

Aggregative Efficiency of Bayesian Learning in Networks*

Krishna Dasaratha[†] Kevin He[‡]

First version: November 21, 2019

This version: July 23, 2021

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 a log-linear function of observations and admit a signal-counting interpretation of accuracy. This generates a fine-grained ranking of networks based on their *aggregative efficiency* index. Networks where agents observe multiple neighbors but not their common predecessors *confound* information, and we show confounding can make learning very inefficient. In a class of networks where agents move in generations and observe the previous generation, aggregative efficiency is a simple function of network parameters: increasing in observations and decreasing in confounding. Generations after the first contribute very little additional information due to confounding, even when generations are arbitrarily large.

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

*We thank Nageeb Ali, Aislinn Bohren, Matt Elliott, Mira Frick, Drew Fudenberg, Ben Golub, Sanjeev Goyal, Faruk Gul, Suraj Malladi, 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.

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

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

1 Introduction

In many economic environments, information about an unknown state of the world is dispersed among many agents. As people take actions based on their private signals and their observations of social neighbors, the process of social learning gradually aggregates their decentralized information into a group consensus.

We ask how the underlying social network influences the efficiency of this information aggregation. Understanding social-learning dynamics in different observation networks is especially relevant today, as communication technology drastically reshapes our networks. Social media platforms like Facebook and Twitter, for instance, expand our social neighborhoods far beyond the friends and family with whom we interact face-to-face.

Starting with [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#), the economic theory literature contains a large body of work on Bayesian models of sequential social learning, where privately informed individuals move in turn and draw rational inferences from their observations. Yet much of this literature has focused on settings where individuals see all predecessors or peers (i.e., the complete observation network). Less is known about how learning compares across different networks, and the existing results in this area tell us that rational agents will *eventually* learn completely in all networks within reasonable parameters.

The primary open questions concern how various social network structures affect the efficiency of signal aggregation (i.e., the rate of learning), as [Golub and Sadler \(2016\)](#)'s recent survey points out:

“A significant gap in our knowledge concerns short-run dynamics and rates of learning in these models. [...] The complexity of Bayesian updating in a network makes this difficult, but even limited results would offer a valuable contribution to the literature.”

This paper studies the impact of the social network on the efficiency of private-signal aggregation. We work with the canonical sequential social-learning model, but make two assumptions to make our analysis tractable. First, we assume the state is binary and agents have Gaussian private signals about the 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 strips away some other obstructions to efficient learning (considered by [Harel, Mossel, Strack, and Tamuz 2021](#); [Molavi, Tahbaz-Salehi, and Jadbabaie 2018](#); [Rosenberg and Vieille 2019](#) and others) and isolates the effect of the social network. Our analysis

¹The simplest example is that agents choose actions equal to their posterior beliefs given their information. This framework also applies to any other decision problem where actions fully communicate beliefs.

provides a fine-grained ranking of networks based on their efficiency for social learning, and generates rich and interpretable comparative statics results about how network parameters influence learning.

We emphasize that *informational confounds*, which appear in almost all realistic social networks, can cause nearly total information loss in social learning. Suppose an agent only observes the actions of a pair of neighbors who have both seen the action of an even earlier predecessor. From the agent’s perspective, this unobserved early action confounds the informational content of her two neighbors’ behavior, as the observation network makes it impossible to fully incorporate the neighbors’ private information without over-weighting the early mover’s private information. Even with rich actions, rational agents must solve a signal-extraction problem to decide how to draw inferences from multiple neighbors’ behavior in light of confounding. Networks differ in the severity of such informational confounds, so Bayesian social learning can vary in its efficiency and welfare properties across different networks that all lead to eventual complete learning. Our main results show that the information loss generated by confounds can be arbitrarily large.

To formalize these findings, we first describe several general properties of the social-learning model that allow us to define and calculate the efficiency of learning. The unique equilibrium of the social-learning game has a log-linear form. We characterize the equilibrium strategy profile that solves agents’ signal-extraction problems and give a procedure to compute every agent’s accuracy in any network. The equilibrium action of each agent is distributed as if she saw some (possibly non-integer) number of independent private signals.² This lets us define *aggregative efficiency* as the fraction of private signals in the society that are consolidated in the equilibrium actions. Aggregative efficiency is an index of the network that measures its efficiency for social learning.

As the main application of these general properties, we quantify 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 these networks cannot sustain much learning: even if generation sizes are large, additional generations after the first contribute very little extra information.

We first study the rate of signal aggregation in *maximal generations networks* where each agent in generation t observes the actions of all predecessors in generation $t - 1$. Society learns

²In our model, if an agent acts on $n \in \mathbb{N}_+$ independent private signals with conditional variance σ^2 , then her log-action has the distribution $\mathcal{N}(n \cdot \frac{2}{\sigma^2}, n \cdot \frac{4}{\sigma^2})$ conditional on the binary state being 1. We show that for any agent i in any network, there exists $r \in \mathbb{R}_+$ so that i ’s equilibrium log-action has the distribution $\mathcal{N}(r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2})$ conditional on the state being 1. So, i acts as if she saw “ r independent signals.”




Figure 1: **Left:** A maximal generations network with generation size $K = 3$. An arrow from i to j means i observes j '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.

completely in the long run for every generation size K , but aggregative efficiency is worse with larger K . No matter the size of the generations, social learning accumulates no more than three signals per generation starting with the third generation, and no more than two signals per generation asymptotically. Therefore, aggregative efficiency is arbitrarily close to zero when generations are large, as illustrated in Figure 1. 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.

More generally, consider any *symmetric* intergenerational observation structure — 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. We prove the same long-run bound of two signals aggregated per generations holds for all networks in this class and for all generation sizes K . An arbitrarily small fraction of available signals is included in the social consensus and agents learn arbitrarily slowly relative to the efficient rate. This failure to aggregate information reflects an inefficiency of equilibrium behavior, not an inherent limitation on information flow in the environment: we show that in any strongly connected symmetric generations network, there is a feasible (but non-equilibrium) log-linear strategy profile that eventually aggregates more than K_0 signals per generation for every $K_0 < K$.

We also compare equilibrium 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 the number of signals aggregated per generation increases in the number of neighbors for each agent and decreases in the level of confounding (i.e., the number of common neighbors for pairs of distinct agents), thus quantifying the trade-offs in changing the network. For instance, an improvement in communication tech-

nology that increases the density of the observation network may bring two countervailing effects on learning: it can speed 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.

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. Also, 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.

Aggregative efficiency also helps characterize social learning in classes of networks that do not have a generations structure. We apply our techniques to analyze learning in the canonical random network where each agent observes $d > 1$ predecessors uniformly at random. These random networks, like the symmetric generations networks, also lead to very inefficient learning: almost surely, society learns in the long run but only aggregates a vanishing fraction of the available signals. Unlike in the generations networks, inefficiency is due to physical constraints on the flow of information and not only informational confounds.

Our analysis focuses on aggregative efficiency because prior work has shown that rational agents fully learn the state *in the long run* on all networks satisfying mild conditions (Acemoglu, Dahleh, Lobel, and Ozdaglar, 2011).³ Since all “reasonable” networks lead to long-run learning, the economic questions of interest concern short-run accuracy and rates of learning. Our framework lets us address these questions, while existing techniques in the literature are designed to analyze long-run learning outcomes. Efficiency matters for welfare: agents attain any utility threshold earlier on the more efficient network (provided private signals are not too precise), and the more efficient network is strictly preferred by all sufficiently (but finitely) patient social planners.

1.1 Related Literature

We study rational⁴ learning in a sequential model (as first introduced by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)) with network observations and show in-

³When private signals and actions are coarse, Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) show rational agents can herd on incorrect beliefs. Because we allow unboundedly informative private signals and rich actions, agents learn the true state asymptotically in our framework, given mild conditions on the network.

⁴Several papers calculate speed of learning under naive updating heuristics instead of rational learning, e.g., Ellison and Fudenberg (1993) and Molavi, Tahbaz-Salehi, and Jadbabaie (2018). In the DeGroot updating model, Golub and Jackson (2012) show that speed of learning is determined by a simple network statistic that also measures the amount of homophily in the network.

complete observation networks can lead to very inefficient learning.

There is considerable interest in comparing rational learning across social networks, but fewer analytical results that distinguish learning on different incomplete networks. Several papers deliver the message that network structure matters little for social learning. [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#) and [Lobel and Sadler \(2015\)](#) show that all networks satisfying weak sufficient conditions guarantee complete long-run learning. [Rosenberg and Vieille \(2019\)](#) use a different criterion for good social learning, but come to the same conclusion that “the nature of the feedback on previous choices matters little.”

In addition, several prior examples and numerical simulations suggest network structure can affect learning considerably in finite populations.⁵ [Lobel, Acemoglu, Dahleh, and Ozdaglar \(2009\)](#) compare the rate of rational learning in two specific network structures where each agent has one neighbor. [SgROI \(2002\)](#) shows that networks with autarkic agents can improve social welfare relative to the complete network, and uses a numerical simulation to study comparative statics of welfare as a function of the number of autarkic agents. [Arieli and Mueller-Frank \(2019\)](#) consider multi-dimensional networks of different connectivity $p \in [0, 1]$. They prove networks with large $p < 1$ lead to better learning than the complete network ($p = 1$), and use numerical simulations to compare different incomplete networks. The prevalence of simulations in this literature suggests that comparing social learning on more complex networks is a relevant (if difficult) question. Our framework can analytically compare incomplete networks and networks that lead to the same long-run learning accuracy.

In another related setting, [Board and Meyer-ter-Vehn \(2021\)](#) also study the role of the social network in a continuous-time product adoption model featuring random entry times and perfectly informative private signals. They show that starting from a network where none of i 's direct neighbors share common indirect neighbors, adding links among i 's neighbors always leads to slower adoption for i . These additional links would not affect i 's learning in a sequential social-learning model, since they do not generate what we call informational confounds — that is, multiple neighbors of i learning from a common source that i does not observe.

Several strands of the social learning literature have studied other obstructions to efficient learning. [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

⁵In a different class of non-sequential social-learning models where a finite set of agents repeatedly observe their neighbors in a fixed network and simultaneously choose actions every period, [Gale and Kariv \(2003\)](#) and [Goyal \(2012\)](#) have compared learning dynamics in specific networks to highlight a possible trade-off between the accuracy of the long-run consensus and the speed of convergence to said consensus. In our setting with rich actions, this trade-off is absent as the long-run consensus is correct in all reasonable networks.

signals. The mechanism behind their result, “rational groupthink,” relies on agents’ finite action spaces. In fact, social learning in their environment would proceed at the efficient rate if actions revealed posterior beliefs, as they do in ours. The coarseness of the action space serves as the primary obstruction to the efficient rate of social learning in several other papers, including [Rosenberg and Vieille \(2019\)](#) and [Hann-Caruthers, Martynov, and Tamuz \(2018\)](#).⁶ We highlight a different mechanism for inefficient aggregation of decentralized information: an observation network that generates informational confounds can also lead to rates of learning far below the optimum even in a setting with rich actions.

Another group of papers point out that if signals about the state comes 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, 2018](#); [Lomys, 2020](#); [Liang and Mu, 2020](#)). We abstract away from this source of slow learning by giving agents exogenous signals, following most of the literature on sequential social learning. This allows us to focus on the role of the network structure on social learning.

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. 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 her belief about ω . So, she will pick the action equal to the probability she assigns to the event $\{\omega = 1\}$.

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 normal 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 $i \rightarrow j$ 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 map the structure of this network to the efficiency of information aggregation via social learning.

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

⁶Slow learning from coarse actions is also related at a conceptual level to incomplete learning with coarse actions (e.g., [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#)) and slow learning with noisy observations of others’ actions (e.g., [Vives \(1993\)](#)).

play after observing actions $a_{j(1)}, \dots, a_{j(d_i)}$ from neighbors⁷ $N(i) = \{j(1), \dots, j(d_i)\}$ and when own private signal is s_i .⁸ A Bayesian Nash equilibrium (*equilibrium* for short) is a strategy profile $(A_i^*)_{i \in \mathbb{N}_+}$ so that for all i and for all observations of i , A_i^* maximizes the Bayesian expected utility given the posterior belief about ω .⁹

The sequential nature of the social-learning game and the existence of a unique optimal action at any belief imply there is a unique equilibrium. Agent 1 has no social observations, so must use the same strategy $A_1^*(s_1)$ in all equilibria. This implies agent 2 also only has one equilibrium strategy A_2^* , as the behavior of agent 1 is unique across all equilibria. Proceeding inductively, there is a unique equilibrium profile $(A_i^*)_{i \in \mathbb{N}_+}$.

The quadratic-loss form of the utility functions is not crucial for the results, and all results remain unchanged if actions are “rich” enough to fully reflect beliefs in the following sense. Let each agent i have an arbitrary action set \hat{A}_i and utility function $\hat{u}_i : \hat{A}_i \times \{0, 1\} \rightarrow \mathbb{R}$. Suppose \hat{A}_i and \hat{u}_i are such that $\hat{a}_i^*(p) := \operatorname{argmax}_{\hat{a}_i \in \hat{A}_i} \mathbb{E}_p[\hat{u}_i(\hat{a}_i, \omega)]$ is single-valued for every $p \in (0, 1)$, where \mathbb{E}_p is the expectation under the belief that assigns p chance to $\{\omega = 1\}$. Finally, suppose that $\hat{a}_i^* : (0, 1) \rightarrow \hat{A}_i$ is one-to-one. In equilibrium, agents who have i as a neighbor can exactly infer i ’s belief using the on-path observation of i ’s action in \hat{A}_i , just as they can when i has the quadratic-loss utility function and reports own belief through the action a_i .

3 Equilibrium

3.1 Linearity of Equilibrium

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 equilibrium log-strategy: the (unique) equilibrium map between the log-actions of i ’s neighbors and i ’s own 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.

⁷We use $j(k)$ to indicate the k -th neighbor of i and suppress the dependence of j on i when no confusion arises.

⁸It is without loss for equilibrium analysis to focus on pure strategies, since every belief about the state induces a unique optimal action for each agent.

⁹We will see that in any equilibrium, $s_i \mapsto A_i^*(a_{j(1)}, \dots, a_{j(d_i)}, s_i)$ 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 in equilibrium, and therefore the posterior beliefs are well-defined.

We also show that there exist constants $(r_i)_{i \in \mathbb{N}_+}$ with $1 \leq r_i \leq i$ so that in equilibrium, (a_i, ω) is jointly distributed as if i chose a_i solely based on r_i independent private signals.¹⁰ The constants r_i depend on the network and may be interpreted as the number of signals that social learning in M aggregates by agent i . This gives a sufficient statistic to compare society's short-run accuracy in different networks.

In general, the actions of i 's neighbors are correlated even after conditioning on the state. Intuitively, agent i would like to put enough weight on the actions of her neighbors to incorporate their private signals, but doing so would also over-count the signals of the earlier agents observed by several members of $N(i)$ but not by i . The social network M thus creates an informational confound that generally prevents i from fully extracting the signals of the individuals in $N(i)$. The equilibrium strategy of i represents the optimal aggregation of her neighbors' actions. The next result shows the optimal aggregation is linear and gives an explicit expression for the coefficients. All proofs are in the Appendix.

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 equilibrium 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 for i 's neighbors' log-actions given $\omega = 1$. These coefficients do not depend on the conditional variance of the private signals $1/\sigma^2$.

The interpretation of the inverse covariance matrix in $\vec{\beta}_i$ is that i rationally discounts the actions of two neighbors $j(1)$ and $j(2)$ if their actions are conditionally correlated in equilibrium.

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](#) is the following property of the Gaussian information structure in our model, which ensures that i 's observations have a jointly Gaussian distribution conditional on ω . This permits us to study optimal inference in closed form.

¹⁰The constants r_i need not be integers, and we will formalize the meaning this claim for non-integer r_i in [Definition 1](#).

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

Proposition 1 implies that we may find weights $(w_{i,j})_{j \leq i}$ so that the realizations of equilibrium 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. Since none of the $\vec{\beta}_i$ vectors depends on σ^2 , neither does W .

Proposition 1 leads to an inductive procedure to compute the coefficients in the unique equilibrium profile and the matrix W . We start with the first row of W , $W_1 = (1, 0, 0, \dots)$. Proceeding iteratively, once the first $i - 1$ rows of W have been constructed, we know the weights that each of i 's neighbor's log-actions $\ell_{j(k)}$ puts on different log-signals, hence we can compute $\mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)}) | \omega = 1]$ and $\text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} | \omega = 1]$. We can find $\vec{\beta}_i$ using Proposition 1, and hence construct the i -th row of W .

3.2 Measure of Accuracy

We would like to evaluate networks in terms of their short-run social-learning accuracy, so as to 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 $n \in \mathbb{N}_+$ independent private signals. Then, the Bayesian i would play the log-action equal to the sum of the n log-signals, so by Lemma 1 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 equilibrium 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 a non-equilibrium strategy profile, in general 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 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 equilibrium log-actions always admit a signal-counting interpretation in any network.

Proposition 2. *In any network, every agent's behavior has a signal-counting interpretation. That is, 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.*

We can use $r_i/i \in [0, 1]$ to measure the fraction of all available signals that get incorporated into agent i 's action, with some signals lost during social learning.

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 that induce faster social learning in the long run are equivalently those with higher levels of aggregative efficiency. The limit defining aggregative efficiency need not exist in all networks, but does exist in almost all of the examples we focus on.

The signal-counting interpretation of behavior is closely identified with the rational learning rule. Even if all of i 's predecessors are rational, one can show that i 's log-action does not admit a signal-counting interpretation under “generic” log-linear strategies. Conversely, 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 $i < I$, that is i 's log-action is $\beta_{i,0}\lambda_i + \sum_{k=1}^{d_i} \beta_{i,j(k)}\ell_{j(k)}$ for any constants $(\beta_{i,j(k)})_{k=0}^{d_i}$ where $N(i) = \{j(1), \dots, j(d_i)\}$. If agent I best responds to the strategies of $i < I$, then I 's behavior has a signal-counting interpretation.*

This result provides one way to extend the definitions of r_i and aggregative efficiency to analyze the rate of social learning under any log-linear heuristic. For a given heuristic, consider a rational outside observer who has no private signal and who only sees the action of the i -th heuristic learner. It follows from Corollary 1 that this observer's log-action has the conditional distributions $\mathcal{N}\left(\pm r_i \cdot \frac{2}{\sigma^2}, r_i \cdot \frac{4}{\sigma^2}\right)$ for some r_i . Here r_i measures the informativeness of the heuristic learner i 's behavior in the units of private signals and leads to an upper-bound on i 's utility.

3.3 Long-Run Learning

Before turning to our main results about the efficiency of signal aggregation in different networks, we perform a sanity check on our model by developing 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. But the key innovation of our model is that it is tractable enough to provide a rich ranking of different network structures based on the rate of social learning, which will be the focus of the next section.

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) late enough agents have arbitrarily long observational paths, that is $\lim_{i \rightarrow \infty} \mathcal{PL}(i) = \infty$; (2) there are not infinitely many agents whose neighborhoods are contained in the same finite set, that is $\lim_{i \rightarrow \infty} [\max_{j \in N(i)} j] = \infty$; (3) social learning eventually aggregates arbitrarily many signals, that is $\lim_{i \rightarrow \infty} r_i = \infty$; (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. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) show that expanding observations is necessary and sufficient for long-run learning in a random-networks model with unboundedly informative signals and binary actions. With continuous actions, the same result is a consequence of Proposition 3.

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 very mild. It is of course possible that $\lim_{i \rightarrow \infty} r_i = \infty$ but aggregative efficiency $\lim_{i \rightarrow \infty} (r_i/i)$ is close to 0, which corresponds to a network where agents learn completely but do so very slowly. We will therefore focus on comparing how quickly $(r_i)_{i \geq 1}$ grows across different networks and the aggregative efficiency of different networks. Comparisons of aggregative efficiency also translate into welfare comparisons, as Section 6 will show.

4 Rate of Learning in Generations Networks

This section shows that informational 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*¹¹ and find that they can lead to very inefficient learning due to confounding. 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., $i = 1, \dots, K$) have no neighbors. A collection of *observation sets*, $\Psi_k \subseteq \{1, \dots, K\}$ for $k = 1, \dots, K$ define the network M for agents in later generations. 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

¹¹This 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 (2020).




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

generation). That is, for $i = (t - 1)K + k$ where $t \geq 2$ and $1 \leq k \leq K$, network M has $N(i) = \{(t - 2)K + \psi : \psi \in \Psi_k\}$.¹² Figure 2 shows an example with $K = 3$.

4.1 Full Observations and the Role of Generation Size

We first focus on the *maximal generations network* where $\Psi_k = \{1, \dots, K\}$ for all k , so agents in generation t for $t \geq 2$ have all agents in generation $t - 1$ as their neighbors.¹³ The following result relates the generation size K to the speed of signal aggregation.

Proposition 4. *Consider a maximal generations network with $K \geq 1$. We have $\lim_{i \rightarrow \infty} (r_i/i) = \frac{(2K-1)}{K^2}$, so aggregative efficiency is lower with larger K , and social learning aggregates no more than two signals per generation asymptotically for any K . Starting from the third generation, every generation aggregates at most three signals more than the previous one: for any K and any agents i, i' in generation t and $t - 1$ with $t \geq 3$, $r_i - r_{i'} \leq 3$.*

Proposition 4 contains two parts. First, it shows that even though society learns completely with any K , the aggregative efficiency is lower with higher 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 an asymptotic ranking on the speed of learning, but the same comparative statics about speed also hold numerically for all agents $i \geq 16$ when comparing among $K \in \{2, 3, 4, 5\}$, as shown in Figure 3.

Second, Proposition 4 bounds the number of signals that social learning aggregates per generation in the maximal generations network. The proof of Proposition 3 shows $r_i \geq$

¹²Stolarczyk, 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.

¹³This network is similar to the “multi-file” treatment in the laboratory experiment of Eyster, Rabin, and Weizsacker (2018), except agents only observe the actions of the immediate past generation, not those of all previous generations. In the multi-file treatment, unlike in the maximal generations network, Bayesian agents can perfectly infer the private signals of all previous movers in equilibrium.




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

$\mathcal{P}\mathcal{L}(i) + 1$ in all networks and thus provides a lower bound of 1: each generation must aggregate at least one signal, since each agent i in generation t has $\mathcal{P}\mathcal{L}(i) = t - 1$. Proposition 4 shows this lower bound is not too far from the actual learning rate. No matter how large K is, social learning aggregates fewer than two signals per generation asymptotically. There is also a short-run version of this result: starting with generation 3, fewer than three signals are aggregated per generation for any K . For K large, these bounds of two or three signals per generation constitute an arbitrarily small fraction of available signals.

This result relates to a statistical intuition that says a small number of independent signals can contain more information than an arbitrarily large number of pairwise correlated signals. In the social-learning setting, neighbors' actions serve as signals about the state of the world. The actions of these neighbors are endogenously correlated due to the structure of the maximal generations networks. The correlation is so strong that observing K neighbors' actions is less informative than just one of these actions and one additional private signal.

4.2 Partial Observations and Aggregative Efficiency

Proposition 4 shows that social learning aggregates fewer than two signals per generation asymptotically in maximal generations networks with any K . We now provide an exact expression for the aggregative efficiency in a broad class of generations networks with more general observation sets. In particular, this result will imply the same two signals per generation bound holds for all networks in this larger class.

We only impose one regularity assumption on the observation sets $(\Psi_k)_k$: symmetry.

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, E represents the prominent positions in the society, and agents only observe predecessors in these prominent positions from the past generation. The maximal generations network is the special case of $E = \{1, \dots, K\}$. 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 2 has this structure, with $d = 2$ and $c = 1$.) There remains 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). The next result applies to all such networks.

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 is*

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

with the convention $0/0 = 0$. For $c \geq 1$, the number of signals aggregated per generation is strictly increasing in d and strictly decreasing in c .

Theorem 1 gives the exact aggregative efficiency for a broader class of generations networks and quantifies the information loss due to confounding. The interpretation of the comparative statics result is that more observations speed up the rate of learning per generation but more confounding slows it down, 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 the other typically affects both d and c . Theorem 1 decomposes the repercussions of such changes on the per-generation learning rate into their effects on the two primitive network parameters that have monotonic influences on said rate.

Theorem 1 specializes to the expression for aggregative efficiency in Proposition 4 by letting $d = c = K$. In fact, the maximal generations networks lead to the slowest per-generation

rate of learning among all symmetric networks with the same number of observations per agent, since they feature the most severe confounding. Nevertheless, Theorem 1 provides a uniform learning-rate bound of two signals per generation across all symmetric networks (as $\frac{d^2-d}{d^2-d+c} \leq 1$), even though these networks may involve much less confounding than the maximal generations networks. To provide some intuition for this bound, imagine that instead of observing their predecessors, all agents in generation t observe a common set of n independent private signals, in addition to their own private signal. We can show an agent in generation $t + 1$ who observes d of these generation t predecessors puts a weight of $\frac{n+1}{dn+1}$ on each of their log-actions, and aggregates $\frac{n(2d-1)+1}{nd+1}$ more signals than they do. As $n \rightarrow \infty$, the number of extra signals aggregated approaches $\frac{2d-1}{d} \leq 2$. In any generations network for late enough t , each generation t agent's social observation constitutes a highly informative signal of the state, so it is analogous to observing n independent private signals with $n \rightarrow \infty$.

Formalizing this intuition is more subtle because different agents may observe different predecessors. This somewhat alleviates the informational confounding for generation $t + 1$, but the benefits are limited by the fact that in any symmetric network, actions in the same generation have a conditional correlation approaching 1 when $t \rightarrow \infty$. Even agents with very different observation sets end up observing highly correlated information in the long run, so no symmetric network aggregates more than two signals per generation asymptotically. To prove the correlation coefficient converges to 1, we construct a stochastic process with the state space given by the different network positions and the transition probabilities given by the observation structure. We then apply 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.

This bound of two signals per generation does not always apply to non-symmetric generations networks. Section 4.5 shows there are generations networks with K agents per generation where agents in the first position aggregate up to $K - 1$ signals per generation in the long run. Even if we restrict attention to environments with a generations structure, learning dynamics can depend greatly on the details of the social network.

Finally, we show that the aggregative efficiency in symmetric networks depends only on the generation size. An insight that extends from maximal generations networks to any symmetric networks is that aggregative efficiency is worse with larger generations. Compare the symmetric network from Figure 2 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 informational confounds in equilibrium. It turns out 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 . Therefore, the idea of worse efficiency with larger generations depicted in Figure 3 for maximal generations networks also holds in the broader class of symmetric networks.

Corollary 2. *In any symmetric network with K agents per generation, every agent observing $d \geq 2$ neighbors, and every pair of agents in the same generation sharing $c < d$ common neighbors, $\lim_{i \rightarrow \infty} (r_i/i) = (2 - (1/K)) \cdot \frac{1}{K}$.*

This corollary follows from the fact that the symmetry condition imposes some combinatorial constraints on the feasible (d, c) parameter pairs when we hold the generation size K constant. 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.

4.3 Social Planner's Benchmark

The uniformly slow speed of signal aggregation across all symmetric networks comes from an inefficiency generated by decentralized social learning, not from an inherent limitation on information flow in these networks. To illustrate this point, we study the social planner's problem and show there exists a feasible (but non-equilibrium) log-linear strategy profile such that social learning eventually aggregates more than K_0 signals per generation for every $K_0 < K$.

We study a class of networks where agents can plausibly include most predecessors' signals in their estimates.

Definition 4. The observation sets $(\Psi_k)_k$ are *strongly connected* if for every $1 \leq k_1 \leq k_2 \leq K$, there exist $t_1 < t_2$ so that there is a path from $t_2K + k_2$ to $t_1K + k_1$ in M .

This rules out the cases such as when the second agent in every generation is always excluded from the indirect neighborhood of the first agent of every future generation, which would mean agents in the first position cannot aggregate more than $K - 1$ signals per generation.

A social planner can choose a log-linear strategy that achieves nearly perfect information aggregation:

Proposition 5. *Suppose the observation sets $(\Psi_k)_k$ are strongly connected and symmetric with $c \geq 1$. There is a log-linear (but non-equilibrium) strategy profile such that, for every*

positive real number $K_0 < K$, there exists a corresponding T so that for all $t \geq T$ and $1 \leq k \leq K$, social learning aggregates more than $(t - 1)K_0$ signals by agent $(t - 1)K + k$.

As K grows large, Theorem 1 and Proposition 5 combine to say that in strongly connected and symmetric generations networks with $c \geq 1$, individuals only manage to aggregate an arbitrarily small fraction of the private signals that can be feasibly aggregated by a social planner using a log-linear strategy. The idea behind the construction is that the social planner can counteract the muddling of private signals when a group of individuals share common social observations by asking each individual to put extra weight on her own private signal in choosing her action.¹⁴

4.4 Application 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

¹⁴If non-log-linear strategies are allowed, then the social planner can achieve close to perfect information aggregation in every generation using exotic strategies that encode individuals' signals far into the decimal expansions of their actions, for example.

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 own 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 Proposition 4, 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.

4.5 Application 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*.¹⁵ 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

¹⁵We thank Suraj Malladi for suggesting this application.

specialties.¹⁶ 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.¹⁷

We use a generations network to show that information silos may benefit the organization compared with fully transparent data sharing,¹⁸ 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.

Corollary 4. *In 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, $\Psi_1 = \{2, \dots, K\}$. 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 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 4 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.

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

¹⁶Sethi 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.

¹⁷Arguments 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).

¹⁸Similar 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.




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

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 4.4](#)), 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.

5 Aggregative Efficiency of Random Networks

Much of our analysis so far has focused on networks with a generational structure, but learning can also be very inefficient on quite different networks. In this section, we suppose

the network is random and give a class of networks where, with probability one, there is long-run learning but aggregative efficiency is zero.

Formally, we consider a distribution over the set of directed networks among the agents $\{1, 2, 3, \dots\}$, i.e., those networks where each agent only has links to others that came before her. At the beginning of time, Nature draws a network M from this distribution, and we maintain the assumption that the realization of M is common knowledge among the agents. In this environment, each r_i is a random variable that depends on the network realization. We can still study the asymptotic behavior of the stochastic processes (r_i) and (r_i/i) and investigate their almost-sure limits with respect to the network realization (when these limits exist).

Consider the *fixed-degree random network*, perhaps the most canonical random network in the setting where agents arrive sequentially (see, for example, Chapter 5.1 of [Jackson, 2010](#)). Let $d \geq 2$ and consider the random network where each agent $i \leq d$ is always linked to all predecessors, but each agent $i \geq d + 1$ is linked to d different predecessors sampled uniformly at random. That is, $N(i)$ is equally likely to be any size- d subset of $\{1, \dots, i - 1\}$, and the neighborhoods of different agents are independently drawn. In the fixed-degree random network, almost all agents have out-degree d .

The fixed-degree random network with any $d \geq 2$ almost surely leads to complete but arbitrarily inefficient learning.

Proposition 6. *In any fixed-degree random network, almost surely $\lim_{i \rightarrow \infty} r_i = \infty$ and $\lim_{i \rightarrow \infty} (r_i/i) = 0$.*

An agent who arrives late enough in the sequence has arbitrarily long observation paths with probability near 1, so there is complete learning almost surely. But at the same time, we can also show that this agent’s indirect neighbors (i.e., those that are reachable through paths of any length) account for an arbitrarily small fraction of all her predecessors with probability near 1, so she cannot access a vast majority of the available signals. Unlike in generations networks, inefficient learning in fixed-degree random networks is due to physical constraints and not solely information confounding.

Intuitively, these physical constraints arise because a “typical” neighbor of an agent i has index considerably smaller than i , as neighbors are chosen uniformly at random from all predecessors. For an agent n , no matter how large n is, with high probability “long paths” from n do not stay in the interval $[\epsilon n, n]$. At the same time, we can bound the number of agents reachable through “short paths” from n , and this bound becomes a vanishingly small fraction of n as n grows. So altogether, through paths of any lengths, n can reach no more than 2ϵ fraction of the predecessors.

6 Aggregative Efficiency and Welfare Comparisons

In this section, we show that aggregative efficiency comparisons translate into two kinds of welfare comparisons.

Let $v_i := \mathbb{E}[u_i(a_i^*, \omega)]$ denote the expected equilibrium welfare of agent i , and recall that $-0.25 < v_i < 0$ for every i in any network and with any private signal precision $0 < 1/\sigma^2 < \infty$. If society learns completely in a network, then $\lim_{i \rightarrow \infty} v_i = 0$. Given a threshold level $\underline{v} \in (-0.25, 0)$ of utility, we might ask when does social learning first attain $v_i \geq \underline{v}$. We say social learning *strongly attains* \underline{v} by agent I if I is the smallest integer such that $v_i \geq \underline{v}$ for all $i \geq I$. We say social learning *weakly attains* \underline{v} by agent i if i is the earliest agent with $v_i \geq \underline{v}$ (but the expected utilities of some later agents may fall below \underline{v}).

Proposition 7. *Suppose in networks M and M' , social learning aggregates r_i and r'_i signals by agent i , respectively, with strictly ranked aggregative efficiency $\lim_{i \rightarrow \infty} (r_i/i) > \lim_{i \rightarrow \infty} (r'_i/i) > 0$. For every utility threshold $\underline{v} \in (-0.25, 0)$, M strongly attains \underline{v} strictly earlier than M' weakly attains the same threshold, provided signals are not too precise. That is, there exists a bound $\pi > 0$ on private signal precision so that whenever $0 < 1/\sigma^2 \leq \pi$, social learning strongly attains \underline{v} by agent I in M and weakly attains \underline{v} by agent i' in M' , with $I < i'$.*

Now fix the signal precision and consider the expected welfare profiles $(v_i)_{i \geq 1}$ and $(v'_i)_{i \geq 1}$ in two networks M and M' that both lead to complete social learning. A planner could compare these two profiles through a social welfare function Γ with $\Gamma(v) = \sum_{i=1}^{\infty} \gamma_i v_i + \gamma_{\infty} (\lim_{i \rightarrow \infty} v_i)$, where $\gamma_1, \gamma_2, \dots, \gamma_{\infty} \geq 0$ is a summable sequence of welfare weights that combine utilities across agents. Here γ_{∞} is the welfare weight on “the end of time,” and comparing two networks based on whether they lead to complete social learning corresponds to an “infinitely patient” Γ_{∞} with the weights $\gamma_i = 0$ for all $i \in \mathbb{N}_+$ and $\gamma_{\infty} = 1$. A social welfare function Γ_T is called *T-patient* if $\gamma_i = 0$ for all $i < T$ and $\gamma_i > 0$ for all finite $i \geq T$. That is, the planner is blind to the welfare of the first $T - 1$ agents, but strictly cares about the welfare of all later agents. One example is $\gamma_i = \delta^{i-T}$ for $i \geq T$ where the welfare of agents later than T are discounted at rate $\delta \in (0, 1)$. For large T , we can interpret a *T-patient* social welfare function as corresponding to a “very patient” but not “infinitely patient” planner. The next result implies that all very patient planners will rank M and M' based on their aggregative efficiency, even though the degenerate limiting case of the infinitely patient planner is indifferent between them.

Proposition 8. *Suppose in networks M and M' , social learning aggregates r_i and r'_i signals by agent i , respectively, with strictly ranked aggregative efficiency, $\lim_{i \rightarrow \infty} (r_i/i) > \lim_{i \rightarrow \infty} (r'_i/i)$. There exists a $\underline{T} \in \mathbb{N}_+$ such that for all $T \geq \underline{T}$, any *T-patient* social welfare function Γ_T is strictly higher on M than on M' .*

7 Conclusion

This paper presents a tractable model of sequential social learning that lets us compare social-learning dynamics across different observation networks. Generally, observation networks confound the informational content of neighbors' behavior and slow down learning. Rational agents face an optimal signal-extraction problem, whose solution takes a log-linear form in our environment. The efficiency of learning can be measured in terms of the fraction of available signals incorporated into beliefs. This allows us to make precise comparisons about the rate of learning and welfare across different networks, where additional links may trade off extra observations against the reduced informational content of each observation.

For a class of symmetric networks where agents move in generations, we derive an analytic expression for aggregative efficiency and quantify the information loss due to confounding. Additional observations speed up learning but extra confounding slows it down. Confounding severely limits the rate of signal aggregation — in any network in this class, social learning aggregates no more than two signals per generation in the long run, even for arbitrarily large generations. Inefficient aggregation is not particular to generations networks: learning is even less efficient in fixed-degree random networks, where individuals aggregate a vanishing fraction of private signals.

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 where agents have fewer observations). Studying the trade-offs and/or interactions between network-based information confounds 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. (2018): "Herding with costly information," *Journal of Economic Theory*, 175, 713–729.

- 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.
- BOARD, S. AND M. MEYER-TER-VEHN (2021): “Learning dynamics in social networks,” *Econometrica*, forthcoming.
- 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 (2020): “Learning from neighbors about a changing state,” *Working Paper*.
- ELLISON, G. AND D. FUDENBERG (1993): “Rules of thumb for social learning,” *Journal of Political Economy*, 101, 612–643.
- EYSTER, E., M. RABIN, AND G. WEIZSACKER (2018): “An experiment on social mislearning,” *Working Paper*.
- GALE, D. AND S. KARIV (2003): “Bayesian learning in social networks,” *Games and Economic Behavior*, 45, 329–346.
- 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 M. O. JACKSON (2012): “How homophily affects the speed of learning and best-response dynamics,” *Quarterly Journal of Economics*, 127, 1287–1338.
- GOLUB, B. AND E. SADLER (2016): “Learning in Social Networks,” in *The Oxford Handbook of the Economics of Networks*, Oxford University Press.
- GOYAL, S. (2012): *Connections: An Introduction to the Economics of Networks*, Princeton University Press.
- HANN-CARUTHERS, W., V. V. MARTYNOV, AND O. TAMUZ (2018): “The speed of sequential asymptotic learning,” *Journal of Economic Theory*, 173, 383–409.
- HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2021): “Rational groupthink,” *Quarterly Journal of Economics*, 136, 621–668.
- 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.
- JACKSON, M. O. (2010): *Social and Economic Networks*, Princeton University Press.
- 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. (2020): “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.
- 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.
- 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.
- TETT, G. (2015): *The Silo Effect: The Peril of Expertise and the Promise of Breaking Down Barriers*, New York: Simon and Schuster.
- VIVES, X. (1993): “How fast do rational agents learn?” *Review of Economic Studies*, 60, 329–347.
- WOLITZKY, A. (2018): “Learning from others’ outcomes,” *American Economic Review*, 108, 2763–2801.

Appendix

A Proofs

A.1 Proof of Lemma 1

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

A.2 Proof of 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 equilibrium 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 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 $\vec{\mu} = \mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)})' \mid \omega = 1]$ and $\Sigma = \text{COV}[\ell_{j(1)}, \dots, \ell_{j(d_i)} \mid \omega = 1]$.

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

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

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 equilibrium strategy must also be linear, completing the inductive step. This argument also gives the explicit form of $\vec{\beta}_{I,\cdot}$.

For the final statement, we first prove a lemma.

Lemma A.1. *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 $W_i = ((\vec{\beta}'_{i,\cdot} \times \hat{W}), 1, 0, 0, \dots)$.*

Proof. Suppose $N(i) = \{j(1), \dots, j(d_i)\}$ with $j(1) < \dots < j(d_i)$. By Lemma 1 and construction of \hat{W} , we have $\mathbb{E}[\ell_{j(k)} \mid \omega = 1] = \frac{2}{\sigma^2} \sum_{h=1}^{i-1} \hat{W}_{k,h}$. So, $\mathbb{E}[(\ell_{j(1)}, \dots, \ell_{j(d_i)}) \mid \omega = 1] = \frac{2}{\sigma^2} (\hat{W} \cdot \mathbf{1}_{(i-1)})' = \frac{2}{\sigma^2} \mathbf{1}'_{(i-1)} \hat{W}'$. Also, again by Lemma 1 and construction of \hat{W} , we can calculate that for $1 \leq k_1 \leq k_2 \leq d_i$, $\text{COV}[\ell_{j(k_1)}, \ell_{j(k_2)} \mid \omega = 1] = \frac{4}{\sigma^2} \sum_{h=1}^{i-1} (\hat{W}_{k_1,h} \hat{W}_{k_2,h})$, 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, $W_1 = (1, 0, 0, \dots)$ 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 $\vec{\beta}'_{I+1}$ 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 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 $\vec{\beta}_i$ 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.2. *For any $0 < \epsilon < 0.5$, $\mathbb{P}[a_i > 1 - \epsilon \mid \omega = 1] = 1 - \Phi\left(\frac{\ln(\frac{1-\epsilon}{\epsilon}) - r_i \frac{2}{\sigma^2}}{\sqrt{r_i \frac{2}{\sigma^2}}}\right)$, where Φ is the standard Gaussian distribution function. This expression is increasing in r_i and approaches 1. Also, $\mathbb{P}[a_i < \epsilon \mid \omega = 0] = \Phi\left(\frac{\ln(\frac{1-\epsilon}{\epsilon}) + r_i \frac{2}{\sigma^2}}{\sqrt{r_i \frac{2}{\sigma^2}}}\right)$. 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

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

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.2 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.2 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 in equilibrium, $\ell_j \sim \mathcal{N}\left(\pm r_j \cdot \frac{2}{\sigma^2}, r_j \cdot \frac{4}{\sigma^2}\right)$ 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.2 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.3. *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.4. *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.5. $\beta_t \rightarrow 1/d$.

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.3, $\text{VAR}[\ell_i \mid \omega = 1] = \frac{4}{\sigma^2} + \beta_{t+1}^2 (d \text{VAR}_t + (d^2 - d) \text{COV}_t)$. Using the definition of signal-counting interpretation and Proposition 2, $\mathbb{E}[\ell_j \mid \omega = 1] = \frac{1}{2} \text{VAR}_t$ for each $j \in N(i)$, and so $\mathbb{E}[\ell_i \mid \omega = 1] = \frac{2}{\sigma^2} + d \beta_{t+1} (\frac{1}{2} \text{VAR}_t)$. By the same argument we also have $\text{VAR}[\ell_i \mid \omega = 1] = 2 \cdot \mathbb{E}[\ell_i \mid \omega = 1]$, 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.4, 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 τ 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 :

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

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

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.5, 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 Proposition 4

Proof. Regardless of K , for each agent i in generation t , $\mathcal{PL}(i) = t-1$, so $\lim_{i \rightarrow \infty} \mathcal{PL}(i) = \infty$. By Proposition 3, society learns completely in the long run. The expression for r_i comes from specializing Theorem 1 (whose proof does not depend on Proposition 4) to the case of $d = c = K$. Observe $\frac{(2K-1)}{K^2} \cdot K = (2K-1)/K < 2$ for any $K \geq 1$.

To bound r_i starting with the 3rd generation, 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 $t-1$ in equilibrium. By the linearity of equilibrium (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.6. *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 in equilibrium. 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.6, 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 independent private signals 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.8 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.9 Proof of Proposition 5

Proof. We first construct weights such that agent i 's action need not admit a signal-counting interpretation, but has distribution close to the distribution of the action that aggregates i signals for i large. We will then rescale these weights to obtain a nearby strategy that admits a signal-counting interpretation.

Consider a Markov process with states $\{1, \dots, K\}$ and state transition probabilities $\mathbb{P}[k_1 \rightarrow k_2] = 1/|\Psi_{k_1}|$ if $k_2 \in \Psi_{k_1}$, $\mathbb{P}[k_1 \rightarrow k_2] = 0$ otherwise. (Each Ψ_k is non-empty, since the observation sets are strongly connected.) This process is irreducible by strong connectivity. Also, since the observation sets are symmetric with $c \geq 1$, the proof of Lemma A.4 implies the process is aperiodic. By standard results (see e.g., [Billingsley \(2013\)](#)), there exists a stationary distribution $\pi^* \in \mathbb{R}_{++}^K$ with $\sum_{k=1}^K \pi_k^* = 1$, such that $\lim_{\tau \rightarrow \infty} (M_\Psi)^\tau \vec{e}_k = \pi^*$ for every $1 \leq k \leq K$, where $\vec{e}_k \in \mathbb{R}^K$ is a vector with 1 in position k and 0 in other positions,

and M_Ψ is the stochastic matrix for the Markov process.

For $t \geq 1$, $1 \leq k \leq K$, abbreviate agent $i = (t-1)K + k$ as $[t, k]$. Consider the strategy profile $\vec{\beta}'_i$ where agent $i = [1, k]$ puts weight $1/\pi_k^*$ on her log-signal, while agent $i = [t, k]$ for $t \geq 2$ puts weight $\beta'_{i,j} = 1/|\Psi_k|$ on each observed log-action and weight $1/\pi_k^*$ on her log-signal. We will call the actions under this strategy profile ℓ'_i .

The weight that $[t, k]$ puts on the log-signal of $[t', k']$ for $t' < t$ is $(1/\pi_{k'}^*) \cdot ((M_\Psi)^{t-t'} \vec{e}_k)_{k'}$. Noting this quantity only depends on the difference $t - t'$ and on k, k' , we abbreviate it as $c_{t-t', k, k'}$ and observe that $\max_{k, k'} |c_{\tau, k, k'} - 1| \rightarrow 0$ as $\tau \rightarrow \infty$, since $\lim_{\tau \rightarrow \infty} (M_\Psi)^\tau \vec{e}_k = \pi^*$ for every k .

We show that under this strategy profile, ℓ'_i with $i = [t, k]$ has the conditional distributions $\mathcal{N}(\pm((t-1)K + o(i))\frac{2}{\sigma^2}, ((t-1)K + o(i))\frac{4}{\sigma^2})$. Let $\epsilon > 0$ be given. We show for all large enough $i = [t, k]$, $|\mathbb{E}[\ell'_i | \omega = 1]/(2/\sigma^2) - ((t-1)K)| < \epsilon i$. This is because there is T so that $\max_{k, k'} |c_{\tau, k, k'} - 1| < \epsilon/4$ for all $\tau \geq T$, which shows

$$|\mathbb{E}[\ell'_i | \omega = 1]/(2/\sigma^2) - ((t-1)K)| \leq (\epsilon/4)(t-1-T)K + \max_{k, k', \tau < T} |c_{\tau, k, k'} - 1| \cdot (TK) + 1/\pi_k^*.$$

Because there are finitely many values of $c_{\tau, k, k'}$ with $\tau < T$, the maximum $\max_{k, k', \tau < T} |c_{\tau, k, k'} - 1|$ is constant in i . Thus the bound is a constant term in i plus a term no larger than $(\epsilon/4) \cdot i$. By similar reasoning,

$$|\text{Var}[\ell'_i | \omega = 1]/(4/\sigma^2) - ((t-1)K)| \leq (\epsilon/2 + \epsilon^2/16)(t-1-T)K + \max_{k, k', \tau < T} |c_{\tau, k, k'}^2 - 1| \cdot (TK) + (1/\pi_k^*)^2.$$

The bound is a constant term in i plus a term no larger than $(2\epsilon/3) \cdot i$ for ϵ near 0.

We next use ℓ'_i to construct a strategy profile such that actions admit a signal-counting interpretation and continue to have the conditional distributions $\mathcal{N}(\pm((t-1)K + o(i))\frac{2}{\sigma^2}, ((t-1)K + o(i))\frac{4}{\sigma^2})$. Define

$$b_i = \frac{2\mathbb{E}[\ell'_i | \omega = 1]}{\text{Var}[\ell'_i | \omega = 1]}$$

and $\ell_i = b_i \ell'_i$, so that

$$2\mathbb{E}[\ell_i | \omega = 1] = \text{Var}[\ell_i | \omega = 1].$$

Thus, the actions ℓ_i admit a signal-counting interpretation. Because the log-actions ℓ'_i with $i = [t, k]$ have conditional distributions $\mathcal{N}(\pm((t-1)K + o(i))\frac{2}{\sigma^2}, ((t-1)K + o(i))\frac{4}{\sigma^2})$, the constants $b_i \rightarrow 1$. Therefore, the log-actions ℓ_i with $i = [t, k]$ have conditional distributions $\mathcal{N}(\pm((t-1)K + o(i))\frac{2}{\sigma^2}, ((t-1)K + o(i))\frac{4}{\sigma^2})$ as well. Therefore, given $K_0 < K$, there exists $T > 0$ such that social learning aggregates at least $(t-1)K_0$ signals by agent $(t-1)K + k$ for all $t \geq T$ and all $1 \leq k \leq K$.

Finally, we observe that because ℓ_i is a linear transformation of ℓ'_i for each i , these log-actions correspond to a feasible strategy profile for the social planner. In particular, agent i in position k puts weight $b_i \beta'_{i,j} \frac{1}{b_j}$ on ℓ_j (i.e., weight $b_i \beta'_{i,j}$ on ℓ'_j) for each neighbor j , and weight b_i / π_k^* on own log-signal. This is equal to $b_i \ell'_i$ because we know that $\ell'_i = (1/\pi_k^*) \lambda_i + \sum_{j \in N(i)} \beta'_{i,j} \ell'_j$. \square

A.10 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.11 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 Proposition 4.

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 $\lim_{t \rightarrow \infty} \frac{r_{(t-1)K+1}}{t} = \sum_{n=1}^N \frac{2|S_n|-1}{|S_n|}$.

A.12 Proof of Proposition 6

Proof. Let $\epsilon > 0$. We show that for all n sufficiently large, with sufficiently high probability, n 's indirect neighborhood contains no more than ϵn agents, so $r_n/n \leq \epsilon$.

For each n , let k_n be the largest even integer such that $d^{k_n} < \epsilon n/4$. Also, let δ_n satisfy $(1 - \delta_n)^{k_n/2} = \epsilon/2$. Note that $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. Let $A_n(\epsilon)$ be the event there is a path of length k_n from agent n to an agent in $[\epsilon n/2, n]$. We bound the probability of the event $A_n(\epsilon)$, which in turns will give a bound on the probability of the event $\{r_n/n \geq \epsilon\}$.

For a path of length k_n from agent n that ends at an agent in $[\epsilon n/2, n]$, there must exist at least $k_n/2$ agents i along the path with each such i observing an agent in $[(1 - \delta_n)i, i - 1]$. This is because the first time an agent i observes someone earlier than $(1 - \delta_n)i$ along the path, the observed neighbor has the index $(1 - \delta_n)n$ or earlier, since $i \leq n$. When this happens for a second time, the observed neighbor has index $(1 - \delta_n)^2 n$ or earlier, since the observer has an index of no larger than $(1 - \delta_n)n$, and so forth. If there were more than $k_n/2$ agents i who observe a neighbor earlier than $(1 - \delta_n)i$ along this path, then the path would end at an agent earlier than $(1 - \delta_n)^{k_n/2} n = \epsilon n/2$, by the construction of δ_n .

Because each agent's neighbors are chosen uniformly at random from the predecessors, the probability of a uniformly randomly chosen neighbor of i being earlier than $(1 - \delta_n)i$ has probability at least $1 - 2\delta_n$. Consider a stochastic process where we place a particle at n on the realized network, move the particle to a random neighbor of n , then move it again to a random neighbor of this neighbor, and so on until the particle reaches agent 1. The probability that the particle stays within $[\epsilon n/2, n]$ for at least k_n steps is at most $\binom{k_n}{k_n/2} (2\delta_n)^{k_n/2}$ — this is because by the analysis above, at least $k_n/2$ of the particle's jumps along the path must be from some i to a random neighbor no earlier than $(1 - \delta_n)i$.

There are at most d^{k_n} paths of length k_n beginning with agent n , so the expected number of such paths ending at an agent in $[\epsilon n/2, n]$ is at most $d^{k_n} \binom{k_n}{k_n/2} (2\delta_n)^{k_n/2}$. By Stirling's approximation applied to the factorials $k!$ and $(k/2)!$, there exists a constant $C_1 > 0$ so that the following holds for every δ and k :

$$\begin{aligned} d^k \binom{k}{k/2} (2\delta)^{k/2} &\leq C_1 \cdot \frac{\sqrt{2\pi k} \cdot (k/e)^k}{(\sqrt{2\pi k/2} \cdot ((k/2)/e)^{k/2})^2} \cdot (2\delta d^2)^{k/2} \\ &= \frac{C_1}{\sqrt{\pi k/2}} 2^k \cdot (2\delta d^2)^{k/2} \\ &= \frac{C_1}{\sqrt{\pi k/2}} (8\delta d^2)^{k/2}. \end{aligned}$$

By our choice of k , we have

$$k_n \geq \log_d(\epsilon n/4) - 1 \geq \log_d(\epsilon n/4) \log_d(e)/2$$

for n sufficiently large. Substituting into the exponent,

$$\frac{C_1}{\sqrt{\pi k_n/2}} (8\delta_n d^2)^{k_n/2} \leq \frac{C_1}{\sqrt{\pi k_n/2}} [(8\delta_n d^2)^{\log_d(e)/4}]^{\log_d(\epsilon n/4)}.$$

When the quantity $(8\delta_n d^2)^{\log_d(e)/4}$ is small enough, $[(8\delta_n d^2)^{\log_d(e)/4}]^{\log_d(\epsilon n/4)}$ tends to 0

faster than $1/n^2$. Since $\delta_n \rightarrow 0$ with n , there exists some N so that for $n \geq N$, the expected number of paths of length k_n from agent n to agents in $[\epsilon n/2, n]$ is smaller than $1/n^2$. The probability of the event $A_n(\epsilon)$ is bounded above by the said expectation, so for large enough n , $\mathbb{P}[A_n(\epsilon)] \leq 1/n^2$.

We show that for large enough n , outside of the event $A_n(\epsilon)$, n 's indirect neighborhood contains no more than ϵn agents. If the event $A_n(\epsilon)$ does not hold, then n does not have any path with length k_n or longer that ends at an agent in $[\epsilon n/2, n]$ (any such path with length larger than k_n contains a subpath of length k_n that ends in this region). So agent n 's indirect neighbors are either reachable by paths of length $k_n - 1$ or shorter, or else they belong to the set $[1, \epsilon n/2]$. There are at most $\sum_{j=1}^{k_n-1} d^j \leq d^{k_n}$ indirect neighbors of the first kind, and at most $\epsilon n/2$ of the second kind. By our choice of k_n , we have $d^{k_n} < \epsilon n/4$. So for n large, n has no more than ϵn indirect neighbors on the complement of the event $A_n(\epsilon)$, which means $r_n/n \leq \epsilon$. Also, for n large, $\mathbb{P}[A_n(\epsilon)] \leq 1/n^2$. By the Borel-Cantelli lemma, this implies that almost surely, $\limsup_{n \rightarrow \infty} (r_n/n) \leq \epsilon$. Since this holds for every $\epsilon > 0$, we must have the almost sure convergence of $r_n/n \rightarrow 0$.

We now show $r_i \rightarrow \infty$ for all $d \geq 2$. For each network realization, Proposition 3 shows that $r_i \rightarrow \infty$ if and only if for every k , there are only finitely many agents who have no direct neighbors outside of $\{1, \dots, k\}$. For fixed k , let E_n be the event that $N(n) \subseteq \{1, \dots, k\}$. For $n \geq k + 1$, $\mathbb{P}[E_n] \leq \frac{k}{n} \cdot \frac{k-1}{n-1}$ since $d \geq 2$. This shows $\sum_n \mathbb{P}[E_n] < \infty$, so by the Borel-Cantelli lemma, the set of networks where $N(n) \subseteq \{1, \dots, k\}$ for infinitely many n has probability 0. This result holds for every k , so except on a zero-probability set of network realizations, $r_i \rightarrow \infty$. \square

A.13 Proof of Proposition 7

We first show that expected utility is increasing in r_i .

Lemma A.7. *Agent i 's expected utility is a strictly increasing function of r_i .*

Proof. Let $r_i > r'_i \geq 1$. Consider an agent j who observes two conditionally independent Gaussian signals of the state, s_A and s_B . When $\omega = 1$, $s_A \sim \mathcal{N}(1, \sigma^2/r'_i)$ and $s_B \sim \mathcal{N}(1, \sigma^2/(r_i - r'_i))$. When $\omega = 0$, $s_A \sim \mathcal{N}(-1, \sigma^2/r'_i)$ and $s_B \sim \mathcal{N}(-1, \sigma^2/(r_i - r'_i))$. If this agent chooses an action a_j using only s_A , then the conditional distributions of the log-action are $\ell_j \sim \mathcal{N}(\pm r'_i \cdot \frac{2}{\sigma^2}, r'_i \cdot \frac{4}{\sigma^2})$. If the agent instead chooses an action a_j^* using both s_A and s_B , then the conditional distributions of the log-action are $\ell_j^* \sim \mathcal{N}(\pm r_i \cdot \frac{2}{\sigma^2}, r_i \cdot \frac{4}{\sigma^2})$, by the conditional independence of s_A and s_B . Action a_j^* gives strictly higher expected utility to j than action a_j since it is based on an extra informative signal, and this implies i has strictly higher expected utility when social learning aggregates r_i instead of r'_i signals. \square

We now prove Proposition 7.

Proof. From the hypotheses, there exist $0 < \rho_L < \rho_H$ and a finite \underline{I} so that $r_i \geq \rho_H i$ and $r'_i \leq \rho_L i$ for all $i \geq \underline{I}$. Without loss we can choose $\underline{I} > \frac{\rho_H}{\rho_H - \rho_L}$. Let $R := \max_{i \leq \underline{I}} r'_i < \infty$.

We choose π so that for any $0 < 1/\sigma^2 \leq \pi$, an agent who aggregates R signals has expected utility strictly lower than \underline{v} . To see this is possible, note that we can choose $\epsilon > 0$ small enough so that $-(1 - \epsilon)(0.5 - \epsilon)^2 < \underline{v}$. Find $\zeta > 0$ so that if $y \leq \zeta$, then $\frac{\exp(y)}{1 + \exp(y)} \leq 0.5 + \epsilon$. Suppose some agent j 's log-action has the conditional distributions $\ell_j \sim \mathcal{N}(\pm R \cdot \frac{2}{\sigma^2}, R \cdot \frac{4}{\sigma^2})$. Then $\mathbb{P}[\ell_j > \zeta \mid \omega = 1] \rightarrow 0$ as $1/\sigma^2 \rightarrow 0$, since ζ is $\frac{\sigma^2 \delta - 2R}{2R\sigma}$ standard deviations above the mean when $\omega = 1$, a quantity that tends to infinity as $\sigma \rightarrow \infty$. But whenever $\mathbb{P}[\ell_j > \zeta \mid \omega = 1] \leq \epsilon$, j 's conditional expected payoff when $\omega = 1$ is bounded above by $\mathbb{P}[a_j \leq 0.5 + \epsilon \mid \omega = 1] \cdot (-(0.5 - \epsilon)^2) \leq -(1 - \epsilon)(0.5 - \epsilon)^2$, and symmetrically the same goes for j 's conditional expected payoff when $\omega = 0$.

For a given $1/\sigma^2 \leq \pi$, let i'' be the least integer in the set $\{\underline{I} + 1, \underline{I} + 2, \dots\}$ such that $\rho_L i''$ signals lead to an expected utility of at least \underline{v} . This i'' exists since $\rho_L > 0$. Utility \underline{v} is weakly attained by no earlier than i'' in network M' . This is because M' cannot weakly attain \underline{v} before agent $\underline{I} + 1$ by construction of π , while agents $i' \geq \underline{I} + 1$ and later aggregate no more than $\rho_L i'$ signals in network M' and their utilities are strictly increasing in the number of signals aggregated by Lemma A.7. On the other hand, M strongly attains \underline{v} by no later than $I = i'' - 1$. This is because $\rho_H(i'' - 1) - \rho_L i'' = (\rho_H - \rho_L)i'' - \rho_H \geq (\rho_H - \rho_L)\underline{I} - \rho_H > 0$ by choice of \underline{I} , so $r_i \geq \rho_L i''$ for all $i \geq i'' - 1$. We again appeal to Lemma A.7 to deduce all agents $i'' - 1$ and later in M have expected utilities at least \underline{v} . \square

A.14 Proof of Proposition 8

Proof. As in the proof of Proposition 7, there exists some I so that $r_i > r'_i$ for all $i \geq I$. Now let $T = I$. Since welfare is a strictly increasing function in r by Lemma A.7, network M leads to strictly higher welfare than M' for all agents $i \geq I$. \square