Network Structure and Social Learning

KRISHNA DASARATHA
Yale University
and
KEVIN HE
University of Pennsylvania

We describe results from Dasaratha and He [DH21a] and Dasaratha and He [DH20] about how network structure influences social learning outcomes. These papers share a tractable sequential model that lets us compare learning dynamics across networks. With Bayesian agents, incomplete networks can generate informational confounding that makes learning arbitrarily inefficient. With naive agents, related forces can lead to mislearning.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Economics, Theory

Additional Key Words and Phrases: Social networks, Sequential social learning, Bayesian learning, Confounding

1. INTRODUCTION

Information about a decision-relevant state of nature is often dispersed among a group of agents. When people try to evaluate a new scientific theory or decide which of two consumer products is better, each individual's personal experience provides a small amount of information. While it is often unrealistic for a central institution to elicit everyone's private information and then broadcast a synthesized prediction about the state, society may nevertheless aggregate different people's signals through a process of decentralized *social learning.* In social learning, people make inferences about the state from their private information and their social observations — other people's opinions on the new theory or others' product adoption choices provide information about the state. More private signals gradually become incorporated into the social consensus as agents act in turn.

In such social-learning settings, interpersonal ties often matter for the flow of information. Typically, people do not observe the choices of all predecessors who have faced the same decision problem in the past, but only those of their "neighbors" in a social network. These neighbors traditionally include the family members, friends, and colleagues with whom one regularly interacts on a face-to-face basis, but increasingly also include one's virtual neighbors on social-media platforms. How do structural properties of the observation network affect social learning, and how do changes to the network (e.g., due to developments in communication technology) improve or hinder learning?

As an example of how the network can matter for learning, many realistic network structures obstruct learning through a *confounding* mechanism. Suppose an agent

Authors' addresses: `krishnadasaratha@gmail.com`, `hesichao@gmail.com`

observes several neighbors, but not some earlier predecessors who influenced these neighbors. Then the agent cannot fully separate the neighbors' private information from their (potentially common) social information. This network-generated confounding can seriously harm learning even for fully rational agents, and leads to even worse consequences for naive agents who rely on behavioral heuristics.

To analyze confounding and other effects of the social network, we describe a tractable social-learning model designed to compare social-learning dynamics across networks. We start with the standard sequential learning environment [Ban92; BHW92], and then equip agents with rich actions that fully convey beliefs and Gaussian private signals. These assumptions let us focus on obstructions to learning that arise from the observation network (as opposed to obstructions due to coarseness of the action space or signal space, e.g., Rosenberg and Vieille [RV19], Harel, Mossel, Strack, and Tamuz [HMST21]).

This model lets us analyze Bayesian learning [DH21a] and naive learning [DH20], and we will discuss results under both regimes. Bayesian agents will learn the true state on all "reasonable" network structures (as in Acemoglu, Dahleh, Lobel, and Ozdaglar [ADLO11]), but we show this learning can be very inefficient. Naive agents will often herd on the incorrect state, and we compute the probability of such mistaken herds. In both cases, we move beyond binary criteria (e.g., do agents eventually learn the true state?) and define richer measures to capture learning outcomes. We compare these measures as we vary features of the social network, delivering a detailed comparison of how changes in the network affect learning.

2. MODEL

Consider a sequence of agents $i = 1, 2, \ldots$ who act in turn, learning about a common binary state $\omega \in \{0, 1\}$. The two states are equally likely ex-ante. Each agent receives an i.i.d. conditionally Gaussian private signal s_i, with the signal drawn from the normal distribution $\mathcal{N}(1, \sigma^2)$ when $\omega = 1$ and $\mathcal{N}(-1, \sigma^2)$ when $\omega = 0$. Agent i also observes the actions $a_j \in [0, 1]$ of her neighbors j in the neighborhood $N(i) \subseteq \{1, 2, \ldots, i-1\}$. The neighborhoods, which are common knowledge, define an observation network. Agents choose actions a_i to maximize the expectation of their utility $u_i(a_i, \omega) = -(a_i - \omega)^2$, so that

$$a_i = \mathbb{E}[\,\omega \mid \text{agent } i\text{'s beliefs}\,].$$

This model that we work with is the standard sequential learning model with two notable assumptions: (1) actions are rich enough to fully reflect beliefs, and (2) private signals are Gaussian.

3. BAYESIAN SOCIAL LEARNING

3.1 General Results

This section follows Dasaratha and He [DH21a] and describes our results with Bayesian agents. We begin with several general results that facilitate calculations in the model. First, actions are log-linear — that is, the log-likelihood $\log(\frac{a_i}{1-a_i})$ is a linear combination of agent i's private signal log-likelihood and her neighbors' log-likelihoods:

$$\log\left(\frac{a_i}{1-a_i}\right) = \log\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right) + \sum_{j\in N_i} \beta_{ij}\log\left(\frac{a_j}{1-a_j}\right)$$

for some coefficients β_{ij}.

Second, we can assign a signal-counting interpretation to the accuracy of every agent's action. If agent i's only information is $n \in \mathbb{N}_+$ independent signals, then her action has the conditional distribution $\log(\frac{a_i}{1-a_i}) \sim \mathcal{N}\left(\pm n \cdot \frac{2}{\sigma^2}, n \cdot \frac{4}{\sigma^2}\right)$, where the sign depends on the state ω. We extend this fact to define a measure of accuracy by interpolating between integer values of n:

Definition 3.1. Social learning **aggregates** $r \in \mathbb{R}_+$ **signals by agent** i if agent i's log-action satisfies $\log(\frac{a_i}{1-a_i}) \sim \mathcal{N}\left(\pm r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2}\right)$.

When agents use arbitrary log-linear strategies, their log-actions may not satisfy the functional-form restriction in the above definition for any r (even though they will be conditionally normally distributed). But the next proposition shows the functional-form restriction is always satisfied for Bayesian agents.

PROPOSITION 3.2. *There exist $(r_i)_{i\geq 1}$ so that for Bayesian agents, social learning aggregates r_i signals by agent i. These $(r_i)_{i\geq 1}$ depend on the network, but not on σ^2.*

This result says we can always quantify a rational agent's accuracy in units of independent private signals. The number r_i is a sufficient statistic for agent i's accuracy, and we can compare learning across networks by comparing their r_i's.

Agents learn the state completely in the long-run if and only if $r_i \to \infty$. It is straightforward to show that there is long-run learning if and only if the mild "expanding neighborhoods" condition from Acemoglu, Dahleh, Lobel, and Ozdaglar [ADLO11] is satisfied. This implies that long-run learning is not a very useful criterion for comparing Bayesian learning across networks, and we will instead focus on how efficiently agents learn.

Definition 3.3. The **aggregative efficiency** of the network is $\lim_{i\to\infty}(r_i/i)$.

Aggregative efficiency measures the fraction of available signals that get incorporated into people's behavior on a certain network. We now show that even networks that enable long-run learning may have arbitrarily low aggregative efficiency.

3.2 Generations Networks

We illustrate the effect of confounding in a simple class of social networks where agents eventually learn the state, but this learning can be arbitrarily inefficient. In *maximal generations networks*, agents arrive in generations of K individuals. Each agent in a generation t observes all agents in the previous generation $t-1$ and has no other neighbors.

We can characterize Bayesian learning outcomes in such networks:

PROPOSITION 3.4. *The aggregative efficiency of a maximal generations network is*

$$\lim_{i\to\infty}(r_i/i) = \frac{(2K-1)}{K^2}.$$

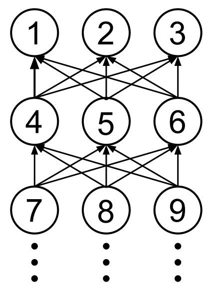

Fig. 1. A maximal generations network with generation size $K = 3$. An arrow from i to j means i observes j's action.

In the long run, social learning aggregates fewer than 2 *signals per generation with any* K*. After generation* 2*, social learning aggregates fewer than* 3 *signals per generation with any* K*.*

Learning aggregates at most three signals per generation beginning with generation 3 and at most two signals per generation asymptotically. This bound holds for any generation size, so learning can be *arbitrarily inefficient*: as K grows large, only a vanishing fraction of the available signals gets aggregated.

The basic intuition is that these networks create severe confounding. An agent in generation t observes many neighbors from the previous generation, who in turn have both independent private signals and common social information from observing generation $t - 2$. The generation t agent trades off overweighting the common social information and underweighting the recent private signals. When t is large, the Bayesian estimate puts almost no weight on the recent private signals. Therefore, very little of the recent information becomes incorporated into beliefs.

Dasaratha and He [DH21a] also allow more general observation structures between generations: each agent may only observe a subset of the previous generation. Under a symmetry assumption, we express aggregative efficiency explicitly in terms of the network parameters. Each generation continues to aggregate at most two additional signals asymptotically, so learning remains very inefficient when generations are large. We also show how structural parameters of the network influence efficiency: all else equal, aggregative efficiency increases when agents have more neighbors and decreases when neighborhoods overlap more (i.e., when there is more confounding).

4. NAIVE SOCIAL LEARNING

We now turn from Bayesian updating to naive updating and present results from Dasaratha and He [DH20]. We maintain the model from Section 2, but ask how the social network structure affects learning outcomes for agents with the following behavioral bias:

Assumption 4.1. Each agent wrongly believes that each predecessor chooses an action to maximize her expected payoff based only on her private signal, and not on her observation of other agents.

This assumption was first studied by Eyster and Rabin [ER10]. Agents are otherwise Bayesian and maximize their expected utility. A key implication is that agents fail to account for the correlation in their observed actions, and instead treat these actions as independent.

As with Bayesian social learning, actions are log-linear for naive agents. Because agents (incorrectly) believe their neighbors' actions are identically distributed, the weights on all observed actions are in fact equal:

$$\log\left(\frac{a_i}{1-a_i}\right) = \log\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right) + \sum_{j\in N_i} \log\left(\frac{a_j}{1-a_j}\right).$$

This implies a simple expression for actions in terms of private signal realizations and the adjacency matrix of the network:

Proposition 4.2. *Suppose the observation network between the first n agents has the adjacency matrix M. The actions of these agents satisfy*

$$\begin{pmatrix} \log\left(\frac{a_1}{1-a_1}\right) \\ \vdots \\ \log\left(\frac{a_n}{1-a_n}\right) \end{pmatrix} = (I-M)^{-1} \cdot \begin{pmatrix} \log\left(\frac{\mathbb{P}[\omega=1|s_1]}{\mathbb{P}[\omega=0|s_1]}\right) \\ \vdots \\ \log\left(\frac{\mathbb{P}[\omega=1|s_n]}{\mathbb{P}[\omega=0|s_n]}\right) \end{pmatrix}$$

where I is the $n \times n$ identity matrix.

The proposition says that actions are a log-linear combination of private signal realizations, and the total weight that agent i places on an earlier agent j's signal is equal to the number of paths from i to j in the observation network.

Unless the network is quite sparse, there are many paths to early agents, so Proposition 4.2 implies that agents will put too much weight on early signals. This can lead to *mislearning*: agents converge to believing strongly in the wrong state. We compare the probability of such mislearning across network structures.

For example, we show the probability of mislearning is higher when the observation network is denser.[1] When the network is denser, there are more paths to early agents, so a few misleading early signals will generate a wrong consensus. A sparser network allows more independent information to accumulate before a consensus forms.

This theoretical result gives a testable prediction about network structure and learning accuracy, and we confirm this prediction in a social-learning experiment with Amazon Mechanical Turk workers [DH21b]. Subjects play a social-learning game on either a sparse or a dense observation network, depending on the treatment. We find that later subjects' guesses about the state are significantly more accurate in the sparse treatment.[2] Despite providing more information, denser networks lead to worse social-learning outcomes because confounding is more severe.

[1] Formally, we prove this result for *weighted networks*, defined in Dasaratha and He [DH20]. We also provide simulation evidence for the same result on random networks.

[2] The experimental setup follows the model from Section 2 except for the action space, which we took to be binary for simplicity.

5. CONCLUSION

We have presented a tractable model of sequential social learning that lets us compare learning across different networks, a topic with important social and economic consequences but limited analytic results in the literature. When agents are rational Bayesians, we show that actions admit a signal-counting interpretation of accuracy, derive analytic expressions for how learning efficiency changes with network parameters, and quantify the extent of information loss due to confounding. When agents learn using a naive heuristic that neglects correlation in predecessors' behavior, we compute the exact probabilities of mislearning in different networks and generate testable predictions that motivate laboratory experiments.

These results attest to the effectiveness and flexibility of the rich-signals, rich-actions framework we use to study social learning in networks. We hope that this framework can provide a complementary set of tools to the current techniques in the social-learning literature, and enable future theoretical and empirical work on which networks are most conducive to social learning.

REFERENCES

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

Abhijit V Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, 1992.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.

Krishna Dasaratha and Kevin He. Network structure and naive sequential learning. *Theoretical Economics*, 15(2):415–444, 2020.

Krishna Dasaratha and Kevin He. Aggregative efficiency of bayesian learning in networks. 2021.

Krishna Dasaratha and Kevin He. An experiment on network density and sequential learning. *Games and Economic Behavior*, 128:182–192, 2021.

Erik Eyster and Matthew Rabin. Naive herding in rich-information settings. *American Economic Journal: Microeconomics*, 2(4):221–243, 2010.

Matan Harel, Elchanan Mossel, Philipp Strack, and Omer Tamuz. Rational groupthink. *Quarterly Journal of Economics*, 136(1):621–668, 2021.

Dinah Rosenberg and Nicolas Vieille. On the efficiency of social learning. *Econometrica*, 87(6):2141–2168, 2019.