

Aggregative Efficiency of Bayesian Learning in Networks*

Krishna Dasaratha[†] Kevin He[‡]

First version: November 21, 2019

This version: April 3, 2025

Abstract

When individuals in a social network learn about an unknown state from private signals and neighbors' actions, the network structure often causes information loss. We consider rational agents and Gaussian signals in the canonical sequential social-learning problem and ask how the network changes the efficiency of signal aggregation. Rational actions in our model are log-linear functions of observations and admit a signal-counting interpretation of accuracy. Networks where agents observe multiple neighbors but not their common predecessors *confound* information, and even a small amount of confounding can lead to much lower accuracy. In a class of networks where agents move in generations and observe the previous generation, we quantify the information loss with an *aggregative efficiency* index. Aggregative efficiency is a simple function of network parameters: increasing in observations and decreasing in confounding. Later generations contribute little additional information, even with arbitrarily large generations.

Keywords: social networks, sequential social learning, Bayesian learning, confounding

*We thank Nageeb Ali, Simon Board, Aislinn Bohren, Matt Elliott, Mira Frick, Drew Fudenberg, Ben Golub, Sanjeev Goyal, Faruk Gul, Rachel Kranton, Suraj Malladi, Moritz Meyer-ter-Vehn, Pooya Molavi, Mallesh Pai, Andy Postlewaite, Matthew Rabin, Evan Sadler, Philipp Strack, Omer Tamuz, Fernando Vega-Redondo, Leeat Yariv, and various conference and seminar participants for helpful comments. Kevin He thanks the California Institute of Technology for hospitality when some of the work on this paper was completed.

[†]Boston University. Email: krishnadasaratha@gmail.com

[‡]University of Pennsylvania. Email: hesichao@gmail.com

1 Introduction

Consider an environment where people learn about an unknown state of the world through their private information and others’ past choices. There is a signal-extraction problem when people do not observe everyone’s actions, but only those of their neighbors in a network. For instance, suppose an agent observes multiple neighbors who have all been influenced by one person’s past choice. If the agent knows this person’s action, she would rationally “anti-imitate” it to subtract out its duplicate effect and infer her other neighbors’ private information (Eyster and Rabin, 2014). But when the agent does not observe her neighbors’ shared source, the observation network generates an obstruction to learning which we term *information confounding*: it is impossible to fully incorporate the private information of the neighbors without over-weighting the private information of their common influence.

This paper shows that the observation network can severely obstruct social learning through the channel of information confounding. We work with the canonical sequential social-learning model, which features a binary state, but make two assumptions to make our analysis tractable. First, we assume that agents have Gaussian private signals about this binary state. Second, we suppose that agents have sufficiently informative actions so that their behavior fully reveal their beliefs.¹ This *rich-signals, rich-actions* world removes some other obstructions to efficient learning² and isolates the effect of the social network. Our main results apply this framework to a class of networks where agents move in different generations and observe some or all of the previous generation. In these networks, we provide a closed-form expression for the efficiency of learning and describe how this rate depends on interpretable network parameters. A key implication of this expression is that when generations are large, information confounding makes learning arbitrarily slow.

To formalize these findings, we first describe several general properties of the social-learning model. The unique rational strategy profile of the social-learning game has a log-linear form. We characterize the strategy profile that solves agents’ signal-extraction problems and give a procedure to compute every agent’s accuracy in any network. In our model, it turns out the action of each rational agent is distributed as if she saw some (possibly non-integer) number of independent private signals. This signal-counting interpretation gives a simple measure of accuracy in the binary-state setting studied in Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and much of the subsequent sequential social-learning literature. We can quantify each agent’s learning outcome in any network in units

¹The simplest example is that agents choose actions equal to their posterior beliefs given their information. Our analysis also applies more generally to any decision problem where actions fully communicate beliefs.

²These obstructions are studied by Harel, Mossel, Strack, and Tamuz (2021), Molavi, Tahbaz-Salehi, and Jadbabaie (2018), Rosenberg and Vieille (2019), and others.

of private signals.

We demonstrate the power of network-based confounding with several examples in finite networks. The leading example considers a network where an agent has many neighbors who in turn share some common predecessors that the agent does not observe. We show that observing any number of neighbors who share just one common predecessor is less informative than observing four independent signals. So even a small amount of confounding can lead to arbitrarily large information loss.

Moving beyond finite networks, we define a measure of the rate of social learning on a network which we call *aggregative efficiency*. This measure relies on the signal-counting interpretation of accuracy to determine what fraction of private signals are incorporated into agents' beliefs and what fraction are effectively lost. For example, aggregative efficiency is $\frac{1}{3}$ if later agents learn about as well as if they saw $\frac{1}{3}$ of their predecessors' private signals (and no other information). Our main application computes the aggregative efficiency and hence quantifies the information loss due to confounding in a class of *generations networks*. Agents are arranged into generations of size K and each agent in generation t observes some subset of her generation $t - 1$ predecessors. This network structure could correspond to actual generations in families or countries, or successive cohorts in organizations like firms or universities. A broad insight is that a class of these networks cannot sustain much learning: even if generation sizes are large, later generations contribute little extra information.

We consider *symmetric* observation structures between generations: all agents observe the same number of neighbors and all pairs of distinct agents in the same generation share the same number of common neighbors. Society learns completely in the long run for every generation size, but this learning can be arbitrarily slow. No matter the size of the generations, social learning accumulates no more than two signals per generation asymptotically. Therefore, aggregative efficiency is arbitrarily close to zero when generations are large. A large number of endogenously correlated observations, such as the actions of all predecessors from the previous generation, can be less informative than a small number of independent signals. This conclusion holds even for networks where one's neighbors have large and almost non-overlapping observation sets, such as when they see many distinct predecessors and each pair of neighbors only shares one predecessor in common.

We can say more in the special case of *maximal generations networks* where each agent in generation t observes the actions of all predecessors in generation $t - 1$. Aggregative efficiency is worse with larger generation sizes, as illustrated in Figure 1. We also show that even early generations learn slowly in maximal generations networks. Social learning accumulates no more than three signals per generation starting with the third generation. If everyone in the first generation observes a single additional common ancestor, then the same bound also

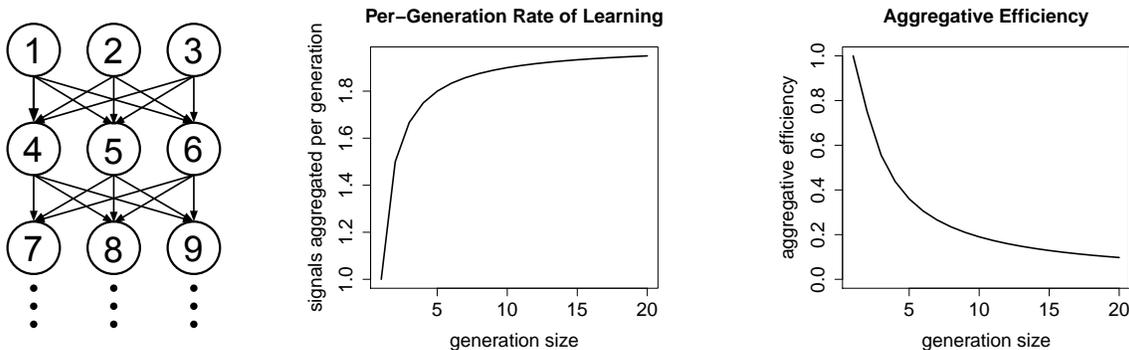


Figure 1: **Left:** A maximal generations network with a generation size of three. An arrow from i to j means j observes i 's action. **Middle:** Number of signals aggregated per generation asymptotically in maximal generations networks, as a function of generation size. **Right:** Aggregative efficiency in maximal generations networks, as a function of generation size.

holds for all generations.

We also compare rational social-learning dynamics across different symmetric generations networks. We derive a simple formula for aggregative efficiency as a function of the network parameters. This expression shows aggregative efficiency is higher when agents have more observations but lower when pairs of agents have more common neighbors, thus quantifying the trade-offs as we change the network. For instance, increasing the density of the observation network has two countervailing effects on learning: it speeds up the per-generation learning rate by adding more social observations, but also slow it down by lowering the informational content of each observation through extra confounding.

The next result connects aggregative efficiency to welfare. If signals are sufficiently precise or the planner is sufficiently impatient, welfare will depend primarily on the learning outcomes of the first few agents. We relate aggregative efficiency to welfare outside of these cases: networks with higher aggregative efficiency lead to higher welfare when signals are not too precise and the social welfare function is sufficiently patient. We also give an example showing the arbitrarily large information loss in generations networks can have arbitrarily large welfare consequences.

Finally, we discuss two economic applications of our results to organization structure. First, we demonstrate that opening up new channels of communication in an organization, such as starting a mentorship program where seniors share their private signals with newcomers, can have large benefits for organizational learning. We show that if mentors can share their private information with newcomers, this can fully correct large learning inefficiencies. Second, we show that *information silos* — partitioning some employees into insular groups that do not communicate with each other — improve executives' information aggregation at

the expense of workers’ learning. Asymmetric networks can improve learning in some parts of organizations by degrading learning in other parts, which may be beneficial if certain agents make particularly consequential decisions.

As we discuss in more detail in Section 7, a number of recent papers study other obstructions to rational social learning. Perhaps most relevant to the current work, several do allow network observations (including Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Rosenberg and Vieille (2019), and Dasaratha, Golub, and Hak (2023)). These papers show that in broad classes of sufficiently connected networks, long-run learning outcomes are determined by the precision and diversity of agents’ private information. Their results do not relate details of the observation network to learning outcomes (except by requiring enough links for learning to occur). For example, Rosenberg and Vieille (2019) argue that “the nature of the feedback on previous choices matters little”. We find that when the outcome of interest is instead short-run accuracy or the rate of learning, network structure can be a crucial obstruction.

2 Model

There are two equally likely states of the world, $\omega \in \{0, 1\}$. An infinite sequence of agents indexed by $i \in \mathbb{N}_+$ move in order, each acting once. (In some examples, we work with a finite subset $\{1, \dots, n\}$ of this infinite sequence.) On her turn, agent i observes a *private signal* $s_i \in \mathbb{R}$ and the actions of her *neighbors*, $N(i) \subseteq \{1, \dots, i - 1\}$. Agent i then chooses an *action* $a_i \in (0, 1)$ to maximize the expectation of $u_i(a_i, \omega) := -(a_i - \omega)^2$ given all of her information. So her action is equal to her belief about the probability that $\omega = 1$.³

We consider a Gaussian information structure where private signals (s_i) are conditionally i.i.d. given the state. We have $s_i \sim \mathcal{N}(1, \sigma^2)$ when $\omega = 1$ and $s_i \sim \mathcal{N}(-1, \sigma^2)$ when $\omega = 0$, where $\mathcal{N}(a, b^2)$ is the Gaussian distribution with mean a and variance b^2 , and $0 < 1/\sigma^2 < \infty$ is the private-signal precision.

Neighborhoods of different agents define a deterministic network M , where there is a directed link $j \rightarrow i$ if and only if $j \in N(i)$. We also view M as the adjacency matrix, with $M_{i,j} = 1$ if $j \in N(i)$ and $M_{i,j} = 0$ otherwise. Since $N(i) \subseteq \{1, \dots, i - 1\}$, M is upper triangular. The network M is common knowledge. The goal of this paper is to explore to what extent network structures can hinder efficient information aggregation.

With the network M fixed, let $d_i := |N(i)|$ denote the number of i ’s neighbors. A *strategy* for agent i is a function $A_i : (0, 1)^{d_i} \times \mathbb{R} \rightarrow (0, 1)$, where $A_i(a_{j(1)}, \dots, a_{j(d_i)}, s_i)$ specifies i ’s

³The quadratic-loss form of the utility functions is not crucial, and our results on learning remain unchanged if actions are “rich” enough to fully reflect beliefs (see Ali (2018b) for details).

play after observing actions $a_{j(1)}, \dots, a_{j(d_i)}$ from neighbors⁴ $N(i) = \{j(1), \dots, j(d_i)\}$ and when her own private signal is s_i .⁵ There is a unique strategy profile $(A_i^*)_{i \in \mathbb{N}_+}$ consistent with common knowledge of rationality at the interim stage: for all i and for all observations of i , A_i^* maximizes Bayesian expected utility given the posterior belief about ω .⁶ Uniqueness of this profile follows from the sequential nature of the social-learning game and the existence of a unique optimal action at any belief. Agent 1 has no social observations, so there is a unique rational strategy $A_1^*(s_1)$. This implies agent 2 also has a unique rational best response A_2^* , as we have fixed the behavior of a rational agent 1. Proceeding in this way, there is a unique strategy profile $(A_i^*)_{i \in \mathbb{N}_+}$ consistent with common knowledge of rationality at the interim stage, which we abbreviate as “rational.”

3 Basic Results about Beliefs and Learning

In this section, we show that rational actions are log-linear and satisfy a signal-counting interpretation. We then use these properties to demonstrate information confounding in several examples. The final subsection gives a condition for long-run learning and defines an asymptotic measure of how efficiently information is aggregated.

3.1 Optimal Actions Are Log-Linear

As is common in analyzing social-learning problems, we will find it convenient to work with the following log-transformations of variables: $\lambda_i := \ln\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right)$, $\ell_i := \ln\left(\frac{a_i}{1-a_i}\right)$. We call λ_i the *log-signal* of i and ℓ_i the *log-action* of i . These changes are bijective, so it is without loss to use the log versions. Write $\mathcal{L}_i^*(\ell_{j(1)}, \dots, \ell_{j(d_i)}, \lambda_i)$ for i 's rational log-strategy: the (unique) rational map from i 's neighbors' log-actions and i 's log-signal to i 's log-action. In this section, we show that every \mathcal{L}_i^* is a linear function of its arguments, with coefficients that only depend on the network M and not on the precision of private signals.

The following result shows the optimal aggregation is linear in log-signals and log-actions (*log-linear*) and gives an explicit expression for the coefficients. All proofs are in the Appendix.

⁴We use $j(k)$ to indicate the k -th neighbor of i and suppress the dependence of j on i when no confusion arises.

⁵It is without loss to focus on pure strategies, since every belief about the state induces a unique optimal action for each agent.

⁶We will see that in the rational strategy profile, $s_i \mapsto A_i^*(a_{j(1)}, \dots, a_{j(d_i)}, s_i)$ is a surjective function onto $(0, 1)$ for all i and $a_{j(1)}, \dots, a_{j(d_i)}$. So all observations are on-path and the posterior beliefs are well-defined.

Proposition 1. For each agent i with $N(i) = \{j(1), \dots, j(d_i)\}$, there exist constants $(\beta_{i,j(k)})_{k=1}^{d_i}$ so that i 's rational log-strategy is given by

$$\mathcal{L}_i^*(\ell_{j(1)}, \dots, \ell_{j(d_i)}, \lambda_i) = \lambda_i + \sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)}.$$

The vector of coefficients $\vec{\beta}_i$ is given by

$$\vec{\beta}_i = 2 \left(\mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)}) \mid \omega = 1] \times \text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} \mid \omega = 1]^{-1} \right),$$

where $\text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} \mid \omega = 1]^{-1}$ is the inverse of the conditional covariance matrix. These coefficients do not depend on the signal variance σ^2 .

For general private-signal distributions, models of Bayesian updating in networks have tractability issues, as Golub and Sadler (2016) point out. The key lemma to proving Proposition 1 shows that given our Gaussian information structure, agent i 's observations have a jointly Gaussian distribution conditional on ω . This permits us to study optimal inference in closed form. The interpretation of the inverse covariance matrix that appears in the coefficients $\vec{\beta}_i$ is that i rationally discounts the actions of two neighbors $j(1)$ and $j(2)$ if their actions are conditionally correlated.

3.2 Measure of Accuracy

We would like to evaluate networks in terms of their social-learning accuracy so that we can compare the rates of Bayesian learning in different networks. Towards a measure of accuracy, imagine that agent i 's only information about ω consists of $n \in \mathbb{N}_+$ conditionally independent private signals. Then, the Bayesian i would play the log-action equal to the sum of the n log-signals, and it turns out (by Lemma A.1 in the Appendix) her behavior would follow the conditional distributions $\ell_i \sim \mathcal{N}\left(\pm n \cdot \frac{2}{\sigma^2}, n \cdot \frac{4}{\sigma^2}\right)$, with the positive and negative means in states $\omega = 1$ and $\omega = 0$ respectively. We quantify learning accuracy using distributions of this form that allow for non-integer n , thus denominating accuracy in the units of private signals.

Definition 1. Social learning aggregates $r \in \mathbb{R}_+$ signals by agent i if the rational log-action ℓ_i has the conditional distributions $\mathcal{N}\left(\pm r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2}\right)$ in the two states. If this holds for some $r \in \mathbb{R}_+$, then we say i 's behavior has a *signal-counting interpretation*.

When agents use an arbitrary strategy profile, the conditional distributions of ℓ_i need not equal $\mathcal{N}\left(\pm r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2}\right)$ for any r , even when the strategy profile is log-linear (i.e., each

agent’s log-action is a linear function of her log-signal and neighbors’ log-actions). Indeed, if this profile results in i putting certain weights $(w_{i,j})_{j \leq i}$ on log-signals $(\lambda_j)_{j \leq i}$, then ℓ_i has a signal-counting interpretation if and only if $\sum_{j=1}^i w_{i,j} = \sum_{j=1}^i w_{i,j}^2$.

But as the next result shows, the rational log-actions always admit a signal-counting interpretation in any network.

Proposition 2. *There exist $(r_i)_{i \geq 1}$ so that social learning aggregates r_i signals by agent i . These $(r_i)_{i \geq 1}$ depend on the network M , but not on private-signal precision.*

The signal-counting interpretation gives a way to compare agents’ accuracy and welfare across different networks or positions in a given network in a binary-state setting. Rather than comparing the full distributions of beliefs, we can compare the summary statistics r_i . A consequence is that agents’ beliefs, which *a priori* are multi-dimensional objects, are in fact ranked in the standard Blackwell ordering: a higher value of r_i implies a weakly higher expected utility for any decision problem.

Such comparisons are straightforward in a different framework with a Gaussian state and Gaussian signals (e.g., [Morris and Shin \(2002\)](#) and, in the context of social learning, [Dasaratha, Golub, and Hak \(2023\)](#)). In these models, Bayesian agents’ beliefs are ranked by their precisions and the analogous number of signals aggregated is simply proportional to precision. We use Proposition 2 to study the binary-state model used in most of the sequential social-learning literature (beginning with [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#)). But we expect information confounding to have similar effects, including severely obstructing learning in some networks, with a Gaussian state.⁷

The signal-counting interpretation of behavior is closely identified with the rational learning rule. Indeed, a rational agent’s behavior always admits a signal-counting interpretation even when her predecessors use arbitrary non-rational log-linear strategies.

Corollary 1. *Fix arbitrary log-linear strategies for agents $i < I$. If agent I best responds to these strategies, then I ’s behavior has a signal-counting interpretation.*

For a rational agent, Definition 1 gives a summary statistic for the accuracy of her beliefs and her utility—even if her observations are not generated by rational behavior. If an agent is not updating beliefs or choosing actions rationally, however, her utility need not be determined by such a summary statistic and can depend on a more complex action distribution.

⁷Indeed, one can show the example networks in Section 3.3 give the same results in either model.

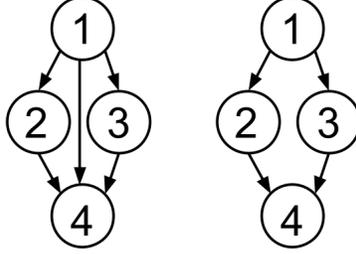


Figure 2: **Left:** The shield network with four agents. **Right:** The diamond network with four agents.

3.3 Examples of Information Confounding in Networks

We say the network causes *information confounding* if there are multiple paths between i and j in the network but the later agent j does not observe i directly. We begin with several examples of social learning in finite networks that cause information confounding.⁸ These examples illustrate how information confounding obstructs social learning, highlight the extent of possible information loss due to confounding, and provide intuition for our main results on generations networks.

Example 1 (The Shield and the Diamond). Consider two network structures with four agents, as shown in Figure 2. In a *shield network*, agent 4 observes agents 1, 2, and 3 while agents 2 and 3 observe agent 1. In a *diamond network*, agent 4 observes agents 2 and 3 while agents 2 and 3 observe agent 1.⁹

In a shield network, agent 4 observes all predecessors and can compute the private signals of all agents. To see this, note that $\ell_1 = \lambda_1$, $\ell_2 = \lambda_1 + \lambda_2$, and $\ell_3 = \lambda_1 + \lambda_3$. So

$$\ell_4 = \lambda_4 + \ell_2 + \ell_3 - \ell_1 = \sum_{j=1}^4 \lambda_j$$

is the optimal action given her private signal and those of her three predecessors, and $r_4 = 4$. As in [Eyster and Rabin \(2014\)](#), the optimal action involves anti-imitating, or placing a negative weight on, agent 1's action.

In a diamond network, however, agent 4 observes the actions of agents 2 and 3 that combine their private signals with agent 1's signal, which agent 4 does not observe. Agent 4 faces an unavoidable tradeoff between overweighting agent 1's signal and underweighting agents 2 and 3's signals. As in the shield network, we have $\ell_1 = \lambda_1$, $\ell_2 = \lambda_1 + \lambda_2$, and

⁸While our model deals with an infinite sequence of agents, we can apply our model to settings with finitely many agents by only looking at the learning of the first n agents.

⁹Our terminology follows [Eyster and Rabin \(2014\)](#), who focus on rational learning in networks without diamonds.

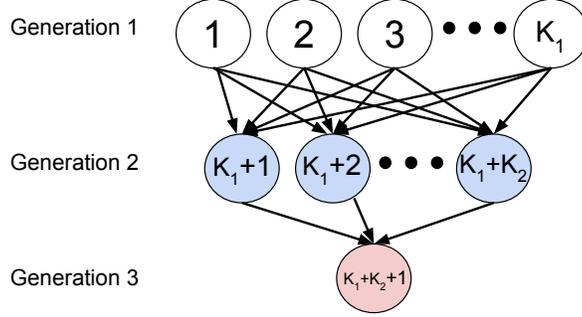


Figure 3: A three-generation network with K_1 agents in generation 1, K_2 agents in generation 2, and one agent in generation 3.

$\ell_3 = \lambda_1 + \lambda_3$. Using Proposition 1, we can calculate that now

$$\ell_4 = \lambda_4 + \frac{2}{3}\ell_2 + \frac{2}{3}\ell_3 = \frac{4}{3}\lambda_1 + \frac{2}{3}\lambda_2 + \frac{2}{3}\lambda_3 + \lambda_4,$$

and therefore $r_4 = \frac{11}{3} < 4$. Even though agent 4 is Bayesian and optimally adjusts for the confounding signal that she does not observe, some information is lost. This information loss is not too severe here, but the next example shows it can be much worse with more agents.

Example 2. To see that confounding can lead to more severe information loss, we next generalize the diamond network to allow more agents. Consider a network with agents in three generations, shown in Figure 3.¹⁰ In the first generation, agents $1, 2, \dots, K_1$ have no neighbors. In the second generation, agents $K_1 + 1, K_1 + 2, \dots, K_1 + K_2$ observe all agents in the first generation. Finally, the third generation consists of a single agent who observes all agents in the second generation but does not observe the first generation. The purpose of this example is to study the beliefs of an agent with many neighbors who all observe a common confound.¹¹

The agent in generation 3 rationally calculates the log-likelihood of state $\omega = 1$ by taking a weighted sum of the log-actions of generation 2 agents and her own signal, where the weights depend on K_1 and K_2 . As in Example 1, the final agent faces an unavoidable tradeoff between overweighting generation 1's private signals and underweighting generation 2's private signals. Using Proposition 1, we can compute that the optimal action places weight $\frac{1+K_1}{1+K_1K_2}$ on each generation 2 action (see Appendix A.1 for details of the calculations in this example). When generation 2 is large, this weight is close to $1/K_2$: it is optimal for the final agent to severely underweight the private signals of generation 2.

¹⁰In Section 4 we will study generations of equal size.

¹¹We could equivalently relax the assumption of i.i.d. signals and replace the first generation with a single agent with a (potentially) more precise signal.

We can also show that the actions of the agents in generation 2 are distributed as if they see $1 + K_1$ conditionally independent private signals, while the action of the final agent is distributed as if she sees $1 + \frac{K_2 + K_1 K_2}{1 + K_1 K_2} \cdot (1 + K_1)$ such signals. The difference between the accuracy of generation 2 and 3's actions is just

$$1 + \frac{K_2 + K_1 K_2}{1 + K_1 K_2} \cdot (1 + K_1) - (1 + K_1) = 1 + \frac{(K_2 - 1)(K_1 + 1)}{K_1 K_2 + 1} < 3$$

private signals, for any values of K_1 and K_2 . So there is always very little learning between generations 2 and 3, even when the size of generation 2 is large and many private signals arrive in that generation. The idea is that confounding significantly limits how much information the final agent can extract from arbitrarily many neighbors' actions.

We emphasize that the first generation can generate substantial confounding even when it is small. For example, if there is a single agent in the first generation ($K_1 = 1$), then the action of the agent in generation 3 will be less accurate than that of someone who saw just five independent private signals. But if generation 1 were empty, then the action of the generation 3 agent would be equivalent to $K_2 + 1$ private signals. So even a small confound can prevent almost all information aggregation. Also, a simple calculation shows that the difference between the accuracy of generation 2 and generation 3 strictly decreases in K_1 and strictly increases in K_2 , provided $K_2 \geq 2$. That is, the incremental amount of learning in the final generation decreases with confounding (a larger K_1) but increases with the number of observations (a larger K_2). We will later see that the same comparative statics hold for a class of infinite networks where agents move in generations.

3.4 Long-Run Learning and Aggregative Efficiency

We now return to studying infinite networks. We begin with a benchmark result providing necessary and sufficient conditions for long-run learning. These conditions turn out to be similar to those in the existing literature, which shows our model is comparable to the standard models on this dimension. A key contribution of our model is ranking networks where agents learn in the long run based on the rate of this learning, and this section concludes by defining a measure that will provide such a ranking for a class of generation networks.

We say society *learns completely in the long run* if (a_i) converges to ω in probability. For a given network M , write $\mathcal{PL}(i) \in \mathbb{N}$ to refer to the length of the longest path in M originating from i (this length is 0 if $N(i) = \emptyset$).

Proposition 3. *The following are equivalent: (1) $\lim_{i \rightarrow \infty} \mathcal{PL}(i) = \infty$; (2) $\lim_{i \rightarrow \infty} [\max_{j \in N(i)} j] = \infty$; (3) $\lim_{i \rightarrow \infty} r_i = \infty$; (4) society learns completely in the long run.*

Condition (2) is the analog of [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#)’s *expanding observations* property for a deterministic network. It says if we consider the most recent neighbor observed by each agent, then this sequence of most recent neighbors tends to infinity. It is straightforward to see the expanding observations condition is necessary for long-run learning, and [Acemoglu et al. \(2011\)](#) show it is also sufficient in a random-networks model with unboundedly informative signals and binary actions.¹² With continuous actions, [Proposition 3](#) states the same result. The intuition is that each agent learns at least as well as if she optimally combined her most accurate social observation with her private signal.

The key takeaway message from [Proposition 3](#) is that whether society learns in the long run is not a useful criterion for comparing different networks in this setting, as the conditions (1) and (2) that guarantee long-run learning are quite mild. It is of course possible that agents learn completely but do so very slowly. We now define a measure of the efficiency of learning, which can evaluate learning outcomes when this occurs:

Definition 2. If $\lim_{i \rightarrow \infty} (r_i/i)$ exists, it is called the *aggregative efficiency* of the network.

Aggregative efficiency measures the fraction of signals in the entire society that individuals manage to aggregate under social learning. Networks with higher levels of aggregative efficiency induce faster social learning in the long run. The limit defining aggregative efficiency need not exist in all networks, but does exist in almost all of the examples we focus on.

The next section will compare how quickly $(r_i)_{i \geq 1}$ grows across a class of generations networks and the aggregative efficiency of these networks. Comparisons of aggregative efficiency also translate into welfare comparisons, as [Section 5](#) will show.

4 Generations Networks

This section shows that information confounding can lead to arbitrarily large information losses and derives a closed-form expression for how confounding influences learning in a class of networks. We study *generations networks*¹³ and find that they can lead to very inefficient learning due to confounding, even when one’s network neighbors observe almost disjoint sets of predecessors. We also compare aggregative efficiency across these networks.

Agents are sequentially arranged into generations of size K , with agents within each generation placed into *positions* 1 through K . Agents in the first generation (i.e., $i = 1, \dots, K$)

¹²With boundedly informative signals and binary actions, however, long run learning fails (see also [Smith and Sørensen \(2000\)](#)).

¹³This class of networks follows a strand of social-learning literature where agents move in generations, for instance [Wolitzky \(2018\)](#), [Banerjee and Fudenberg \(2004\)](#), [Burguet and Vives \(2000\)](#), and [Dasaratha, Golub, and Hak \(2023\)](#).

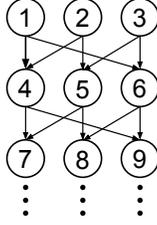


Figure 4: A generations network with $K = 3$ agents per generation and the observation sets $\Psi_1 = \{1, 2\}$, $\Psi_2 = \{2, 3\}$, and $\Psi_3 = \{1, 3\}$.

have no neighbors. A collection of *observation sets* $\Psi_k \subseteq \{1, \dots, K\}$ for $k = 1, \dots, K$ define the network M among the agents. The agent in position k in generation $t \geq 2$ observes agents in positions Ψ_k from generation $t - 1$ (and no agents from any other generation). That is, for $i = (t - 1)K + k$ where $t \geq 2$ and $1 \leq k \leq K$, network M has $N(i) = \{(t - 2)K + \psi : \psi \in \Psi_k\}$.¹⁴ Figure 4 shows an example with $K = 3$.

We focus on observation sets $(\Psi_k)_k$ satisfying a symmetry condition:

Definition 3. The observation sets are *symmetric* if all agents observe $d \geq 1$ neighbors and all pairs of agents in the same generation share c common neighbors, i.e. $|\Psi_k| = d$ for every $1 \leq k \leq K$ and $|\Psi_{k_1} \cap \Psi_{k_2}| = c$ for distinct positions $1 \leq k_1 < k_2 \leq K$.

A generations network defined by symmetric observation sets is called a *symmetric network*. To give a class of examples of symmetric networks, fix any non-empty subset $E \subseteq \{1, \dots, K\}$, and let $(\Psi_k)_k$ be such that for all $1 \leq k \leq K$, $\Psi_k = E$. Here we have $d = c = |E|$. To interpret, the set E represents the prominent positions in the society, and agents only observe predecessors in these prominent positions from the past generation. We call the special case of $E = \{1, \dots, K\}$ the *maximal generations network*, where agents in generation t for $t \geq 2$ have all agents in generation $t - 1$ as their neighbors. For another class of examples, suppose $K \geq 2$ and each agent observes a different subset of $K - 1$ predecessors from the previous generation. Specifically, $\Psi_k = \{1, \dots, K\} \setminus \{k - 1\}$ for $2 \leq k \leq K$, and $\Psi_1 = \{1, \dots, K - 1\}$. This network is symmetric with $d = K - 1$ and $c = K - 2$. (The network in Figure 4 has this structure, with $d = 2$ and $c = 1$.) There are also a large variety of other symmetric networks that are not covered by these two classes of examples: one enumeration shows there are at least 103 pairs of feasible (d, c) parameters in the range of $3 \leq d \leq 41$ and $1 \leq c \leq d - 2$ that correspond to at least one symmetric network, typically with multiple non-isomorphic networks for each feasible parameter pair (Mathon and Rosa, 1985).

¹⁴Stolarczyk, Bhardwaj, Bassler, Ma, and Josić (2017) study a related model where only the first generation observes private signals. Their main results characterize when no information gets lost between generations, i.e., social learning is completely efficient.

4.1 Aggregative Efficiency in Symmetric Generations Networks

We provide an exact expression for the aggregative efficiency in symmetric generations networks to quantify the information loss due to confounding.

Theorem 1. *Given any symmetric observation sets $(\Psi_k)_k$ where every agent observes d neighbors and every pair of agents in the same generation share c common neighbors, aggregative efficiency is¹⁵*

$$\lim_{i \rightarrow \infty} (r_i/i) = \left(1 + \frac{d^2 - d}{d^2 - d + c}\right) \frac{1}{K}$$

The number of signals aggregated per generation asymptotically ($\lim_{i \rightarrow \infty} r_{i+K} - r_i$) is less than 2. For $c \geq 1$, this number is strictly increasing in d and strictly decreasing in c .

Theorem 1 calculates the aggregative efficiency in any symmetric generations network in terms of the parameters d and c . The expression on the right-hand side normalizes by the size of the generation K , so the term in the parentheses provides a uniform learning-rate upper bound of two signals per generation across all symmetric networks (as $\frac{d^2-d}{d^2-d+c} \leq 1$).

The interpretation of the comparative statics result in d and c is that more observations speed up the rate of learning per generation but more common neighbors slow it down by worsening confounding, all else equal. This result lets us compare learning dynamics across different symmetric networks characterized by different (d, c) parameter pairs. Changing from one network to another, we can change both d and c (along with the generation size K). Theorem 1 decomposes the repercussions of such changes on the *per-generation* learning rate (i.e., after normalizing by $1/K$) into their effects on the two primitive network parameters that have monotonic influences on said rate.

The main content of the theorem is the uniform bound on the learning rate, which implies learning is very inefficient in large symmetric generations networks. The proof of Proposition 3 provides a lower bound of one signal aggregated per generation, since agents could always optimally combine their private signal with one observed action. Theorem 1 shows this lower bound is not too far from the actual learning rate, which is fewer than two signals per generation.

For maximal generations networks (i.e., agents observe all predecessors from the previous generation), the basic intuition for this bound is similar to Example 2, which tells us that when any number of agents observe one or more common signals in addition to their private signals, a successor who observes all of these agents cannot improve on their accuracy by more than three signals worth of information. The successor must balance overweighting

¹⁵In the case $d = 1$ and $c = 0$, we adopt the convention $0/0 = 0$.

the common confound and underweighting her neighbors' private signals, and the optimal weights severely underweight recent signals.

Extending this intuition beyond maximal generations networks is more subtle, because different agents in a generation may observe different predecessors whose actions may be less correlated. This can alleviate the information confounding in early generations, but we show the benefits are limited: even if agents in the same generation have almost disjoint observation sets, actions become highly correlated in later generations. To prove this, we use a mixing argument to show that the actions of two agents in the same generation are influenced in very similar ways by the signal realizations of their common ancestors from many generations ago. An implication is that recent signals are severely underweighted, as in the maximal generations case: the *total* weight an agent places on private signals from the previous generation converges to one, while in the absence of network-based confounds the agent would place a weight of one on *each* signal.

Perhaps surprisingly, an implication of Theorem 1 is that aggregative efficiency in symmetric networks only depends on the generation size. Compare the symmetric network from Figure 4 with $d = 2, c = 1, K = 3$ with the maximal generations network with $K = 3$. Theorem 1 implies they have the same aggregative efficiency. The extra social observations in the second network exactly cancel out the reduced informational content of each observation, due to the more severe information confounds. Our next result shows that more generally, any symmetric network with parameters (d, c, K) where $d \geq 2, c < d$ has the same aggregative efficiency as the maximal generations network with the same generation size K .

Corollary 2. *In any symmetric network with $d \geq 2$, aggregative efficiency is $\lim_{i \rightarrow \infty} (r_i/i) = (2 - (1/K)) \cdot \frac{1}{K}$.*

This corollary follows from the fact that the symmetry condition imposes some combinatorial constraints on the feasible (d, c, K) parameter triplets. It turns out these constraints allow us to simplify the expression in Theorem 1 when we know the generation size. While Corollary 2 gives a simple expression of aggregative efficiency that just depends on K , Theorem 1 lets us compare networks that differ in d and c (and possibly also K) more transparently.

In particular, these results imply that the aggregative efficiency of maximal generations network with K agents per generation is $\lim_{i \rightarrow \infty} (r_i/i) = \frac{(2K-1)}{K^2}$, which decreases with K . Indeed, if $K = 1$, then every agent perfectly incorporates all past private signals and the speed of social learning is the highest possible. Not only does this result about the aggregative efficiency imply that asymptotically fewer signals are aggregated by the same agent in networks with larger K , but the same comparative statics also hold numerically for all agents $i \geq 16$ when comparing among $K \in \{2, 3, 4, 5\}$, as shown in Figure 5.

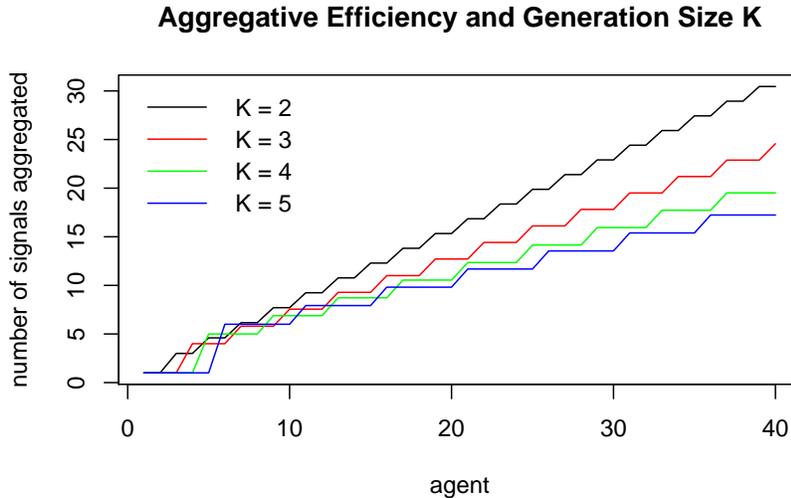


Figure 5: Number of signals aggregated by social learning in maximal generations networks with different generation sizes, $K \in \{2, 3, 4, 5\}$.

We next discuss the role of the symmetric generations structure in the bound of two additional signals per generation. We have implicitly imposed three restrictions on the network, beyond the basic generations structure: the observation structures are the same across generations, all generations are the same size, and observations are symmetric within generations. We assume the same observation structure across generations primarily for simplicity of exposition, and could bound aggregative efficiency with the same techniques while allowing different symmetric observation structures in different generations. The key step in extending the proof is the Markov chain mixing result, which must be replaced by a mixing result for non-homogeneous Markov chains (for examples of such results, see [Blackwell \(1945\)](#) and [Tahbaz-Salehi and Jadbabaie \(2006\)](#)). We could also allow different sized generations and obtain bounds on aggregative efficiency. For example, the logic of [Example 2](#) would extend to maximal generation networks with generations of varying sizes.

The assumption of symmetry within generations is more substantive, and our bound of two signals of additional information per generation does not always apply to non-symmetric generations networks. For example, suppose $\Psi_k = \{k\}$ for some position k . The accuracy of agents in this position will only increase by one additional signal per generation, but these agents can aggregate independent information that can be valuable to agents in other positions (see [Section 6.2](#) for details). An unresolved question is whether asymmetric observation structures could let *all* agents aggregate more than two signals per generation.

4.2 Finite-Population Learning

Our framework not only allows us to study the asymptotic rate of learning, but also lets us derive finite-population bounds that apply from early generations. Theorem 1 tells us social learning aggregates fewer than two signals per generation asymptotically. There is also a short-run version of this result in the maximal generations network: starting with generation 3, fewer than three signals are aggregated per generation for any K .

Proposition 4. *In any maximal generations network, for any agents i, i' in generation t and $t - 1$ with $t \geq 3$, $r_i - r_{i'} \leq 3$.*

We find an even starker bound on r_i if we consider a modified version of the maximal generations network: there is a zeroth generation with only one agent, and all subsequent generations contain K agents each. Agents in generation $t \geq 1$ observe all predecessors from generation $t - 1$.

Proposition 5. *In this modified maximal generations network, $r_i \leq 3t - 1$ for every i in every generation $t \geq 1$.*

Similar to Example 2, the single agent before the first generation causes significant information confounding. With this additional agent, there is a uniform bound on every generation's accuracy across all generation sizes K .

5 Aggregative Efficiency and Welfare

In this section, we relate aggregative efficiency comparisons to welfare comparisons. When signals are precise enough for agents to learn well, welfare will depend largely on learning outcomes of finite networks (such as the examples in Section 3.3) rather than asymptotic quantities. Higher aggregative efficiency implies higher welfare, however, when signals are sufficiently imprecise.

Let $v_i^M := \mathbb{E}[u_i(a_i^*, \omega)]$ denote the expected welfare of agent i in network M , and recall that $-0.25 < v_i^M < 0$ for every i in any network and with any private signal precision $0 < 1/\sigma^2 < \infty$. It turns out that whenever the aggregative efficiency of a network M is strictly positive, $v_i^M \rightarrow 0$ and this convergence happens at an exponential rate. This implies the undiscounted sum of utilities of all agents, $\sum_i v_i^M$, is convergent.

We show that if two networks are ranked by aggregative efficiency, then the undiscounted sums of all agents' utilities on these networks follow the same ranking whenever private signals are sufficiently imprecise. The same result also applies to the discounted sums of

utilities, provided the discounting does not weigh the welfare of the earliest agents too heavily.

Proposition 6. *Suppose networks M and M' have aggregative efficiencies $AE_M > AE_{M'} > 0$. Then there exists some $\underline{\sigma}^2 > 0$ such that for any signal variance $\sigma^2 \geq \underline{\sigma}^2$, we have $\sum_i \delta^{i-1} v_i^M > \sum_i \delta^{i-1} v_i^{M'}$ for δ sufficiently close to or equal to 1.*

This result provides a foundation for the aggregative efficiency measure in terms of a conventional social welfare function: the (un)discounted sum of utilities. The result applies to arbitrary networks, and does not require the generations structure from Section 4.

The arbitrarily large information loss we have highlighted in Section 4 can have large welfare consequences. To illustrate this, we give an example comparing complete networks to maximal generations networks with large generations.

Example 3. Let M be the maximal generations network with generation size K . We will let K grow large and set signal variance $\sigma^2 = \sigma_0^2/K$ for a constant $\sigma_0^2 > 0$. That is, we increase the generation size but fix the total informativeness of a single generation's private signals. Let M' be the complete network (or any other network with $r_i = i$ for all i). Then

$$\lim_{K \rightarrow \infty} \frac{\sum_i v_i^M}{\sum_i v_i^{M'}} = \infty.$$

(We provide an argument in Appendix A.) Since utilities are negative for all agents, the limit implies that the total disutility in the maximal generations network is unboundedly larger than when agents can extract all previous private signals.

6 Organizational Applications

6.1 Value of Mentorship

We provide an economic application of our results in terms of the value of mentors who share their private signals with mentees in the next generation.

Many organizations with cohort structures, such as universities and firms, have mentorship programs that pair newcomers with members of a previous cohort. Our results suggest that one benefit of such programs is that mentors provide information that helps newcomers interpret others' actions, thus increasing the speed of learning within the organization.

Formally, we model a mentor as someone who shares her private signal with a mentee in the subsequent generation. Equivalently, the mentor could share a sufficient statistic describing her best estimate of the state based on her social observations. If we begin with

the maximal generations network and add mentorship relationships in this way, learning is nearly efficient.

Corollary 3. *Suppose each agent observes the actions of all members of the previous generation and the private signal of one member of the previous generation. Then social learning aggregates more than $i - K$ signals by every agent i , so aggregative efficiency is 1.*

If an agent observes the actions of the previous generation along with one of their private signals, she can calculate the common confounding information and fully compensate for this confound. In networks with large K , showing each agent just one extra signal (of someone from the previous generation) increases aggregative efficiency from nearly 0 to 1.

In the context of the application, incumbents in the organization act based on private information and shared organizational knowledge. A newcomer ignorant of the organizational knowledge cannot fully separate these two forces that shape others' behavior. But by describing her perspective, a mentor can help a newcomer interpret everyone else's behavior. This removes the informational confound facing the newcomer and lets her extract the private information underlying these predecessors' actions. A related force is described in management literature:

“Mentors can be powerful socializing agents as an individual adjusts to a new job or organization. As protégés learn about their roles within the organization, mentors can help them correctly interpret their experiences within the organization's expectations and culture.” – Chao (2007)

Our result formalizes this intuition in a social-learning environment. Our stylized model of mentorship abstracts away from many of its other benefits (e.g., the expertise of the mentor in terms of being able to generate more precise signals than the mentee), and shows how the “interpretive” value of mentorship improves learning within the organization.

If each mentor instead generates a new, independent private signal for their mentee, rather than sharing the realization of their private signal from the past, then social learning does not speed up very much. Compared to a world without mentoring, this intervention would at most double the number of signals aggregated by each agent. Using Theorem 1, this limits the organization to eventually aggregating at most four signals per cohort. In organizations with large cohorts, mentors who share their personal experience increase the rate of social learning much more than mentors who generate new signals. The comparison between the signal-sharing mentors and the signal-generating mentors shows that Corollary 3 relies critically on the “interpretative” channel of mentoring: almost all of the additional learning under mentoring comes not from the mentor giving the mentee an extra signal about

the state, but from the mentee using the mentor’s past experience to clarify other people’s behavior and to extract substantially more information from said behavior.

6.2 Information Silos

Within some organizations, information is fragmented among various subgroups (departments, product divisions, trading desks) that fail to communicate with each other, creating *information silos*.¹⁶ These silos have a number of causes: compensation structures that discourage collaboration between teams, different subunits storing information in mutually incompatible databases, or technical language barriers that stop ideas from flowing between specialties.¹⁷ Tett (2015) documents the prevalence of information silos in government bureaucracies, technology firms, and banks, noting that many of these silos persist for many decades. She joins a number of other authors and management consultants in arguing that information silos are a necessary evil for running a complex workforce, but they hurt the organization by obstructing internal information exchange.¹⁸

We use a generations network to show that information silos may benefit the organization compared with fully transparent data sharing,¹⁹ when the organization’s success primarily depends on the actions of a few executives who can observe and process the behavior in all the silos. Consider a generations network with $K \geq 2$ agents per generation. Suppose positions $\{2, \dots, K\}$ are partitioned into $N \geq 1$ silos S_1, \dots, S_N so that each position k only observe predecessors in the same silo, $\Psi_k = S_n$ for $k \in S_n$, while agents in the first position can observe all of the silos, $\Psi_1 = \{2, \dots, K\}$.

Corollary 4. *Agents in the first position eventually aggregate $\lim_{t \rightarrow \infty} \frac{r^{(t-1)K+1}}{t} = \sum_{n=1}^N \frac{2|S_n|-1}{|S_n|}$ signals per generation. Agents in position $k \in S_n$ in silo n eventually aggregate $\lim_{t \rightarrow \infty} \frac{r^{(t-1)K+k}}{t} = \frac{2|S_n|-1}{|S_n|}$ signals per generation.*

The agents in positions S_n form an information silo for each $1 \leq n \leq N$. As a new cohort of workers join the organization, each newcomer learns by observing their seniors

¹⁶We thank Suraj Malladi for suggesting this application.

¹⁷Sethi and Yildiz (2016) show that a silo-like information segregation may become endogenously entrenched in an organization as each agent learns the subjective perspectives of the people she talks to most often. This encourages the agent to keep consulting the same people’s opinions in the future, as she can better account for their subjective biases and extract more precise information from their opinions.

¹⁸Arguments in favor of eliminating information silos are common in the popular press: see for example, Gleeson and Rozo (2013) and Casciaro, Edmondson, and Jang (2019).

¹⁹Similar network structures can also improve social learning in experimentation settings. In a model where a sequence of short-lived, behavioral agents take turns interacting with a multi-armed bandit, Immorlica, Mao, Slivkins, and Wu (2020) show that an observation structure featuring many information silos ensures at least one silo produces a large amount of information about the payoffs of each bandit arm, thus improving the welfare of later agents who observe all the information generated in every silo.

from the same silo, and information does not flow across different silos. Agents in position 1 are executives who observe all predecessors in silos. Corollary 4 shows that executives can aggregate up to $K - 1$ signals per generation, depending on the sizes of the silos. Figure 6 shows an example with two information silos that contain one and two agents respectively in each generation. Social learning can aggregate more than three signals per generation for the executives when there are silos. By contrast, the executives' information improves by no more than three signals per generation starting from generation 3 in the maximal generations network (Proposition 4), which represents an organization with full data transparency.

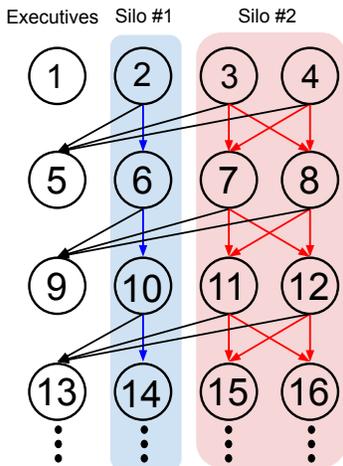


Figure 6: A generations network with executives in the first position and two information silos, $S_1 = \{2\}$ and $S_2 = \{3, 4\}$.

If the organization's payoff is closely identified with the utility of its executive's action in each generation, then information silos can improve the organization's welfare. Such an organization structure provides less confounded information to the key decision-makers by sacrificing the rate of learning within silos. Indeed, behavior in different silos are conditionally independent of each other. If the organization is instead one where every member's action significantly contributes to its welfare, then information silos are detrimental to the organization. Newcomers could learn better by observing all incumbents in the organization, instead of only those in the same silo.

The negative case studies that [Tett \(2015\)](#) and others use to advocate breaking down silos mostly involve workers in silos who take actions that severely harm the company, or executives who are unable to process the data from multiple silos. For instance, [Tett \(2015\)](#) discusses two product divisions of Sony simultaneously producing two very similar music players that ended up competing with each other on the market, a situation where the organization's welfare depends on the actions taken within the silos rather than the action of a single executive who oversees all silos.

An important qualification is that if employees communicate private signals along with actions (perhaps as in the application in Section 6.1), then information silos will be harmful for social learning. When full information sharing is possible, information silos will lead to less informed workers without meaningfully improving executives’ actions.

7 Related Literature

We study rational social learning in a sequential model (as first introduced by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)) where agents only observe some predecessors. Our work contributes to the social learning literature by quantifying how parameters of the network structure affect the efficiency of social learning through the information confounding channel. This leads us to the new conclusion that a small amount of confounding can generate arbitrarily inefficient social learning, even when agents perfectly observe their neighbors’ beliefs.

Our paper continues a literature on sequential learning when private signals can generate unboundedly strong beliefs. Smith and Sørensen (2000) show that there is complete long-run learning on the complete network with such signals, and Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015) extend this result to all networks satisfying weak sufficient conditions (see Proposition 3). Rosenberg and Vieille (2019) consider several networks without information confounding²⁰ and reach a similar conclusion that “the nature of the feedback on previous choices matters little” (under a different criterion for good learning). Indeed, the networks that Rosenberg and Vieille (2019) study would all have the same aggregative efficiency in our setting. Our results instead show that networks with different levels of information confounding can lead to substantial variations in short-run accuracy and rates of learning.²¹

The most closely related network-based obstructions to learning appear in Eyster and Rabin (2014), who mention the possibility of such confounds but restrict their analysis to networks where rational agents can fully correct for correlations in observations via anti-imitation. They note that relaxing this restriction to allow confounds would lead to “distributional complications”; our framework and results resolve these complications and study

²⁰In Rosenberg and Vieille (2019), there is at most one path in the observation network between i and any predecessor j whom i does not directly observe. This rules out, for example, i observing two agents who both observe j . The type of confounding we focus on arises when i overweights j ’s private signal because there are multiple paths between i and j in the network.

²¹A precedent to our comparison of learning rates across networks is Lobel, Acemoglu, Dahleh, and Ozdaglar (2009), who examine two particular networks, both involving each agent seeing exactly one neighbor. Our results allow us to compare networks that vary along richer dimensions, including the number of neighbors that agents have.

the implications of the confounds. Related obstructions are also present in [Dasaratha, Golub, and Hak \(2023\)](#), which studies learning failures in network structures similar to our generations networks but has no formal results about how learning differs across networks. They focus instead on how private signal precisions and the evolution of the state, which changes over time, affect learning. Indeed, they argue that network structure matters much less than the state and information structures in their setting. By contrast, we show that in a standard fixed-state environment learning can be quite efficient on some networks and highly confounded on others. Variations in the network structure can trace out a wide range of learning efficiencies, including nearly total information loss, which highlights the power of the confounding.

The previous two paragraphs discussed settings with unboundedly informative signals, but sequential models with boundedly informative signals can provide an alternative setting for asking how network structure affects learning. Complete long-run learning fails when signals are boundedly informative ([Smith and Sørensen, 2000](#)), and a largely open question is how the probability of failures depend on the network structure. Several papers show that incomplete networks where some agents do not observe others lead to better long-run outcomes than the complete network ([SgROI, 2002](#); [Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011](#); [Arieli and Mueller-Frank, 2019](#)). These papers each compare a specific class of incomplete networks with the complete network, but they do not allow comparisons of different incomplete networks (except through numerical simulations). We show a framework with Gaussian signals has useful properties such as log-linear actions and the signal-counting measure of accuracy and do not consider bounded signals, but we think further analytic results on the role of network structure in settings with bounded signals could be an interesting direction for future work.

In a paper combining diffusion and social learning literatures, [Board and Meyer-ter Vehn \(2021\)](#) study a product adoption model where agents arrive at random times, observe network neighbors' adoption choices, and can pay for a fully revealing private signal. Networks matter through different channels in their setting than the confounding mechanism that we focus on: indeed, networks that cause information confounding do not appear (or have vanishing probability) in [Board and Meyer-ter Vehn \(2021\)](#). The main force is instead that agents infer from the absence of product adoptions, and this inference can depend on network structure.

Finally, a different strand of the literature examines other obstructions to efficient social learning in settings where agents observe all possible peers rather than only neighbors in a social network. [Harel, Mossel, Strack, and Tamuz \(2021\)](#) study a social-learning environment with coarse communication and find, as in our generations network, that agents learn at the same rate as they would if they perfectly observed an arbitrarily small fraction of private

signals. The mechanism behind their result (“rational groupthink”) is not related to an observation network preventing some agents from seeing others’ social information, but rather relies on agents’ finite action spaces obscuring all information about their private signals for many periods.²² Another group of papers point out that if signals about the state come from myopic agents’ information-acquisition choices, then individuals can make socially inefficient choices and slow down learning (Burguet and Vives, 2000; Mueller-Frank and Pai, 2016; Ali, 2018a; Lomys, 2024; Liang and Mu, 2020). We assume rich action spaces and exogenous signals to abstract from these obstructions and focus on the role of the network.

8 Conclusion

This paper presents a tractable model of sequential social learning that lets us compare social-learning dynamics across different observation networks. In our environment, rational actions are a log-linear function of observations and admit a signal-counting interpretation. Thus, we can measure the efficiency of learning in terms of the fraction of available signals incorporated into beliefs asymptotically (“aggregative efficiency”) and make precise comparisons about the rate of learning and welfare across different networks.

The network causes information confounding when an agent does not see an early predecessor whose action influenced several of the agent’s neighbors. We show that confounding can be a powerful obstacle to social learning: even little confounding can cause almost total information loss. For a class of symmetric networks where agents move in generations, we derive a simple expression for aggregative efficiency. For any network in this class, social learning aggregates no more than two signals per generation in the long run, even for arbitrarily large generations. We also compute comparative statics of learning with respect to network parameters, finding that additional observations speed up learning but extra confounding slows it down.

We have focused on how the network structure affects social learning and abstracted away from many other sources of learning-rate inefficiency. These other sources may realistically co-exist with the informational-confounding issues discussed here and complicate the analysis. For instance, even though the complete network allows agents to exactly infer every predecessor’s private signal, it could lead to worse informational free-riding incentives in settings where agents must pay for the precision of their private signals (compared to networks

²²Huang, Strack, and Tamuz (2024) extend the results of Harel et al. (2021) to give a uniform bound on the rate of learning across strongly connected networks. The obstruction to learning continues to be coarse actions and not network structure, however. Indeed, Huang et al. (2024)’s result on general networks (which may introduce additional confounding) allows faster learning than the bound on the complete network from Harel et al. (2021).

where agents have fewer observations). Studying the trade-offs and/or interactions between network-based information confounding and other obstructions to fast learning could lead to fruitful future work.

References

- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): “Bayesian learning in social networks,” *Review of Economic Studies*, 78, 1201–1236.
- ALI, S. N. (2018a): “Herding with costly information,” *Journal of Economic Theory*, 175, 713–729.
- (2018b): “On the role of responsiveness in rational herds,” *Economics Letters*, 163, 79–82.
- ARIELI, I. AND M. MUELLER-FRANK (2019): “Multidimensional social learning,” *Review of Economic Studies*, 86, 913–940.
- BANERJEE, A. AND D. FUDENBERG (2004): “Word-of-mouth learning,” *Games and Economic Behavior*, 46, 1–22.
- 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.
- BILLINGSLEY, P. (2013): *Convergence of Probability Measures*, John Wiley and Sons.
- BLACKWELL, D. (1945): “Finite non-homogeneous chains,” *Annals of Mathematics*, 594–599.
- BOARD, S. AND M. MEYER-TER VEHN (2021): “Learning dynamics in social networks,” *Econometrica*, 89, 2601–2635.
- BURGUET, R. AND X. VIVES (2000): “Social learning and costly information acquisition,” *Economic Theory*, 15, 185–205.
- CASCIARO, T., A. C. EDMONDSON, AND S. JANG (2019): “Cross-silo leadership,” *Harvard Business Review*, <https://hbr.org/2019/05/cross-silo-leadership>.

- CHAO, G. (2007): “Mentoring and organizational socialization: Networks for work adjustment,” in *The Handbook of Mentoring at Work*, ed. by B. R. Ragins and K. E. Kram, Sage Publications, chap. 7, 179–196.
- DASARATHA, K., B. GOLUB, AND N. HAK (2023): “Learning from neighbours about a changing state,” *Review of Economic Studies*, 90, 2326–2369.
- EYSTER, E. AND M. RABIN (2014): “Extensive imitation is irrational and harmful,” *Quarterly Journal of Economics*, 129, 1861–1898.
- GLEESON, B. AND M. ROZO (2013): “The silo mentality: How to break down the barriers,” *Forbes*, <https://www.forbes.com/sites/brentgleeson/2013/10/02/the-silo-mentality-how-to-break-down-the-barriers>.
- GOLUB, B. AND E. SADLER (2016): “Learning in Social Networks,” in *The Oxford Handbook of the Economics of Networks*, Oxford University Press.
- HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2021): “Rational groupthink,” *Quarterly Journal of Economics*, 136, 621–668.
- HUANG, W., P. STRACK, AND O. TAMUZ (2024): “Learning in repeated interactions on networks,” *Econometrica*, 92, 1–27.
- IMMORLICA, N., J. MAO, A. SLIVKINS, AND Z. S. WU (2020): “Incentivizing Exploration with Selective Data Disclosure,” in *Proceedings of the 21st ACM Conference on Economics and Computation*, New York, NY, USA: Association for Computing Machinery, EC ’20, 647–648.
- LIANG, A. AND X. MU (2020): “Complementary information and learning traps,” *Quarterly Journal of Economics*, 135, 389–448.
- LOBEL, I., D. ACEMOGLU, M. DAHLEH, AND A. OZDAGLAR (2009): “Rate of convergence of learning in social networks,” *Proceedings of the American Control Conference*.
- LOBEL, I. AND E. SADLER (2015): “Information diffusion in networks through social learning,” *Theoretical Economics*, 10, 807–851.
- LOMYS, N. (2024): “Collective search in networks,” *Working Paper*.
- MATHON, R. AND A. ROSA (1985): “Tables of parameters of BIBDs with $r \leq 41$ including existence, enumeration, and resolvability results,” in *Annals of Discrete Mathematics (26)*:

- Algorithms in Combinatorial Design Theory*, ed. by C. Colbourn and M. Colbourn, North-Holland, vol. 114 of *North-Holland Mathematics Studies*, 275 – 307.
- MOLAVI, P., A. TAHBAZ-SALEHI, AND A. JADBABAIE (2018): “A theory of non-Bayesian social learning,” *Econometrica*, 86, 445–490.
- MORRIS, S. AND H. S. SHIN (2002): “Social value of public information,” *American Economic Review*, 92, 1521–1534.
- MUELLER-FRANK, M. AND M. M. PAI (2016): “Social learning with costly search,” *American Economic Journal: Microeconomics*, 8, 83–109.
- ROSENBERG, D. AND N. VIEILLE (2019): “On the efficiency of social learning,” *Econometrica*, 87, 2141–2168.
- RYSER, H. J. (1963): *Combinatorial Mathematics*, American Mathematical Society.
- SETHI, R. AND M. YILDIZ (2016): “Communication with unknown perspectives,” *Econometrica*, 84, 2029–2069.
- SGROI, D. (2002): “Optimizing information in the herd: Guinea pigs, profits, and welfare,” *Games and Economic Behavior*, 39, 137–166.
- SMITH, L. AND P. SØRENSEN (2000): “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398.
- STOLARCZYK, S., M. BHARDWAJ, K. E. BASSLER, W. J. MA, AND K. JOSIĆ (2017): “Loss of information in feedforward social networks,” *Journal of Complex Networks*, 6, 448–469.
- TAHBAZ-SALEHI, A. AND A. JADBABAIE (2006): “On consensus over random networks,” in *44th Annual Allerton Conference*, Citeseer.
- TETT, G. (2015): *The Silo Effect: The Peril of Expertise and the Promise of Breaking Down Barriers*, New York: Simon and Schuster.
- WOLITZKY, A. (2018): “Learning from others’ outcomes,” *American Economic Review*, 108, 2763–2801.

Appendix

A Proofs

A.1 Details on Example 2

We show that for the network in Figure 3, agent $K_1 + K_2 + 1$'s rational log-strategy puts weight $\frac{1+K_1}{1+K_1K_2}$ on each neighbor's log-action, and hence $r_{K_1+K_2+1} = 1 + \frac{K_2+K_1K_2}{1+K_1K_2} \cdot (1 + K_1)$.

For $1 \leq j \leq K_2$, we have $\ell_{K_1+j} = \sum_{i=1}^{K_1} \lambda_i + \lambda_{K_1+j}$. So, $\mathbb{E}[\ell_{K_1+j} \mid \omega = 1] = (K_1 + 1) \cdot \frac{2}{\sigma^2}$, $\text{Var}[\ell_{K_1+j} \mid \omega = 1] = \frac{4}{\sigma^2}(K_1+1)$, while $\text{Cov}[\ell_{K_1+j}, \ell_{K_1+j'} \mid \omega = 1] = \frac{4}{\sigma^2}K_1$ for $1 \leq j < j' \leq K_2$. By Proposition 1, the vector of weights that the final agent's rational log-strategy puts on neighbors' log actions is given by

$$2 \cdot \frac{2}{\sigma^2} \cdot \begin{bmatrix} K_1 + 1 & K_1 + 1 & \cdots \end{bmatrix} \cdot \frac{\sigma^2}{4} \cdot \underbrace{\begin{bmatrix} K_1 + 1 & K_1 & \cdots & K_1 \\ K_1 & K_1 + 1 & \cdots & K_1 \\ \vdots & \vdots & \ddots & \vdots \\ K_1 & K_1 & \cdots & K_1 + 1 \end{bmatrix}^{-1}}_{K_2 \text{ by } K_2}.$$

The matrix inverse is equal to

$$\frac{1}{K_1K_2 + 1} \underbrace{\begin{bmatrix} (K_2 - 1)K_1 + 1 & -K_1 & \cdots & -K_1 \\ -K_1 & (K_2 - 1)K_1 + 1 & \cdots & -K_1 \\ \vdots & \vdots & \ddots & \vdots \\ -K_1 & -K_1 & \cdots & (K_2 - 1)K_1 + 1 \end{bmatrix}}_{K_2 \text{ by } K_2}.$$

Therefore, weight on each neighbor is $\frac{1+K_1}{1+K_1K_2}$. Also, since $\mathbb{E}[\ell_{K_1+j} \mid \omega = 1] = (K_1 + 1) \cdot \frac{2}{\sigma^2}$ for each neighbor $K_1 + j$ and there are K_2 neighbors, we get $\mathbb{E}[\ell_{K_1+K_2+1} \mid \omega = 1] = [1 + \frac{K_2+K_1K_2}{1+K_1K_2} \cdot (1 + K_1)] \cdot \frac{2}{\sigma^2}$. By the signal counting interpretation, $r_{K_1+K_2+1} = 1 + \frac{K_2+K_1K_2}{1+K_1K_2} \cdot (1 + K_1)$.

A.2 Proof of Proposition 1

We first prove a lemma about the conditional distributions of the log-signals.

Lemma A.1. *For each i , the log-signal λ_i has a Gaussian distribution conditional on ω , with $\mathbb{E}[\lambda_i \mid \omega = 0] = -2/\sigma^2$, $\mathbb{E}[\lambda_i \mid \omega = 1] = 2/\sigma^2$, and $\text{VAR}[\lambda_i \mid \omega = 0] = \text{VAR}[\lambda_i \mid \omega = 1] = 4/\sigma^2$.*

Proof. We show that $\lambda_i = \frac{2}{\sigma^2} s_i$. This is because

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

The result then follows from scaling the conditional distributions of s_i , $(s_i | \omega = 1) \sim \mathcal{N}(1, \sigma^2)$ and $(s_i | \omega = 0) \sim \mathcal{N}(-1, \sigma^2)$. \square

Now we prove Proposition 1.

Proof. Agent 1 does not observe any predecessors, so clearly $\mathcal{L}_1^*(\lambda_1) = \lambda_1$. Suppose by way of induction that the rational strategies of all agents $j \leq I - 1$ are linear. Then each ℓ_j for $j \leq I - 1$ is a linear combination of $(\lambda_h)_{h=1}^I$, which by Lemma A.1 are conditionally Gaussian with conditional means $\pm 2/\sigma^2$ in states $\omega = 1$ and $\omega = 0$ and conditional variance $4/\sigma^2$ in each state. This implies $(\ell_{j(1)}, \dots, \ell_{j(n_I)})$ have a conditional joint Gaussian distribution with $(\ell_{j(1)}, \dots, \ell_{j(n_I)}) \sim \mathcal{N}(\vec{\mu}, \Sigma)$ conditional on $\omega = 1$, and $(\ell_{j(1)}, \dots, \ell_{j(n_I)}) \sim \mathcal{N}(-\vec{\mu}, \Sigma)$ conditional on $\omega = 0$, where $\vec{\mu} = \mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)})' | \omega = 1]$ and $\Sigma = \text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} | \omega = 1]$.

From the the multivariate Gaussian density, (writing $(\ell_{j(1)}, \dots, \ell_{j(n_I)})' = \vec{a}$),

$$\begin{aligned}\ln \left(\frac{\mathbb{P}[\ell_{j(1)}, \dots, \ell_{j(n_I)} | \omega = 1]}{\mathbb{P}[\ell_{j(1)}, \dots, \ell_{j(n_I)} | \omega = 0]} \right) &= \ln \left(\frac{\exp(-\frac{1}{2}(\vec{a} - \vec{\mu})' \Sigma^{-1} (\vec{a} - \vec{\mu}))}{\exp(-\frac{1}{2}(\vec{a} + \vec{\mu})' \Sigma^{-1} (\vec{a} + \vec{\mu}))} \right) \\ &= \vec{a}' \Sigma^{-1} \vec{\mu} + \vec{\mu}' \Sigma^{-1} \vec{a}\end{aligned}$$

which is $2(\vec{\mu}' \Sigma^{-1}) \cdot (\ell_{j(1)}, \dots, \ell_{j(n_I)})'$ because Σ is symmetric. This then shows agent I 's rational strategy must also be linear, completing the inductive step. This argument also gives the explicit form of $\vec{\beta}_I, \dots$

For the final statement, we prove another lemma. The argument so far implies that we may find weights $(w_{i,j})_{j \leq i}$ so that the realizations of rational log-actions are related to the realizations of log-signals by $\ell_i = \sum_{j=1}^i w_{i,j} \lambda_j$. Let W be the matrix containing all such weights.

Lemma A.2. *Let \hat{W} be the submatrix of W with rows $N(i)$ and columns $\{1, \dots, i - 1\}$. Then $\vec{\beta}_i = \mathbf{1}_{(i-1)}' \times \hat{W}' (\hat{W} \hat{W}')^{-1}$ and the i -th row of W is $W_i = \left((\vec{\beta}_i' \times \hat{W}), 1, 0, 0, \dots \right)$.*

Proof. Suppose $N(i) = \{j(1), \dots, j(d_i)\}$ with $j(1) < \dots < j(d_i)$. By Lemma A.1 and construction of \hat{W} , we have $\mathbb{E}[\ell_{j(k)} | \omega = 1] = \frac{2}{\sigma^2} \sum_{h=1}^{i-1} \hat{W}_{k,h}$. So, $\mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)}) | \omega = 1] = \frac{2}{\sigma^2} (\hat{W} \cdot \mathbf{1}_{(i-1)})' = \frac{2}{\sigma^2} \mathbf{1}_{(i-1)}' \hat{W}'$. Also, again by Lemma A.1 and construction of \hat{W} , we can calculate that for $1 \leq k_1 \leq k_2 \leq d_i$, $\text{COV}[\ell_{j(k_1)}, \ell_{j(k_2)} | \omega = 1] = \frac{4}{\sigma^2} \sum_{h=1}^{i-1} (\hat{W}_{k_1,h} \hat{W}_{k_2,h})$, meaning

$\text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} \mid \omega = 1] = \frac{4}{\sigma^2} \hat{W} \hat{W}'$. It then follows from what we have shown above that $\vec{\beta}_{i,\cdot} = 2 \cdot \frac{2}{\sigma^2} \mathbf{1}'_{(i-1)} \hat{W}' \times \left[\frac{4}{\sigma^2} \hat{W} \hat{W}' \right]^{-1} = \mathbf{1}'_{(i-1)} \times \hat{W}' (\hat{W} \hat{W}')^{-1}$.

Since i puts weight 1 on λ_i and weights $\vec{\beta}_{i,\cdot}$ on $(\ell_{j(1)}, \dots, \ell_{j(d_i)})' = \hat{W} \times (\lambda_1, \dots, \lambda_{i-1})'$, this shows the first $i - 1$ elements in the row W_i must be $\vec{\beta}'_{i,\cdot} \cdot \hat{W}$ while the i -th element is 1. \square

To prove the final statement of Proposition 1, first observe that $W_1 = (1, 0, 0, \dots)$ does not depend on σ^2 . The same applies to $\vec{\beta}_{1,\cdot}$. By way of induction, suppose rows W_i and vectors $\vec{\beta}_{i,\cdot}$ do not depend on σ^2 for any $i \leq I$. If \hat{W} is the submatrix of W with rows $N(I + 1)$, then since $N(I + 1) \subseteq \{1, \dots, I\}$, by the inductive hypothesis \hat{W} must be independent of σ^2 . Thus the same independence also applies to $\vec{\beta}'_{I+1,\cdot}$, since this vector only depends on \hat{W} by the result just derived. In turn, since W_{I+1} is only a function of $\vec{\beta}'_{I+1,\cdot}$ and \hat{W} , and these terms are independent of σ^2 as argued before, same goes for W_{I+1} , completing the inductive step. \square

A.3 Proof of Proposition 2

Proof. It suffices to show that $\mathbb{E}[\ell_i \mid \omega = 1] = \frac{1}{2} \text{VAR}[\ell_i \mid \omega = 1]$. By Proposition 1, $\ell_i = \lambda_i + \sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)}$. From Lemma A.1, we have $\mathbb{E}[\lambda_i \mid \omega = 1] = \frac{1}{2} \text{VAR}[\lambda_i \mid \omega = 1]$. Furthermore, λ_i is independent from $\sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)}$, as the latter term only depends on $\lambda_1, \dots, \lambda_{i-1}$. So we need only show $\mathbb{E}[\sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)} \mid \omega = 1] = \frac{1}{2} \text{VAR}[\sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)} \mid \omega = 1]$

Let $\vec{\mu} = \mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)})' \mid \omega = 1]$ and $\Sigma = \text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} \mid \omega = 1]$. Using the expression for $\vec{\beta}_{i,\cdot}$ from Proposition 1, $\mathbb{E}[\sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)} \mid \omega = 1] = 2 (\vec{\mu}' \Sigma^{-1}) \cdot \vec{\mu}$. Also,

$$\text{VAR} \left[\sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)} \mid \omega = 1 \right] = (2 \vec{\mu}' \Sigma^{-1}) \Sigma (2 \vec{\mu}' \Sigma^{-1})' = 4 \vec{\mu}' \Sigma^{-1} \vec{\mu}$$

using the fact that Σ is a symmetric matrix. This is twice $\mathbb{E}[\sum_{k=1}^{d_i} \beta_{i,j(k)} \ell_{j(k)} \mid \omega = 1]$ as desired. \square

A.4 Proof of Corollary 1

Proof. When $i < I$ use log-linear strategies, each ℓ_i is some linear combination of $(\lambda_h)_{h \leq I-1}$. Thus, $(\ell_j)_{j \in N(I)}$ are conditionally jointly Gaussian, $(\ell_j)_{j \in N(I)} \mid \omega \sim \mathcal{N}(\pm \vec{\mu}, \Sigma)$. This is sufficient for the the proofs of Propositions 1 and 2 to go through, implying that the ℓ_I maximizing I 's expected utility using the information in $(\ell_j)_{j \in N(I)}$ is a log-linear strategy and has a signal-counting interpretation. \square

A.5 Proof of Proposition 3

We first state and prove an auxiliary lemma.

Lemma A.3. *For any $0 < \epsilon < 0.5$, $\mathbb{P}[a_i > 1 - \epsilon \mid \omega = 1] = 1 - \Phi\left(\frac{\ln\left(\frac{1-\epsilon}{\epsilon}\right) - r_i \frac{2}{\sigma^2}}{\sqrt{r_i} \frac{2}{\sigma}}\right)$, where Φ is the standard Gaussian distribution function. This expression is increasing in r_i and approaches 1. Also, $\mathbb{P}[a_i < \epsilon \mid \omega = 0] = \Phi\left(\frac{\ln\left(\frac{1-\epsilon}{\epsilon}\right) + r_i \frac{2}{\sigma^2}}{\sqrt{r_i} \frac{2}{\sigma}}\right)$. This expression is increasing in r_i and approaches 1.*

Proof. Note that $a_i > 1 - \epsilon$ if and only if $\ell_i > \ln\left(\frac{1-\epsilon}{\epsilon}\right) > 0$. Given that $(\ell_i \mid \omega = 1) \sim \mathcal{N}\left(r_i \cdot \frac{2}{\sigma^2}, r_i \cdot \frac{4}{\sigma^2}\right)$ by Proposition 2, the expression for $\mathbb{P}[a_i > 1 - \epsilon \mid \omega = 1]$ follows. To see that it is increasing in r_i , observe that $\frac{d}{dr_i} \frac{\ln\left(\frac{1-\epsilon}{\epsilon}\right) - r_i \frac{2}{\sigma^2}}{\sqrt{r_i} \frac{2}{\sigma}}$ has the same sign as

$$-\frac{2}{\sigma^2}(\sqrt{r_i} \frac{2}{\sigma^2}) - (\ln\left(\frac{1-\epsilon}{\epsilon}\right) - r_i \frac{2}{\sigma^2})\left(\frac{1}{2} r_i^{-0.5} \frac{2}{\sigma}\right) = -\frac{2}{\sigma^3} \sqrt{r_i} - \ln\left(\frac{1-\epsilon}{\epsilon}\right) r_i^{-0.5} \frac{1}{\sigma} < 0.$$

Also, it is clear that $\lim_{r_i \rightarrow \infty} \frac{\ln\left(\frac{1-\epsilon}{\epsilon}\right) - r_i \frac{2}{\sigma^2}}{\sqrt{r_i} \frac{2}{\sigma}} = -\infty$, hence $\lim_{r_i \rightarrow \infty} \mathbb{P}[a_i > 1 - \epsilon \mid \omega = 1] = 1$. The results for $\mathbb{P}[a_i < \epsilon \mid \omega = 0]$ follow from analogous arguments. \square

We now turn to the proof of Proposition 3.

Proof. By Proposition 2, there exist $(r_i)_{i \geq 1}$ so that social learning aggregates r_i signals by agent i . We first show that (3) and (4) in Proposition 3 are equivalent. Let $\epsilon' > 0$ be given and suppose $\lim_{i \rightarrow \infty} r_i = \infty$. Putting $\epsilon = \min(\epsilon', 0.4)$, we get that $\mathbb{P}[|a_i - \omega| < \epsilon \mid \omega = 1] \rightarrow 1$ and $\mathbb{P}[|a_i - \omega| < \epsilon \mid \omega = 0] \rightarrow 1$ since the two expressions in Lemma A.3 increase in r_i and approach 1, hence also $\mathbb{P}[|a_i - \omega| < \epsilon'] \rightarrow 1$. So society learns completely in the long run. Conversely, if we do not have $\lim_{i \rightarrow \infty} r_i = \infty$, then for some $K < \infty$ we have $r_i < K$ for infinitely many i . By Lemma A.3 we will get that $\mathbb{P}[|a_i - \omega| < 0.1 \mid \omega = 1]$ are bounded by $1 - \Phi\left(\frac{\ln(9) - K \frac{2}{\sigma^2}}{\sqrt{K} \frac{2}{\sigma}}\right)$ for these i , hence society does not learn completely in the long run.

Next, we show that Conditions (1) and (2) in the proposition are both equivalent to Condition (3), $\lim_{i \rightarrow \infty} r_i = \infty$.

Condition (1): $\lim_{i \rightarrow \infty} \mathcal{P}\mathcal{L}(i) = \infty$.

Necessity: Suppose $\lim_{i \rightarrow \infty} r_i = \infty$. For $h \in \mathbb{N}$, let $I(h) := \{i : \mathcal{P}\mathcal{L}(i) = h\}$. We show by induction that $I(h)$ is finite for all $h \in \mathbb{N}$. For every $i \in I(0)$, $r_i = 1$, so $\lim_{i \rightarrow \infty} r_i = \infty$ implies $|I(0)| < \infty$. Now suppose $|I(h)| < \infty$ for all $h \leq L$. If $i \in I(L+1)$, then every j that can be reached along M from i must belong to $I(h)$ for some $h \leq L$. The subnetwork containing i is therefore a subset of $\cup_{h=0}^L I(h)$, a finite set by the inductive hypothesis. Thus $r_i \leq 1 + \sum_{h=0}^L |I(h)|$ for all $i \in I(L+1)$. So $\lim_{i \rightarrow \infty} r_i = \infty$ implies $I(L+1)$ is finite, completing the inductive step and proving $I(h)$ is finite for all h . Hence $\lim_{i \rightarrow \infty} \mathcal{P}\mathcal{L}(i) = \infty$.

Sufficiency: First note if $j \in N(i)$, then $r_i \geq r_j + 1$. This is because $\ell_j \sim \mathcal{N}\left(\pm r_j \cdot \frac{2}{\sigma^2}, r_j \cdot \frac{4}{\sigma^2}\right)$ conditional on the two states, and furthermore ℓ_j is conditionally independent of s_i . So, $\ell_j + \lambda_i$ is a possibly play for i , which would have the conditional distributions $\mathcal{N}\left(\pm(r_j + 1) \cdot \frac{2}{\sigma^2}, (r_j + 1) \cdot \frac{4}{\sigma^2}\right)$ in the two states. If $r_i < r_j + 1$, then i would have a profitable deviation by choosing $\ell_i = \ell_j + \lambda_i$ instead, since it follows from Lemma A.3 that a log-action that aggregates more signals leads to higher expected payoffs.

Condition (2): $\lim_{i \rightarrow \infty} \left[\max_{j \in N(i)} j \right] = \infty$.

Necessity: If Condition (2) is violated, there exists some $\bar{j} < \infty$ so that there exist infinitely many i 's with $N(i) \subseteq \{1, \dots, \bar{j}\}$. The subnetwork containing any such i is a subset of $\{1, \dots, \bar{j}\}$, so $r_i \leq \bar{j} + 1$. We cannot have $\lim_{i \rightarrow \infty} r_i = \infty$.

Sufficiency: Construct an increasing sequence $C_1 \leq C_2 \leq \dots$ as follows. Condition (2) implies there exists C_1 so that $\max_{j \in N(i)} j \geq 1$ for all $i \geq C_1$. So, $\mathcal{PL}(i) \geq 1$ for all $i \geq C_1$. Suppose $C_1 \leq \dots \leq C_n$ are constructed with the property that $\mathcal{PL}(i) \geq k$ for all $i \geq C_k$, $k = 1, \dots, n$. Condition (2) implies there exists C_{n+1} so that $\max_{j \in N(i)} j \geq C_n$ for all $i \geq C_{n+1}$. But since all $j \geq C_n$ have $\mathcal{PL}(j) \geq n$ by the inductive hypothesis, all $i \geq C_{n+1}$ must have $\mathcal{PL}(i) \geq n + 1$, completing the inductive step. This shows $\lim_{i \rightarrow \infty} \mathcal{PL}(i) = \infty$. By the sufficiency of Condition (1) for $\lim_{i \rightarrow \infty} r_i = \infty$, we see that Condition (2) implies the same. \square

A.6 Proof of Theorem 1

Proof. If $d = 1$, then exactly one signal is aggregated per generation so $r_i/K \rightarrow 1$ as required. Also, if $c = 0$, then we must have $d = 1$. From now on we assume $d \geq 2$ and $c \geq 1$.

Lemma A.4. *For $d \geq 2$, each generation t and each $i \neq i'$ in generation t , $\text{VAR}[\ell_i \mid \omega = 1]$ and $\text{COV}[\ell_i, \ell_{i'} \mid \omega = 1]$ depend only on t and not on the identities of i or i' , which we call VAR_t and COV_t , respectively. Similarly, for i in generation t and each $j \in N(i)$, the weight $\beta_{i,j}$ depends only on t , which we call β_t .*

Proof. The results hold by inductively applying the symmetry condition. Clearly they are true for $t = 2$. Suppose they are true for all $t \leq T$. For an agent i in generation $t = T + 1$, the inductive hypothesis implies $\text{VAR}[\ell_j \mid \omega = 1]$ is the same for all $j \in N(i)$, and all pairs $j, j' \in N(i)$ with $j \neq j'$ have the same conditional covariance. Also, using Proposition 2, $\mathbb{E}[\ell_j \mid \omega = 1]$ is the same for all $j \in N(i)$. Thus by Proposition 1, i places the same weight, say β_t , on all neighbors. Using the fact that $\ell_i = \lambda_i + \sum_{j \in N(i)} \beta_t \ell_j$, we have the recursive expressions $\text{VAR}[\ell_i \mid \omega = 1] = \frac{4}{\sigma^2} + \beta_t^2(d\text{VAR}_{t-1} + (d^2 - d)\text{COV}_{t-1})$ for all i in generation t , and $\text{COV}[\ell_i, \ell_{i'} \mid \omega = 1] = \beta_t^2(c\text{VAR}_{t-1} + (d^2 - c)\text{COV}_{t-1})$ for all agents $i \neq i'$ in generation t . This shows the claims for $t = T + 1$, and completes the proof by induction. \square

Taking the difference of the two expressions for VAR_t and COV_t gives:

$$\text{VAR}_t - \text{COV}_t = \frac{4}{\sigma^2} + \beta_t^2(d - c)(\text{VAR}_{t-1} - \text{COV}_{t-1}). \quad (1)$$

We now require two auxiliary lemmas.

Lemma A.5. *Consider the Markov chain on $\{1, \dots, K\}$ with state transition matrix p , with $p_{i,j} = \mathbb{P}[i \rightarrow j] = 1/d$ if $j \in \Psi_i$, 0 otherwise. Suppose $(\Psi_k)_k$ is symmetric with $c \geq 1$. Then $p_i^\infty := \lim_{t \rightarrow \infty} (p^t)_i \in [0, 1]^K$ exists, and it is the same for all $1 \leq i \leq K$.*

Proof. For existence of p_i^∞ , consider the decomposition of the Markov chain into its communication classes, $C_1, \dots, C_L \subseteq \{1, \dots, K\}$. Without loss suppose the first L' communication classes are closed and the rest are not.

We show that each closed communication class is aperiodic when $(\Psi_k)_k$ is symmetric and $c, d \geq 1$. Let $i \in C_m$ for $1 \leq m \leq L'$. Let $\Psi_i = \{j_1, \dots, j_d\}$. If $i \in \Psi_i$, then i 's periodicity is 1. Otherwise, $\Psi_i \subseteq C_m$ since C_m is closed, so for every $1 \leq h \leq d$ there exists a cycle of some length Q_h starting at i , where the h -th such cycle is $i \rightarrow j_h \rightarrow \dots \rightarrow i$. Since $c \geq 1$, i and j_1 share a common neighbor, which must be j_{h^*} for some $1 \leq h^* \leq d$. We can therefore construct a cycle of length $Q_{h^*} + 1$ starting at i , $i \rightarrow j_1 \rightarrow j_{h^*} \rightarrow \dots \rightarrow i$. Since cycle lengths Q_{h^*} and $Q_{h^*} + 1$ are coprime, i 's periodicity is 1.

By standard results (see e.g., Billingsley (2013)) there exist $\nu_m^*, 1 \leq m \leq L'$, so that $\lim_{t \rightarrow \infty} (p^t)_i = \nu_m^*$ whenever $i \in C_m$. If $i \notin \cup_{1 \leq m \leq L'} C_m$, then starting the process at i , almost surely the process enters one of the closed communication classes eventually. This shows $\lim_{t \rightarrow \infty} (p^t)_i$ exists and is equal to $\sum_{m=1}^{L'} q_m \nu_m^*$, where q_m is the probability that the process started at i enters C_m before any other closed communication class.

To prove that p_i^∞ is the same for all i , we inductively show that for all $i \neq j$, $\|p_i^\infty - p_j^\infty\|_{\max} \leq \left(\frac{d-c}{d}\right)^t$ for all $t \geq 1$. Since $c \geq 1$, this would show that in fact $p_i^\infty = p_j^\infty$ for all i, j .

For the base case of $t = 1$, enumerate $\Psi_i = \{n_1, \dots, n_c, n_{c+1}, \dots, n_d\}$, $\Psi_j = \{n_1, \dots, n_c, n'_{c+1}, \dots, n'_d\}$ where all $n_1, \dots, n_d, n'_{c+1}, \dots, n'_d \in \{1, \dots, K\}$ are distinct. Then

$$p_i^\infty = \frac{1}{d} \left(\sum_{k=1}^c p_{n_k}^\infty \right) + \frac{1}{d} \left(\sum_{k=c+1}^d p_{n_k}^\infty \right),$$

$$p_j^\infty = \frac{1}{d} \left(\sum_{k=1}^c p_{n_k}^\infty \right) + \frac{1}{d} \left(\sum_{k=c+1}^d p_{n'_k}^\infty \right), \text{ so}$$

$$\|p_i^\infty - p_j^\infty\|_{\max} \leq \frac{1}{d} \sum_{k=c+1}^d \|p_{n_k}^\infty - p_{n'_k}^\infty\|_{\max} \leq \frac{d-c}{d} \cdot 1$$

where the 1 comes from $\|x - y\|_{\max} \leq 1$ for any two distributions x, y .

The inductive step just replaces the bound $\|x - y\|_{\max} \leq 1$ with

$$\|p_{n_k}^\infty - p_{n'_k}^\infty\|_{\max} \leq \left(\frac{d-c}{d}\right)^{t-1}$$

from the inductive hypothesis. □

Lemma A.6. $\beta_t \rightarrow 1/d$.

Proof. For i in generation $t + 1$, $\ell_i = \lambda_i + \beta_{t+1} \sum_{j \in N(i)} \ell_j$, so as in the proof of Lemma A.4, $\text{VAR}[\ell_i | \omega = 1] = \frac{4}{\sigma^2} + \beta_{t+1}^2 (d \text{VAR}_t + (d^2 - d) \text{COV}_t)$. Using the definition of the signal-counting interpretation and Proposition 2, $\mathbb{E}[\ell_j | \omega = 1] = \frac{1}{2} \text{VAR}_t$ for each $j \in N(i)$, and so $\mathbb{E}[\ell_i | \omega = 1] = \frac{2}{\sigma^2} + d\beta_{t+1}(\frac{1}{2} \text{VAR}_t)$. By the same argument we also have $\text{VAR}[\ell_i | \omega = 1] = 2 \cdot \mathbb{E}[\ell_i | \omega = 1]$, and this lets us solve out

$$\beta_{t+1} = \frac{\text{VAR}_t}{\text{VAR}_t + (d-1)\text{COV}_t} \geq \frac{1}{d}.$$

It is therefore sufficient to show that $\text{VAR}_t/\text{COV}_t \rightarrow 1$. The weight $w_{i,i'}$ that an agent i in generation t places on the private signal of an agent i' in generation $t - \tau$ is equal to the product of $\prod_{j=1}^{\tau} \beta_{t+1-j}$ and the number of paths from i to i' in the network M .

We can compute the number of paths as follows. Consider a Markov chain with states $\{1, \dots, K\}$ and state transition probabilities $\mathbb{P}[k_1 \rightarrow k_2] = 1/d$ if $k_2 \in \Psi_{k_1}$, $\mathbb{P}[k_1 \rightarrow k_2] = 0$. The number of paths from i in generation t to j in generation $t - \tau$ is equal to d^τ times the probability that the state is j after τ periods.

By Lemma A.5, there exists a stationary distribution $\pi^* \in \mathbb{R}_+^K$ with $\sum_{k=1}^K \pi_k^* = 1$ of the Markov chain. Given $\epsilon > 0$, we can choose τ_0 such that the number of paths from i in generation t to $j = (\tau - 1)K + k$ in generation $t - \tau$ is in $[d^\tau(\pi_k^* - \epsilon), d^\tau(\pi_k^* + \epsilon)]$ for all t and all $\tau \geq \tau_0$.

Fixing distinct agents i and i' in generation t :

$$\text{VAR}_t = \frac{4}{\sigma^2} + \frac{4}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^K w_{i,(t-\tau)K+k}^2 \quad \text{and} \quad \text{COV}_t = \frac{4}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^K w_{i,(t-\tau)K+k} w_{i',(t-\tau)K+k}.$$

We want to show that

$$\text{VAR}_t/\text{COV}_t = \frac{1 + \sum_{\tau=1}^{t-1} \sum_{k=1}^K w_{i,(t-\tau)K+k}^2}{\sum_{\tau=1}^{t-1} \sum_{k=1}^K w_{i,(t-\tau)K+k} w_{i',(t-\tau)K+k}} \rightarrow 1.$$

Take $\epsilon > 0$ smaller than π_k^* for all k . For $\tau \geq \tau_0$, we have

$$w_{i,(t-\tau)K+k} w'_{i',(t-\tau)K+k} \geq (d^\tau \prod_{j=1}^{\tau} \beta_{t+1-j})^2 (\pi_k^* - \epsilon)^2 \text{ and } w_{i,(t-\tau)K+k}^2 \leq (d^\tau \prod_{j=1}^{\tau} \beta_{t+1-j})^2 (\pi_k^* + \epsilon)^2$$

The covariance grows at least linearly in t since each $\beta \geq 1/d$, while the contribution from periods $t - \tau_0 + 1, \dots, t$ is bounded and therefore lower order. Thus,

$$\limsup_{t \rightarrow \infty} \text{VAR}_t / \text{COV}_t \leq \limsup_{t \rightarrow \infty} \frac{\sum_{k=1}^K \sum_{\tau=\tau_0}^{t-1} (d^\tau \prod_{j=1}^{\tau} \beta_{t+1-j})^2 (\pi_k^* + \epsilon)^2}{\sum_{k=1}^K \sum_{\tau=\tau_0}^{t-1} (d^\tau \prod_{j=1}^{\tau} \beta_{t+1-j})^2 (\pi_k^* - \epsilon)^2} \leq \max_{1 \leq k \leq K} \frac{(\pi_k^* + \epsilon)^2}{(\pi_k^* - \epsilon)^2}.$$

Since ϵ is arbitrary, this completes the proof of the lemma. \square

We return to the proof of Theorem 1. Fix small $\epsilon > 0$. By Lemma A.6, we can choose T such that $\beta_t \leq \frac{1+\epsilon}{d}$ for all $t \geq T$. Therefore, $\beta_t^2(d-c) \leq \frac{(1+\epsilon)^2}{d^2}(d-c)$ for $t \geq T$. Consider the contraction map $\varphi(x) = \frac{4}{\sigma^2} + \frac{(1+\epsilon)^2}{d^2}(d-c)x$. Iterating Equation (1) starting with $t = T$, we find that $\text{VAR}_t - \text{COV}_t \leq \varphi^{(t-T)}(\text{VAR}_T - \text{COV}_T)$, so this shows

$$\limsup_{t \rightarrow \infty} (\text{VAR}_t - \text{COV}_t) \leq \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2 - (1+\epsilon)^2 d + (1+\epsilon)^2 c}$$

where the RHS is the fixed point of φ . Since this holds for all small $\epsilon > 0$, we get $\limsup_{t \rightarrow \infty} (\text{VAR}_t - \text{COV}_t) \leq \frac{4}{\sigma^2} \frac{d^2}{d^2 - d + c}$.

At the same time, $\beta_t \geq \frac{1}{d}$ for all t . Consider the contraction map $\varphi(x) = \frac{4}{\sigma^2} + \frac{1}{d^2}(d-c)x$. Iterating Equation (1) starting with $t = 1$, we find that $\text{VAR}_t - \text{COV}_t \geq \varphi^{(t-1)}(\text{VAR}_1 - \text{COV}_1)$, so this shows

$$\liminf_{t \rightarrow \infty} (\text{VAR}_t - \text{COV}_t) \geq \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2 - d + c}$$

where the RHS is the fixed point of φ . Combining with the result before, we get $\lim_{t \rightarrow \infty} (\text{VAR}_t - \text{COV}_t) = \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2 - d + c}$.

As in the proof of Lemma A.6, for i in generation $t+1$, $\mathbb{E}[\ell_i \mid \omega = 1] = \frac{2}{\sigma^2} + d\beta_{t+1}(\frac{1}{2}\text{VAR}_t)$. Using the definition of signal-counting interpretation and Proposition 2, we have $\text{VAR}_{t+1} = 2 \cdot \mathbb{E}[\ell_i \mid \omega = 1] = 2(\beta_{t+1}d(\text{VAR}_t/2) + 2/\sigma^2)$, so

$$\begin{aligned} \text{VAR}_{t+1} - \text{VAR}_t &= (\beta_{t+1}d - 1)\text{VAR}_t + \frac{4}{\sigma^2} \\ &= \left(\frac{d\text{VAR}_t}{\text{VAR}_t + (d-1)\text{COV}_t} - 1 \right) \text{VAR}_t + \frac{4}{\sigma^2} \\ &= \left(\frac{d\text{VAR}_t}{d\text{VAR}_t - (d-1)(\text{VAR}_t - \text{COV}_t)} - 1 \right) \text{VAR}_t + \frac{4}{\sigma^2} \end{aligned}$$

Using $\lim_{t \rightarrow \infty} (\text{VAR}_t - \text{COV}_t) = \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2-d+c}$, we conclude

$$\begin{aligned} \lim_{t \rightarrow \infty} (\text{VAR}_{t+1} - \text{VAR}_t) &= \lim_{t \rightarrow \infty} \left(\frac{\text{VAR}_t}{\text{VAR}_t - \frac{4}{\sigma^2} \frac{d^2-d}{d^2-d+c}} - 1 \right) \text{VAR}_t + \frac{4}{\sigma^2}. \\ &= \lim_{t \rightarrow \infty} \left(\frac{4}{\sigma^2} \frac{d^2-d}{d^2-d+c} \cdot \frac{\text{VAR}_t}{\text{VAR}_t - \frac{4}{\sigma^2} \frac{d^2-d}{d^2-d+c}} \right) + \frac{4}{\sigma^2} \end{aligned}$$

Since $\text{VAR}_t \rightarrow \infty$, the asymptotic increase in conditional variance across successive generations is $\lim_{t \rightarrow \infty} (\text{VAR}_{t+1} - \text{VAR}_t) = \frac{4}{\sigma^2} \left(\frac{d^2-d}{d^2-d+c} + 1 \right)$. Since agent i is in generation $\lfloor i/K \rfloor$, we therefore have $r_i = \left(1 + \frac{d^2-d}{d^2-d+c} \right) \frac{i}{K} + o(i)$. So $\lim_{i \rightarrow \infty} (r_i/i) = \left(1 + \frac{d^2-d}{d^2-d+c} \right) \frac{1}{K}$. \square

A.7 Proof of Corollary 2

Proof. When $d \geq 2$ and $c < d$, the collection of symmetric observation sets with these parameters correspond to the collection of symmetric balanced incomplete block designs by Theorem 2.2 from Chapter 8 of Ryser (1963). If there exists at least one symmetric network with parameters (d, c, K) under the previous inequalities, then $K = \frac{d^2-d+c}{c}$ by Equation (3.17) from Chapter 8 of Ryser (1963).

Applying this result to the expression for aggregative efficiency from our Theorem 1, $\lim_{i \rightarrow \infty} (r_i/i) = \left(1 + \frac{d^2-d}{d^2-d+c} \right) \frac{1}{K} = \left(2 - \frac{c}{d^2-d+c} \right) \frac{1}{K} = \left(2 - \frac{1}{K} \right) \cdot \frac{1}{K}$. \square

A.8 Proof of Proposition 4

Proof. We first establish a lemma that expresses $\vec{\beta}_{i,\cdot}$ in closed-form for an agent i in generation $t+1$. Let ℓ_{sum} be the sum of the log-actions played in generation $t-1$. By the log-linearity of rational strategies (Proposition 1), there must exist some $\mu_{\text{sum}}, \sigma_{\text{sum}}^2 > 0$ so that the conditional distributions of ℓ_{sum} in the two states are $\mathcal{N}(\pm\mu_{\text{sum}}, \sigma_{\text{sum}}^2)$.

Lemma A.7. *Each element in $\vec{\beta}_{i,\cdot}$ is $\left(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2} \right) / \left(K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2} \right)$.*

Proof. An application of Proposition 1 shows each agent j in generation t aggregates ℓ_{sum} and own private signal λ_j according to $\ell_j = 2 \cdot \frac{\mu_{\text{sum}}}{\sigma_{\text{sum}}^2} \ell_{\text{sum}} + \lambda_j$.

Next, consider the problem of someone in generation $t+1$ who observes the log-actions ℓ_j of the K agents $j = (t-1)K + k$ for $1 \leq k \leq K$ from generation t . By symmetry, i places

the same weight on these K log-actions. To find this weight, we calculate

$$\begin{aligned}\mathbb{E} \left[\sum_{k=1}^K \ell_{(t-1)K+k} \mid \omega = 1 \right] &= 2K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + 2K \frac{1}{\sigma^2} \\ \text{VAR} \left[\sum_{k=1}^K \ell_{(t-1)K+k} \mid \omega = 1 \right] &= K \cdot \left(4 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + 4 \cdot \frac{1}{\sigma^2} \right) + K \cdot (K-1) \cdot 4 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2}\end{aligned}$$

So by Proposition 1,

$$\beta_{i,j} = \frac{2 \cdot \left(2K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + 2K \frac{1}{\sigma^2} \right)}{K \cdot \left(4 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + 4 \cdot \frac{1}{\sigma^2} \right) + K \cdot (K-1) \cdot 4 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2}} = \frac{\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}}{K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}}$$

for every $j = (t-1)K + k$ for $1 \leq k \leq K$, as desired. \square

Consider an agent i in generation t . From Proposition 2, there is some $x_{\text{old}} > 0$ so that $\ell_i \sim \mathcal{N}(\pm x_{\text{old}}, 2x_{\text{old}})$ conditional on the two states. In fact, from Proposition 1, $x_{\text{old}} = 2 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{2}{\sigma^2}$. For an agent in generation $t+1$, using the same argument and applying the formula for $\vec{\beta}_{i,\cdot}$ from Lemma A.7, we have $x_{\text{new}} = \frac{2K \left(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2} \right)^2}{K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}} + \frac{2}{\sigma^2}$.

A hypothetical agent who observes ℓ_{sum} (the sum of log-actions in generation $t-1$) with conditional distributions $\mathcal{N}(\pm \mu_{\text{sum}}, \sigma_{\text{sum}}^2)$ and three extra independent private signals (in addition to the one private signal usually observed) would play a log-action with conditional distributions $\mathcal{N}(\pm y, 2y)$ where $y = \left[2 \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{6}{\sigma^2} \right] + \frac{2}{\sigma^2}$. We have

$$\begin{aligned}(y - x_{\text{new}}) \cdot \left(K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2} \right) &= \left[2 \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{6}{\sigma^2} \right] \cdot \left[K \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2} \right] - 2K \left(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2} \right)^2 \\ &= (2 + 6K) \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} \cdot \frac{1}{\sigma^2} + \frac{6}{\sigma^4} - 4K \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} \cdot \frac{1}{\sigma^2} - 2K \frac{1}{\sigma^4} \\ &\geq 2K \frac{1}{\sigma^2} \left(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} - \frac{1}{\sigma^2} \right).\end{aligned}$$

We must have $\mathbb{P}[\ell_{\text{sum}} > 0 \mid \omega = 1] \geq \mathbb{P}[\lambda_1 > 0 \mid \omega = 1]$, a probability that just depends on the ratio of the mean and standard deviation. So $\frac{\mu_{\text{sum}}}{\sigma_{\text{sum}}} \geq \frac{1}{\sigma}$, i.e. $\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} \geq \frac{1}{\sigma^2}$. Hence the difference above is positive. This shows $x_{\text{new}} - x_{\text{old}} \leq 3 \cdot \frac{2}{\sigma^2}$. \square

A.9 Proof of Proposition 5

Proof. It is clear that $r_0 = 1$ and that $r_i = 2$ for $1 \leq i \leq K$. By applying the example in Section 2 with $K_1 = 1$, $K_2 = 2$, we see that for every agent j in generation 2, we get

$r_j \leq r_1 + 1 + \frac{(K-1)(1+1)}{K+1} < 5$. For an agent j' in generation $t \geq 3$, the same arguments in the proof of Proposition 4 apply, showing that $r_{j'} - r_{i'} \leq 3$ where i' is any agent in generation $t - 1$. \square

A.10 Proof of Proposition 6

First, we establish a lemma that tells us when aggregative efficiency is in $(0, 1)$, actions converge to the objectively optimal action at an exponential rate.

Lemma A.8. *Suppose the aggregative efficiency of a network is $AE \in (0, 1)$. For every $\epsilon > 0$ so that $0 < AE - \epsilon < AE + \epsilon < 1$,*

$$1 - \mathbb{P}[1 - e^{-i(AE-\epsilon)\cdot(2/\sigma^2)} \leq a_i \leq 1 - e^{-i(AE+\epsilon)\cdot(2/\sigma^2)} \mid \omega = 1] = O(e^{-i})$$

and

$$1 - \mathbb{P}[e^{-i(AE+\epsilon)\cdot(2/\sigma^2)} \leq a_i \leq e^{-i(AE-\epsilon)\cdot(2/\sigma^2)} \mid \omega = 0] = O(e^{-i}).$$

Proof. Conditional on $\omega = 1$, $\ell_i = r_i \cdot \frac{2}{\sigma^2} + z_i$ where $z_i \sim \mathcal{N}(0, r_i \cdot \frac{4}{\sigma^2})$. Since $a_i = \frac{\exp(\ell_i)}{1 + \exp(\ell_i)}$, we get

$$a_i = 1 - \frac{1}{1 + \exp(i \cdot (AE + [(r_i/i) - AE])) \cdot (2/\sigma^2)) \cdot \exp(z_i)}.$$

From this, we have

$$\begin{aligned} a_i &\geq 1 - \frac{1}{\exp(i \cdot (AE + [(r_i/i) - AE])) \cdot (2/\sigma^2)) \cdot \exp(z_i)} \\ &\geq 1 - \frac{1}{\exp(i \cdot (AE - (\epsilon/2)) \cdot (2/\sigma^2)) \cdot \exp(z_i)} \text{ for } i \text{ large} \\ &= 1 - \frac{1}{\exp(i \cdot (AE - \epsilon) \cdot (2/\sigma^2)) \cdot \exp(z_i + i \cdot (\epsilon/2) \cdot (2/\sigma^2))}. \end{aligned}$$

So for large i , $\mathbb{P}[a_i \geq 1 - e^{-i(AE-\epsilon)\cdot(2/\sigma^2)}] \geq \mathbb{P}[z_i + i \cdot (\epsilon/2) \cdot (2/\sigma^2) \geq 0]$. But the mean of $z_i + i \cdot (\epsilon/2) \cdot (2/\sigma^2)$ grows linearly in i and the standard deviation grows at most at the rate of \sqrt{i} , and it is well known that the complement of the Gaussian distribution function $\Phi(x)$ converges to 0 at the rate of e^{x^2} . This shows $\mathbb{P}[z_i + i \cdot (\epsilon/2) \cdot (2/\sigma^2) < 0] = O(e^{-i})$.

For the other direction, we have

$$\begin{aligned} a_i &\leq 1 - \frac{1}{1 + \exp(i \cdot (AE + \epsilon/3) \cdot (2/\sigma^2)) \cdot \exp(z_i)} \text{ for } i \text{ large} \\ &= 1 - \frac{1}{1 + \exp(i \cdot (AE + 2\epsilon/3) \cdot (2/\sigma^2)) \cdot \exp(-i \cdot (\epsilon/3) \cdot (2/\sigma^2) + z_i)} \end{aligned}$$

For large i , $\mathbb{P}[a_i \leq 1 - \frac{1}{1 + \exp(i \cdot (AE + 2\epsilon/3) \cdot (2/\sigma^2))}] \geq \mathbb{P}[-i \cdot (\epsilon/3) \cdot (2/\sigma^2) + z_i \leq 0]$. But the mean of $-i \cdot (\epsilon/3) \cdot (2/\sigma^2) + z_i$ decreases linearly in i and the standard deviation grows at most at the rate of \sqrt{i} , so by the same reason as before $\mathbb{P}[-i \cdot (\epsilon/3) \cdot (2/\sigma^2) + z_i > 0] = O(e^{-i})$. Finally, for large i , $\frac{1}{1 + \exp(i \cdot (AE + 2\epsilon/3) \cdot (2/\sigma^2))} > \frac{1}{\exp(i \cdot (AE + \epsilon) \cdot (2/\sigma^2))}$, so in fact $\mathbb{P}[a_i > 1 - e^{-i(AE + \epsilon) \cdot (2/\sigma^2)}] = O(e^{-i})$ as well.

The claim for $\omega = 0$ is symmetric. \square

Lemma A.8 implies that for any signal variance, the undiscounted infinite sums of the expected utilities $\sum_i v_i^M, \sum_i v_i^{M'} > -\infty$ are convergent. We now show $\sum_i \delta^{i-1} v_i^M > \sum_i \delta^{i-1} v_i^{M'}$ under the hypotheses of Proposition 6.

Proof. Find $\epsilon > 0$ small enough so that $AE_M - \epsilon > (AE_{M'} + \epsilon) \cdot (1 + \epsilon)$. There is some N so that each agent $i \geq N$ aggregates at least $(AE_M - \epsilon) \cdot i$ signals in network M and no more than $(AE_{M'} + \epsilon) \cdot i$ signals in network M' . Let C be the difference in expected utility between getting $AE_M - \epsilon$ signals of variance 1 and $(AE_{M'} + \epsilon) \cdot (1 + \epsilon)$ signals of variance 1. Find $\hat{\sigma}^2 > 0$ large enough so that whenever $\sigma^2 \geq \hat{\sigma}^2$, all agents before N in both networks will get utility so close to -0.25 that the undiscounted difference between the sums of utilities for the first $N - 1$ agents across the two networks is strictly smaller than $C/2$.

Let $\underline{\sigma}^2 = \max(\hat{\sigma}^2, 1/\epsilon, N)$. We will show that there exists a function $\underline{\delta} : [\underline{\sigma}^2, \infty) \rightarrow (0, 1)$ such that $\sum_i \delta^{i-1} v_i^M > \sum_i \delta^{i-1} v_i^{M'}$ whenever $\delta \geq \underline{\delta}(\sigma^2)$. Let $\underline{\delta}(\sigma^2)$ be $\log_{1/2}(\lceil \sigma^2 \rceil)$, so that $\underline{\delta}(\sigma^2) \lceil \sigma^2 \rceil = C/2$. Whenever $\sigma^2 \geq \underline{\sigma}^2$ and $\delta \geq \underline{\delta}(\sigma^2)$, first note $v_i^M > v_i^{M'}$ for every $i \geq N$, and that $\sum_{i=1}^{N-1} \delta^{i-1} v_i^M - \sum_{i=1}^{N-1} \delta^{i-1} v_i^{M'} \geq -C/2$ since $\sigma^2 \geq \underline{\sigma}^2 \geq \hat{\sigma}^2$. Because both $\sum_i \delta^{i-1} v_i^M$ and $\sum_i \delta^{i-1} v_i^{M'}$ are convergent, it suffices to identify one agent $i^* \geq N$ so that $\delta^{i^*-1} (v_{i^*}^M - v_{i^*}^{M'}) \geq C/2$. Consider the agent $i^* = \lceil \sigma^2 \rceil$, where $i^* \geq N$ since $\sigma^2 \geq \underline{\sigma}^2 \geq N$. This agent has more than $\lceil \sigma^2 \rceil \cdot (AE_M - \epsilon)$ signals with variance σ^2 , so $v_{i^*}^M$ is higher than the expected utility of $AE_M - \epsilon$ signals of variance 1. At the same time, $\lceil \sigma^2 \rceil \cdot \frac{1}{\sigma^2} \leq (1 + \epsilon)$ since $\sigma^2 \geq \underline{\sigma}^2 \geq 1/\epsilon$, so $v_{i^*}^{M'}$ is lower than the expected utility of $(AE_{M'} + \epsilon) \cdot (1 + \epsilon)$ signals of variance 1. This shows $(v_{i^*}^M - v_{i^*}^{M'}) \geq C$, and we know $\delta^{i^*-1} (v_{i^*}^M - v_{i^*}^{M'}) \geq \underline{\delta}(\sigma^2) \lceil \sigma^2 \rceil \cdot C = C/2$. \square

A.11 Details on Example 3

We want to show that

$$\lim_{K \rightarrow \infty} \frac{\sum_i v_i^M}{\sum_i v_i^{M'}} = \infty.$$

For each integer $j > 0$, Let $v^{(j)}$ be the expected utility of agent i if $r_i = jK$. Equivalently, this is the expected utility of an agent observing j independent signals with precision σ_0^2 .

On network M , for each agent in generation t we have $r_i \leq K + 3(t - 2)$ for all i by Proposition 4. Since $K + 3(t - 2) \leq 2K$ whenever $t \leq K/3 + 2$, all agents in the first

$\lfloor K/3 + 2 \rfloor$ generations have expected utility at most $v^{(2)}$. So

$$\sum_i v_i^M > K(K/3 + 1)v^{(2)}.$$

On the complete network M' , each agent in generation t has $r_i \geq (t - 1)K$, so

$$\sum_i v_i^{M'} > K \sum_{j=0}^{\infty} v^{(j)}.$$

By Lemma A.8, the sum $\sum_{j=0}^{\infty} v^{(j)}$ is convergent. So we have

$$\sum_i v_i^{M'} > -C'K$$

for some constant $C' > 0$ (which is independent of K).

Combining these bounds, we have

$$\frac{\sum_i v_i^M}{\sum_i v_i^{M'}} > \frac{K(K/3 + 1)v^{(2)}}{-C'K}.$$

The right-hand side diverges to infinity as $K \rightarrow \infty$.

A.12 Proof of Corollary 3

Proof. We claim that for any agent i in generation t , the action ℓ_i is equal to the sum of λ_i and λ_j for all agents j in generations $1, \dots, t - 1$. The proof is by induction on t . The claim holds for the first generation because all agents in the first generation choose $\ell_i = \lambda_i$.

Consider an agent in generation t . By the inductive hypothesis, she observes neighbors' actions $\ell_j = \lambda_j + \sum_{j' \leq (t-2)K} \lambda_{j'}$ for all j in generation $t - 1$ and observes s_j for one such j . Therefore, she can compute $\sum_{j' \leq (t-2)K} \lambda_{j'}$ and λ_j for all j in generation $t - 1$. Since these signals are independent and she has access to no information about other signals from her generation, she chooses $\ell_i = \lambda_i + \sum_{j \leq (t-1)K} \lambda_j$. By induction, we have $r_i = K(t-1) + 1 > i - K$ for all agents in generation t . \square

A.13 Proof of Corollary 4

Each information silo is equivalent to a maximum generations network, so the expression for r_i for agents in information silos follows immediately from Theorem 1.

The actions of agents in separate information silos are conditionally independent. For an agent in position $(t - 1)K + 1$, we have $\frac{r_{(t-1)K+1}}{t} \geq \sum_{n=1}^N \frac{2|S_n| - 1}{|S_n|}$ for t large, because that agent

observes conditionally independent actions of agents with $\lim_t \frac{r_i}{t} = \frac{2|S_n|-1}{|S_n|}$ for $1 \leq n \leq N$. On the other hand, even if agent knew all the actions and private signals of her neighbors, we would have $\frac{r^{(t-1)K+1}}{t} = \sum_{n=1}^N \frac{2|S_n|-1}{|S_n|} + o(t)$, because there a constant number of such signals. This gives an upper bound, so we conclude $\lim_{t \rightarrow \infty} \frac{r^{(t-1)K+1}}{t} = \sum_{n=1}^N \frac{2|S_n|-1}{|S_n|}$.