardless of the network density. We can confirm that 40 rational agents are enough to approach this asymptotic learning limit when $q = \frac{3}{4}$. To do this, we compute a lower bound for the probability of correct learning for each agent i, under parameter values matching our experimental design and assuming all agents are rational Bayesians (see Appendix B.1 for details). This lower bound is based on agent strategies depending only on their private signal and the action of just one neighbor, as in the neighbor choice functions in Lobel and Sadler (2015). This exercise shows that the 33^{rd} rational agent is correct at least 96.8% of the time, with the lower bound on the probability of correct learning continuing to increase up to the 40^{th} agent, who is correct at least 97.5% of the time. In addition to suggesting that the asymptotic result of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) very likely holds by the 40^{th} agent, the fact that this lower bound for accuracy on the dense network is so close to perfect learning proves the 40^{th} rational agent could not perform substantially better on the sparse network.

Finally, we intuitively expect more connections to also help rational agents in the short- and medium-run as they can adjust for potential redundancies in information. For example, on the complete network with continuous actions, rational agents can back out private signals of every predecessor by observing their actions, so every agent i does better on the complete network than under any less dense network structure.

We experimentally test the predictions of the naive and the rational models by evaluating the comparative static as we vary network density. We thus provide indirect evidence for the naive inference assumption, complementing the direct measurement of behavior in Eyster, Rabin, and Weizsacker (2015) and Mueller-Frank and Neri (2015).

Beyond providing another form of evidence, we believe this indirect test is a valuable complement to direct tests of behavior because we use the welfare-relevant outcome, namely the accuracy of beliefs, as our dependent variable. Even if individual behavior tends to match redundancy neglect models in simple or stylized settings, one might be concerned that in practice theoretical results about aggregate learning need not hold for complex environments. So if a policy intervention altering the observation network is feasible, for example, experiments using welfare-relevant outcomes as their dependent variables give more explicit guidance as to the consequences of that change.

6.2 Experimental Design

We conducted our experiment on the online labor platform Amazon Mechanical Turk (MTurk) using Qualtrics survey software.

We pre-registered our experimental protocol and regression specification, including the dependent variable to measure the accuracy of social learning and the target sample size, prior to the start of the experiment in August 2017. The pre-registration document is included in the Online Appendix and may also be accessed via the registry website at https://aspredicted.org/yp6eq.pdf.

We recruited 1040 subjects satisfying the selection criteria described in the Online Appendix. Each subject also needed to complete three comprehension questions (which were scenarios in the game with a dominant choice); MTurk users who incorrectly answered one or more comprehension questions were excluded from the experiment. The experiment was carried out in fall 2017.

Each trial consisted of 40 agents who were asked to each make a binary guess between two *a priori* equally likely states of the world, L (for left) and R (for right). The states were color-coded to make instructions and observations more reader-friendly. Agents are assigned positions in the sequence and move in order. Each MTurk subject participated in 10 trials, all in the same position (depending on when they participated in the experiment). The grouping of subjects into trials was independent across trials. Subjects received \$0.25 for completing the experiment and \$0.25 per correct guess, for a maximum possible payment of \$2.75. Subjects ordinarily took less than 10 minutes to complete their participation and earned on average \$2.08, so the incentives were quite large for an MTurk task.

In each trial, every agent received a private signal, which had the Gaussian distribution $\mathcal{N}(-1, 4)$ in state L and the Gaussian distribution $\mathcal{N}(1, 4)$ in state R. These distributions were presented visually in the instructions. Along with the value of their signal, subjects were told the probability of each state conditional on only their private signal.

Each trial was also associated with a density parameter, either $q = \frac{1}{4}$ or $q = \frac{3}{4}$. A random network was generated for each trial by linking each agent with each predecessor with probability q. Each MTurk subject was assigned into either the "sparse" or the "dense" treatment, then placed into 10 trials either all with $q = \frac{1}{4}$ or all with $q = \frac{3}{4}$. So there were 520 subjects and 130 trials for each treatment. Agents were told the actions of each linked predecessor and the link probability q (but not the full realized network, which could not be presented succinctly).

In each trial agents viewed their private signal and any social observations and were asked to guess the state. States, signals, and networks were independently drawn across trials. An example of a choice screen from a trial is shown in Appendix C.

6.3 Results

Let $y_{i,j}$ be the indicator random variable with $y_{i,j} = 1$ if agent i in trial j correctly guesses the state, $y_{i,j} = 0$ otherwise. Define $\tilde{y}_j := \frac{1}{8}\sum_{i=33}^{40} y_{i,j}$ as the fraction of the last eight agents in trial j who correctly guess the state. We test learning outcomes for the final eight agents because welfare depends on long-run learning outcomes in large societies and these agents better approximate long-run outcomes. By using only her private signal, an agent can correctly guess the state 69.15% of the time.[23] We call $\tilde{y}_j - 0.6915$ the *gain from social learning* in trial j, as this quantity represents improvement relative to the autarky benchmark.

We find that the average gain from social learning is 8.73 percentage points for the

[23]In fact, 94% of the subjects assigned to the first position (who have no social observations) correctly use their private signal.

	(1) FractionCorrect	(2) FractionCorrect
NetworkDensity	-0.0923	-0.0923
	(0.0406)	(0.0406)
Constant	0.802	0.802
	(0.0227)	(0.0218)
Observations	260	260
Adjusted R^2	0.016	0.016

(1) without robust SEs; (2) with robust SEs

Table 1: Regression results for the effect of network density on learning outcomes.

$q = \frac{1}{4}$ treatment and 4.12 percentage points for the $q = \frac{3}{4}$ treatment. Social learning improves accuracy on the sparse networks by twice as much as on the dense networks. To test for statistical significance, we consider the regression

$$\tilde{y}_j = \beta_0 + \beta_1 q_j + \epsilon_j$$

where $q_j \in \{\frac{1}{4}, \frac{3}{4}\}$ is the network density parameter for trial j. Recall that each subject was assigned into ten random trials with the same network density and in the same sequential position. This means for two different trials $j^{'} \neq j^{''}$, the error terms $\epsilon_{j'}$ and $\epsilon_{j''}$ are close to independent since there are likely very few subjects who participated in both trials. Indeed, our estimates are identical whether we use robust standard errors or not.

We estimate $\beta_1 = -0.092$ with a p-value of 0.0239 (see Table 1). These findings are consistent with naive updating but not with rational updating, as discussed in Section 6.1.[24]

This difference in the gains from social learning is not driven by different rates of autarky among the two treatments for the last eight agents. We say an agent *goes*

[24]We pre-registered average accuracy in the last 8 rounds (i.e last 20% of agents) as the dependent variable for the experiment, but the regression result is robust to other definitions of $\tilde{y}_j$. When $\tilde{y}_j$ encodes average accuracy among the last m agents for any $4 \leq m \leq 12$ (i.e. between last 10% and last 30% of the agents), the estimate for β_1 remains negative.

against her signal if she guesses L when her signal is positive or guesses R when her signal is negative. Within the last eight rounds, there are 138 instances of agents going against their signals in the $q = \frac{1}{4}$ treatment, which is very close to the 136 instances of the same under the $q = \frac{3}{4}$ treatment. However, when agents go against their signals in the last eight rounds, they correctly guess the state 81.88% of the time under the $q = \frac{1}{4}$ treatment, but only 71.32% of the time under the $q = \frac{3}{4}$ treatment. This shows the observed difference in accuracy is due to social learning being differentially effective on the two network structures.

However, the $q = \frac{3}{4}$ treatment yields better learning outcomes for early agents. For agents 10 through 20, the average accuracy is 72.24% under the $q = \frac{1}{4}$ treatment and 73.22% under the $q = \frac{3}{4}$ treatment. So, if we replace the dependent variable in the pre-registered regression with overall accuracy $\bar{y}_j := \frac{1}{40}\sum_{i=1}^{40} y_{i,j}$, then we do not find a statistically significant estimate for β_1 (q-value of 0.663). This result is consistent with the naive-learning model: according to the simulations in Section 6.1, for early agents accuracy is higher under $q = \frac{3}{4}$, but eventually accuracy is higher under $q = \frac{1}{4}$. The point of overtaking happens at a later round in practice than in theory, because our experimental subjects rely more on their private signal than predicted by the naive model,[25] consistent with the meta-analysis of Weizsäcker (2010).

We do not directly test alternate behavioral models for two reasons. First, given the complex signal and network structures, such tests will be very noisy in our data. Second, because the spaces of possible networks and actions have very high dimension, determining the action each agent would take assuming common knowledge of rationality is computationally infeasible.

[25] The overall frequency of agents going against their signals was 36.8% of the predicted frequency under the naive model.

7 Conclusion

In this paper, we have explored the influence of network structures on learning outcomes when agents move sequentially and suffer from inferential naiveté. We have compared long-run welfare across networks, both theoretically by deriving the exact probabilities of mislearning on arbitrary networks, and experimentally by measuring learning accuracy on sparse and dense 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 with a given turn order is the same as when that order is certain.

Broadly speaking, our model of naive sequential learning is likely to fit well in settings with these three key features: (1) players' decisions are not too close to simultaneous; (2) decisions are hard to modify; (3) the network is complex. The motivating examples from the introduction, including product adoption by a group of consumers and treatment choice by generations of medieval doctors, are well-approximated by a sequential timing structure. Also, once made, these decisions are difficult to reverse.

The third key feature—a complex network—encourages naiveté in at least two ways. First, naive inference is more likely when the rational action is hard to calculate. Second, in complex networks agents are more likely to learn from only predecessors' actions rather than communicate using arbitrarily complex messages—for example reporting the source of every piece of information they have accumulated so far. This kind of direct communication and source-tagging appear more natural in small, close-knit communities, such as the college-campus community where Mobius, Phan, and Szeidl (2015)'s experiment took place. But in many natural sequential-learning problems players are geographically or temporally dispersed, such as in the two examples above. In these settings, it may be too costly to keep track of the original sources of all pieces of information, and in any event close to impossible to communicate this nuanced information to future generations.

Appendix

A Proofs

A.1 Proof of Lemma 1

Proof. Due to naiveté, i thinks neighbor j must have received signal s_j such that $\tilde{s}_j = \tilde{a}_j$. The log-likelihood ratio of state $\omega = 1$ and state $\omega = 0$ conditional on some signal realizations is:

$$\begin{aligned}
\ln\left(\frac{\mathbb{P}[\omega=1|s_i,(s_j)_{j\in N_i}]}{\mathbb{P}[\omega=0|s_i,(s_j)_{j\in N_i}]}\right) &= \ln\left(\frac{\mathbb{P}[s_i,(s_j)_{j\in N_i}|\omega=1]}{\mathbb{P}[s_i,(s_j)_{j\in N_i}|\omega=0]}\right) \text{ (two states equally likely)} \\
&= \ln\left(\frac{\mathbb{P}[s_i|\omega=1]}{\mathbb{P}[s_i|\omega=0]}\cdot\prod_{j\in N_i}\frac{\mathbb{P}[s_j|\omega=1]}{\mathbb{P}[s_j|\omega=0]}\right) \text{ (by independence)} \\
&= \ln\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right)+\sum_{j\in N_i}\ln\left(\frac{\mathbb{P}[\omega=1|s_j]}{\mathbb{P}[\omega=0|s_j]}\right) \\
&= \tilde{s}_i+\sum_{j\in N_i}\tilde{s}_j.
\end{aligned}$$

A naive agent i chooses an action with $\tilde{a}_i$ equal to the log-likelihood ratio of state $\omega = 1$ and state $\omega = 0$ given private signal s_i and neighbors' signals such that $\tilde{s}_j = \tilde{a}_j$, which is $\tilde{s}_i + \sum_{j\in N_i} \tilde{a}_j$. □

A.2 Proof of Lemma 2

Proof. We check the formula for each of the two interpretations. For the noisy observation interpretation, i observes $\tilde{a}_j + \epsilon_{i,j}$ for each neighbor j and believes that $\tilde{s}_j = \tilde{a}_j$. By Assumption 2 and Lemma 3 (which is proven independently), $\tilde{s}_j$ is distributed as $\mathcal{N}(\frac{2}{\sigma^2}, \frac{4}{\sigma^2})$ when $\omega = 1$, so that i believes their observation has distribution $\mathcal{N}(\frac{2}{\sigma^2}, \frac{1}{M_{i,j}} \cdot \frac{4}{\sigma^2})$ when $\omega = 1$ and $\mathcal{N}(-\frac{2}{\sigma^2}, \frac{1}{M_{i,j}} \cdot \frac{4}{\sigma^2})$ when $\omega = -1$. So agent i believes that $\frac{\sigma^2}{2}(\tilde{a}_j + \epsilon_{i,j}) \sim \mathcal{N}(\pm 1, \frac{\sigma^2}{M_{i,j}})$, depending on the state. Applying Lemma 3 again, we find that i then believes the log-likelihood ratio of $\omega = 1$ and $\omega = 0$ is $M_{i,j}(\tilde{a}_j + \epsilon_{i,j})$. In the case when observations are correct, this simplifies to $M_{i,j}\tilde{a}_j$. The proof now proceeds as in the proof of Lemma 1.

We next show the lemma holds under the generations interpretation. By the law of large numbers a fraction $M_{i,j}$ of the agents in generation i observe an agent in generation j, and the average observed log-action among this group is equal to $\tilde{a}_j$. The result follows from computing each member of generation i's action as in the proof of Lemma 1 and then averaging. □

A.3 Proof of Proposition 1

Proof. By Lemma 2, 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:

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

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

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

A.4 Proof of Lemma 3

Proof. The log-likelihood ratio is

$$\ln\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right)=\ln\left(\frac{\mathbb{P}[s_i|\omega=1]}{\mathbb{P}[s_i|\omega=0]}\right)=\ln\left(\frac{\exp\left(\frac{-(s_i-1)^2}{2\sigma^2}\right)}{\exp\left(\frac{-(s_i+1)^2}{2\sigma^2}\right)}\right)$$

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

A.5 Proof of Lemma 4

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

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

A.6 Proof of Theorem 1

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

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

A.7 Proof of Proposition 2

Proof. Without loss of generality, assume $\omega=1$. (The case of $\omega=0$ is exactly analogous and will be 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 $\tilde{s}$ has finite variance, 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 $\|\vec{b}_n\|_1 \mathbb{E}[\tilde{s}_1|\omega = 1]$ and standard deviation $\|\vec{b}_n\|_2 \mathrm{Std}[\tilde{s}_1|\omega = 1]$. We now use the connectedness assumption to show that $\frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2} \to \infty$.

Find N and $C \leq 1$ as in the connectedness assumption. 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 assumption. This shows for any $j \in \mathbb{N}$ and for $n > M_\epsilon$, 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, $\frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2}$ 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

$$z_n := \frac{\mathbb{E}[\tilde{a}_n|\omega = 1] - K}{\mathrm{Std}[\tilde{a}_n|\omega = 1]} = \frac{\|\vec{b}_n\|_1 \cdot \mathbb{E}[\tilde{s}_1|\omega = 1]}{\|\vec{b}_n\|_2 \cdot \mathrm{Std}[\tilde{s}_1|\omega = 1]} - \frac{K}{\|\vec{b}_n\|_2 \cdot \mathrm{Std}[\tilde{s}_1|\omega = 1]}.$$

Since we assumed the signal structure is non-trivial, $\mathbb{E}[\tilde{s}_1|\omega = 1] > 0$. Together with $\frac{\|\vec{b}_n\|_1}{\|\vec{b}_n\|_2} \to \infty$, this shows the first term converges to infinity. By the connectedness assumption, $\|\vec{b}_n\|_2 \geq C$ for all large enough n, so the second term is bounded. This implies $z_n \to \infty$. By Chebyshev's inequality, $\mathbb{P}[\tilde{a}_n < K|\omega = 1] \leq \frac{1}{z_n^2}$. This shows $\mathbb{P}[\tilde{a}_n < K|\omega = 1] \to 0$. □

A.8 Proof of Proposition 3

Proof. 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

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

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

Now we show $\mathbb{P}[a_n \to 0|\omega = 1] = \Phi\left(-\sqrt{\frac{q+2}{q}}\right)$, which is the same probability as $\mathbb{P}[\tilde{a}_n \to -\infty|\omega = 1]$. The random variable $Y(\vec{s})$ that $X_n(\vec{s})$ converges to a.s. has the distribution $\mathcal{N}(\frac{1}{q}, \frac{1}{q(q+2)})$ when $\omega = 1$, and $\tilde{a}_n$ has the same sign as $X_n(\vec{s})$ with probability converging to 1 for n large. The distribution $\mathcal{N}(\frac{1}{q}, \frac{1}{q(q+2)})$ assigns

$1-\Phi\left(-\sqrt{\frac{q+2}{q}}\right)$ probability to the positive region. The symmetric argument holds for $\omega = 0$. □

A.9 Proof of Proposition 4

Proof. 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_-, ξ_+, ξ_- so that

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

where $\xi_\pm$ 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

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

where $\xi_+ > 1$ and $\xi_- < 0$. By arguments analogous to those in the proof of Proposition 3, we may again establish that a_n converges to 0 or 1 almost surely. We now analyze the probability of mislearning.

Since $\xi_+ > |\xi_-|$, the exponential term with base ξ_+ dominates as n grows large. This shows $c_+ > 0$, since $b_{n,i}$ counts number of weighted paths in a network so must be a positive number. This also shows that $\mathbb{P}[\tilde{a}_n < 0 | \omega = 1] \to \mathbb{P}[\sum_{i=0}^{\infty} (\xi_+)^{-i} \tilde{s}_i < 0 | \omega = 1]$ as $n \to \infty$. Conditional on $\omega = 1$, the sum $\sum_{i=0}^{\infty} (\xi_+)^{-i} \tilde{s}_i$ has the distribution $\mathcal{N}(\frac{2}{\xi_+ - 1}, \frac{4}{(\xi_+ - 1)(\xi_+ + 1)})$, so it is easy to show the probability assigned to the negative region is increasing in ξ_+.

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

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

It is easy to see that $\frac{\partial \xi_+}{\partial q_d} > \frac{\partial \xi_+}{\partial q_s} > 0$ for all $q_s \geq 0$ and $q_d > 0$. □

A.10 Proof of Theorem 2

Proof. We will prove a more general statement where disagreement happens if the signal structure satisfies the following three assumptions:

No Uninformative Signals: Signals s for which $\mathbb{P}[\omega = 1|s] = \mathbb{P}[\omega = 0|s]$ have 0 probability. So, we may decompose (up to a $\mathbb{P}$-null set) $S = S^+ \cup S^-$ where

$$S^+ := \{s \in S : \mathbb{P}[\omega = 1|s] > \mathbb{P}[\omega = 0|s]\}$$

$$S^- := \{s \in S : \mathbb{P}[\omega = 1|s] < \mathbb{P}[\omega = 0|s]\}.$$

Symmetry:

$$\frac{\mathbb{P}[\omega = 1|s \in S^+]}{\mathbb{P}[\omega = 0|s \in S^+]} = \frac{\mathbb{P}[\omega = 0|s \in S^-]}{\mathbb{P}[\omega = 1|s \in S^-]}.$$

Thin Tails:

$$\sum_{n=0}^{\infty} \mathbb{P}\left[s \in S : \ln\left(\frac{\mathbb{P}[\omega = 1|s]}{\mathbb{P}[\omega = 0|s]}\right) > n|\omega = 1\right] < \infty \text{ and}$$

$$\sum_{n=0}^{\infty} \mathbb{P}\left[s \in S : \ln\left(\frac{\mathbb{P}[\omega = 1|s]}{\mathbb{P}[\omega = 0|s]}\right) < -n|\omega = 1\right] < \infty.$$

No Uninformative Signals is for convenience and rules out tie-breaking problems. *Thin Tails* says when the state is $\omega = 1$, extremely informative signals in favor of either state are not too common. This assumption is automatically satisfied if the signal structure is boundedly informative, such as in the case of binary signal structure from Assumption 5.

The Gaussian signal structure from Assumption 2 satisfies *No Uninformative Signals* and *Symmetry*, with $S^+ = \mathbb{R}_{>0}$ and $S^- = \mathbb{R}_{<0}$. For *Thin Tails*, $\ln\left(\frac{\mathbb{P}[\omega=1|s]}{\mathbb{P}[\omega=0|s]}\right) = 2s/\sigma^2$ under the Gaussian assumption, so we are looking at the sum

$$\sum_{n=0}^{\infty} \mathbb{P}[2\mathcal{N}(1,\sigma^2)/\sigma^2 > n].$$

This is convergent because it is known that the Gaussian distribution function tends to 0 faster than geometrically. Similarly we also have $\sum_{n=0}^{\infty} \mathbb{P}[2\mathcal{N}(1,\sigma^2)/\sigma^2 < -n] < \infty$.

So it suffices to show the Theorem under assumptions *No Uninformative Signals*, *Symmetry* and *Thin Tails*. We make use of the following concentration inequality from Hoeffding (1963):

Fact 1. *(Hoeffding) Suppose $X_1, ..., X_n$ are independent random variables on $\mathbb{R}$ such that $a_i \le X_i \le b_i$ with probability 1 for each i. Write $S_n := \sum_{i=1}^{n} X_i$. We have*

$$\mathbb{P}\left[S_n - \mathbb{E}[S_n] \le -t\right] \le \exp(-2t^2/\sum_{i=1}^{n}(b_i - a_i)^2).$$

Define κ to be a naive agent's log-likelihood ratio of state $\omega = 1$ versus state $\omega = 0$ upon observing one neighbor who picks action 1. Then we have:

$$\kappa := \ln\left(\frac{\mathbb{P}[\omega = 1|s \in S^+]}{\mathbb{P}[\omega = 0|s \in S^+]}\right) > 0.$$

By *Symmetry*, the log-likelihood ratio after observing one neighbor who chooses action 0 is $-\kappa$.

Suppose after $2n$ agents have moved, the actions taken so far involve every odd-numbered agent playing 1, every even-numbered agent playing 0. Write $B_{n,q_s} \sim \text{Binom}(n, q_s)$ for the number of agents that $2n+1$ observes from her own group, and $B_{n,q_d} \sim \text{Binom}(n, q_d)$ is the number of agents $2n+1$ observes from the other group.

By Hoeffding's inequality, the probability that $B_{n,q_s} - B_{n,q_d}$ falls below half of its expected value $n(q_s - q_d)$ is:

$$\epsilon_n^{(1)} := \mathbb{P}\left[(B_{n,q_s} - B_{n,q_d}) - n(q_s - q_d) \leq -\frac{1}{2}n(q_s - q_d)\right] \leq \exp(-\frac{1}{2}n^2(q_s - q_d)^2/(2n)).$$

But in the event that $B_{n,q_s} - B_{n,q_d} \geq \frac{1}{2}n(q_s - q_d)$, $2n + 1$ has a log-likelihood ratio of at least $\frac{1}{2}n(q_s - q_d)\kappa$ from observations alone. The probability that private signal s_{2n+1} is so strongly in favor of $\omega = 0$ as to make $2n + 1$ play 0 is:

$$\epsilon_n^{(2)} := \mathbb{P}\left[s \in S : \ln\left(\frac{\mathbb{P}[\omega = 1|s]}{\mathbb{P}[\omega = 0|s]}\right) < \frac{1}{2}n(q_s - q_d)\kappa|\omega = 1\right]$$

We have $\sum_{n=1}^{\infty} \epsilon_n^{(1)} < \infty$ because $\epsilon_n^{(1)} \sim \exp(-n)$, while $\sum_{n=1}^{\infty} \epsilon_n^{(2)} < \infty$ due to *Thin Tails* assumptions. This shows that there is positive probability that every odd-numbered agent plays 1.

By *Symmetry*, an analogous argument establishes that there is also positive probability that every even-numbered agent plays 0. □

B Predictions in the Experimental Environment

B.1 Bounding the Performance of Rational Agents

Consider 40 rational agents on a random network where each agent is linked to each of her predecessors $\frac{3}{4}$ of the time, i.i.d. across link realizations. Agents know their own neighbors but have no further knowledge about the realization of random network. The signal structure and payoff structure match Section 6.2's experimental design.

We provide a lower bound for the accuracy of agents 33 through 40. We first show that when every player uses the rational strategy, all agents learn at least as well as when everyone uses any constrained strategy that only depends on their signal and their final social observation. We then exhibit payoffs under one such strategy, which give a lower bound on rational performance.

Fix an arbitrary sequence of constrained strategies (σ_i) where $\sigma_i : S_i \times \{0, 1, \emptyset\} \to \Delta(\{0, 1\})$ is only a function of i's signal s_i and the action of the most recent predecessor that i observes ($\sigma_i(s_i, \emptyset)$ refers to i's play if i does not observe any predecessor). Let a_i denote i's action induced by this sequence of strategies. Let a'_i denote i's action when all agents use the rational strategy.

Claim 1. For all i, $\mathbb{P}[a'_i = \omega] \geq \mathbb{P}[a_i = \omega]$.

Proof. The proof is by induction on i and the base case of $i = 1$ is clear. Suppose the claim holds for $i \leq n$. Conditional on agent $n + 1$ observing no predecessors, the claim again holds as in the base case, so we can check the claim conditional on N_{n+1} non-empty.

Let j be the final agent in N_{n+1}. Then the rational agent observes s_{n+1}, a'_j for some $j \leq n$, and perhaps some other actions while the constrained agent observes s_{n+1} and a_j, where $\mathbb{P}[a'_j = \omega] \geq \mathbb{P}[a_j = \omega]$ by the inductive hypothesis. By garbling the observed action a'_j, the rational agent could construct a random variable with the same joint distribution with ω as the less accurate action a_j. Ignoring other observed actions, the rational agent $n+1$ could therefore follow a strategy that does as well as agent $n+1$ under the strategy profile (σ_i). So we must have $\mathbb{P}[a'_{n+1} = \omega] \geq \mathbb{P}[a_{n+1} = \omega]$ when everyone uses the rational strategy. □

We then numerically compute the values for $\mathbb{P}[a_i = \omega]$ under the optimal constrained strategy, which are displayed in Table 2.

agent number	33	34	35	36	37	38	39	40
probability correct	0.9685	0.9695	0.9705	0.9714	0.9723	0.9731	0.9739	0.9746

Table 2: Lower bounds on rational performance.

B.2 Performance of Naive Agents

Consider 40 naive agents on a random network where each agent is linked to each of her predecessors with probability q, i.i.d. across link realizations. The signal structure and payoff structure match the experimental design in Section 6.2.

Because naive agents' actions do not depend on the order of predecessors, behavior depends only on the number of agents who have played L and the number of agents who have played R as well as the network. We will compute the distribution over the number of agents from the first n who have played L and the number who have played R recursively.

Assume the state is R. Let $P(k, k')$ be the probability that k of the first n agents play L and k' of the first n agents play R. We define $P(k, k') = 0$ if $k < 0$ or $k' < 0$. The posterior log-likelihood of state R for a naive agent observing one R action (and no signal) is $\ell = \frac{2}{\sigma^2} \cdot \frac{\mu+\sigma\phi(-\mu/\sigma)}{1-\Phi(-\mu/\sigma)}$, where Φ and ϕ are the distribution function and probability density function of a standard Gaussian random variable, respectively.

Then we have the recursive relation

$$\begin{aligned} P(k,k') &= P(k-1,k') \textstyle\sum_{i \le k-1, i' \le k'} B(i,k-1,q)B(i',k',q)\Phi(\frac{\sigma(i-i')\ell-2\mu\sigma}{2}) + \\ & P(k,k'-1) \textstyle\sum_{i \le k, i' \le k'-1} B(i,k,q)B(i',k'-1,q)[1-\Phi(\frac{\sigma(i-i')\ell-2\mu\sigma}{2})], \end{aligned}$$

where $B(i, k, q)$ is the probability mass function of a binomial distribution with parameters k and q at i. The first summand corresponds to agent $k + k'$ choosing L after $k - 1$ predecessors choose L and the remainder choose R, and the second summand corresponds to agent $k + k'$ choosing R after k predecessors choose L and the remainder choose R. The binomial coefficients correspond to the possible network realizations. Here we use naiveté, which implies that only the number of observed agents choosing each action matters for behavior and not their order.

From these distributions $P(\cdot, \cdot)$ we can compute the probability that agent n

chooses the correct action R:

$$\sum_{k=0}^{n} P(k,n-k) \sum_{i\leq k, i'\leq n-k} B(i,k,q)B(i',n-k,q)[1-\Phi(\frac{\sigma(i-i')\ell - 2\mu\sigma}{2})].$$

These probabilities, which we compute numerically, are displayed in Table 3 for agents 33 through 40.

agent number	33	34	35	36	37	38	39	40
accuracy with $q = 1/4$	0.8773	0.8780	0.8786	0.8792	0.8797	0.8801	0.8805	0.8808
accuracy with $q = 3/4$	0.7768	0.7768	0.7768	0.7768	0.7768	0.7768	0.7768	0.7768

Table 3: Naive performance.

C Experimental Details

In addition to comprehension questions, we restricted to subjects located in the United States who had completed at least 50 previous MTurk tasks with a lifetime approval rate of at least 90%. Subjects were not permitted to participate in more than one round of our experiment. There were 15 subjects who completed fewer than 10 trials, implying a completion rate of 98.5%, and these subjects were excluded and replaced by new subjects.

The three comprehension questions were: one question in which subjects observed player 1, who agreed with their private signal; one question in which the subjects had no social observations and their signal was informative; and one question in which subjects observed player 1 and received the uninformative signal 0. Subjects were instructed to answer assuming player 1 followed their private signal.

An example choice follows.

This is round 1 (out of 10).

Your signal: 1.4239.
Based on your signal alone, there is 32.92% chance the direction is LEFT, 67.08% chance the direction is RIGHT.

Your observations:
Player 1 guessed RIGHT
Player 2 guessed LEFT
Player 5 guessed LEFT
Player 6 guessed LEFT
Player 7 guessed LEFT
Player 8 guessed LEFT
Player 9 guessed LEFT

What is your guess about the direction this round?

○ LEFT
○ RIGHT

References

Acemoglu, D., V. Chernozhukov, and M. Yildiz (2016): "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, 11, 187–225.

Acemoglu, D., G. Como, F. Fagnani, and A. Ozdaglar (2013): "Opinion fluctuations and disagreement in social networks," *Mathematics of Operations Research*, 38, 1–27.

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian learning in social networks," *Review of Economic Studies*, 78, 1201–1236.

Acemoglu, D., A. Ozdaglar, and A. ParandehGheibi (2010): "Spread of (mis) information in social networks," *Games and Economic Behavior*, 70, 194–227.

Airoldi, E. M., D. M. Blei, S. E. Feinberg, and E. P. Xing (2008): "Mixed Membership Stochastic Block Models," *Journal of Machine Learning Research*, 9, 1981–2014.

Bala, V. and S. Goyal (1998): "Learning from neighbours," *Review of Economic Studies*, 65, 595–621.

Banerjee, A. V. (1992): "A simple model of herd behavior," *Quarterly Journal of Economics*, 107, 797–817.

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

Bohren, J. A. (2016): "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247.

Bohren, J. A. and D. Hauser (2017): "Bounded Rationality And Learning: A Framework and A Robustness Result," *Mimeo*.

Chandrasekhar, A. G., H. Larreguy, and J. P. Xandri (2015): "Testing models of social learning on networks: Evidence from a lab experiment in the field," *Mimeo*.

Coleman, J. (1958): "Relational analysis: the study of social organizations with survey methods," *Human Organization*, 17, 28–36.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003): "Persuasion bias, social influence, and uni-dimensional opinions," *Quarterly Journal of Economics*, 118, 909–968.

Enke, B. and F. Zimmermann (2018): "Correlation neglect in belief formation," *Review of Economic Studies*.

Eyster, E. and M. Rabin (2010): "Naive herding in rich-information settings," *American Economic Journal: Microeconomics*, 2, 221–243.

——— (2014): "Extensive Imitation is Irrational and Harmful," *Quarterly Journal of Economics*, 129, 1861–1898.

Eyster, E., M. Rabin, and G. Weizsacker (2015): "An Experiment on Social Mislearning," *Mimeo.*

Gagnon-Bartsch, T. and M. Rabin (2016): "Naive Social Learning, Mislearning, and Unlearning," *Mimeo.*

Golub, B. and M. O. Jackson (2010): "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2, 112–149.

——— (2012): "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *Quarterly Journal of Economics*, 127, 1287–1338.

Hoeffding, W. (1963): "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, 58, 13–30.

Lee, I. H. (1993): "On the convergence of informational cascades," *Journal of Economic Theory*, 61, 395–411.

Levy, G. and R. Razin (2018): "Information diffusion in networks with the Bayesian Peer Influence heuristic," *Games and Economic Behavior.*

Lobel, I. and E. Sadler (2015): "Information diffusion in networks through social learning," *Theoretical Economics*, 10, 807–851.

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

Mobius, M., T. Phan, and A. Szeidl (2015): "Treasure hunt: Social learning in the field," *Mimeo.*

Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2018): "A Theory of Non-Bayesian Social Learning," *Econometrica*, 86, 445–490.

Mueller-Frank, M. and C. Neri (2015): "A general model of boundedly rational observational learning: Theory and evidence," *Mimeo*.

Sadler, E. (2017): "Innovation adoption and collective experimentation," .

Sethi, R. and M. Yildiz (2012): "Public disagreement," *American Economic Journal: Microeconomics*, 4, 57–95.

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

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