Aggregative Efficiency of Bayesian Learning in Networks*

Krishna Dasaratha† Kevin He‡

First version: November 21, 2019
This version: October 29, 2020

Abstract

In social-learning settings where individuals receive private signals and observe network neighbors' actions, the network structure often obstructs information aggregation. We consider sequential social learning with rational agents and Gaussian signals and ask how the efficiency of signal aggregation changes with the network. Rational actions in our model admit a signal-counting interpretation of accuracy, which lets us compare the *aggregative efficiency* of social learning across networks. Learning is very inefficient in a class of networks where agents move in generations and observe the previous generation. Generations after the first contribute very little additional information, even when generations are arbitrarily large.

Keywords: social networks, sequential social learning, Bayesian learning

*We thank Aislinn Bohren, Mira Frick, Drew Fudenberg, Ben Golub, Sanjeev Goyal, Faruk Gul, Suraj Malladi, Pooya Molavi, Mallesh Pai, 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.

†University of Pennsylvania. 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 a society of 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) and less is known about how learning compares across different networks. The primary open questions in this area concern how various social network structures affect the *efficiency* of signal aggregation — that is, the fraction of dispersed private signals aggregated through social 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. We assume the state is binary and agents have Gaussian private signals about the state. We also 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 2020; Molavi, Tahbaz-Salehi, and Jadbabaie 2018; Rosenberg and Vieille 2019 and others) and isolates the impact of the social network on learning dynamics.

[1] The simplest example that fits this framework is when 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.

In general, the observation network creates *informational confounds* for social learning even with rich action spaces. Suppose an agent only observes the actions of a pair of neighbors who have both seen the action of an even earlier mover. 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. Rational agents solve a signal-extraction problem to decide how to optimally combine their observations and signals. Networks differ in the severity of such informational confounds, and thus Bayesian social learning can aggregate information more or less efficiently. Our main result shows that these confounds can lead to arbitrarily inefficient information aggregation.

To formalize this result, 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 represents its efficiency for social learning. Networks that lead to faster and better learning are precisely those that incorporate a larger fraction of available signals into the social consensus.

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. How well can agents learn when they only observe the actions from the recent past, but not the choices from long ago or from their contemporaries? 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 speed of learning in *maximal generations networks* where each agent in generation t observes the actions of all predecessors in generation $t-1$. Society learns 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

[2] In 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."

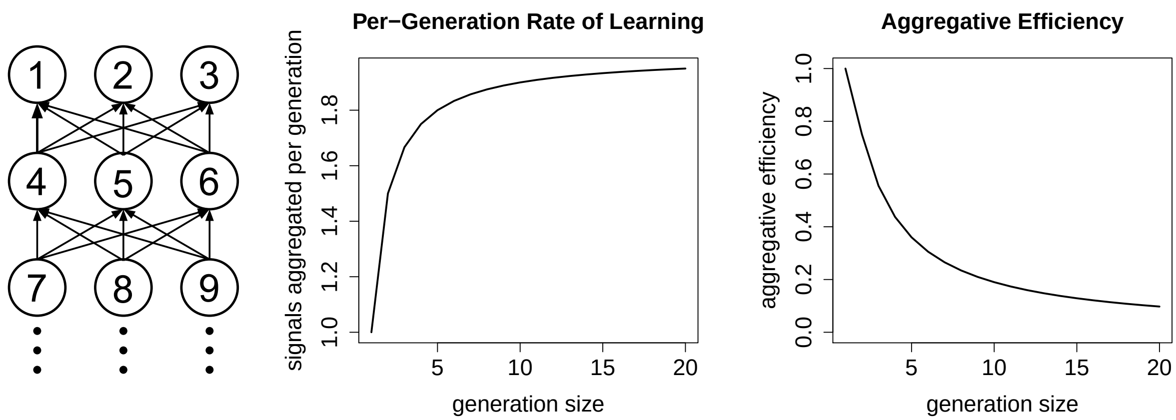

Figure 1: **Left**: A maximal generations network with generation size $K = 3$. An arrow from i to j means i observes j's action. **Middle**: Number of signals aggregated per generation asymptotically in maximal generations networks, as a function of generation size. **Right**: Aggregative efficiency in maximal generations networks, as a function of generation size.

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.

More generally, consider any *symmetric* inter-generational 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 is eventually more accurate than aggregating 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 technology 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.

Our generations network framework can provide basic insights about communication and learning in organizations. We discuss two economic applications of our results to organizational structure: (1) the value of mentorship in speeding up learning within organizations; (2) how information silos can benefit or harm different organizations. The mentorship application demonstrates that institutional structures that open up additional channels of communication, such as mentors who share their private signals with their mentees, can have large benefits for organizational learning. We also 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.

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 any pair of "reasonable" networks both 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, unlike the existing techniques in the literature that are designed to analyze long-run learning outcomes. Comparing aggregative efficiency across different networks leads to a rich analysis of learning dynamics that ranks networks based on their welfare properties for social learning. When one network has a higher aggregative efficiency than another, we show that agents attain any utility threshold earlier on the more efficient network (provided private signals are not too precise), and it is strictly preferred by all sufficiently (but finitely) patient social planners.

We extend our analysis of the generations networks using simulations that explore more general networks and non-Gaussian signal structures. We find that when agents observe multiple past generations, the number of signals aggregated per generation remains a very small fraction of the generation size. We also find that when each agent's information arrives as a large number of conditionally i.i.d. and finitely supported signals (corresponding to cohorts accumulating imprecise private information over time and then acting), then the distribution of equilibrium log-action in every generation approaches its analog when each agent has one Gaussian signal.

[3] Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) show rational agents can herd on incorrect beliefs when private signals and actions are coarse. 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.

1.1 Related Literature

We study rational learning in a sequential model (as first introduced by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)) with network observations. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) and Lobel and Sadler (2015) show that in sequential-learning environments similar to our model, rational agents learn the true state asymptotically under mild conditions on the network. We instead focus on finite-time learning accuracy and the speed of learning in different networks.[4]

Harel, Mossel, Strack, and Tamuz (2020) study a setting where a fixed group of agents repeatedly receive signals and choose actions each period, learning from each other's past actions. As in our generations network, they find that agents learn at the same rate as they would if they perfectly observed an arbitrarily small fraction of private signals. The mechanism behind their result, "rational groupthink," relies on coarse communication — agents have a finite action space and may get trapped in a wrong consensus for an extended period of time, because small changes in individual beliefs that do not lead to taking a different action are unobservable to other group members. In fact, social learning in their environment would proceed at the efficient rate if actions revealed posterior beliefs, as they do in ours. 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.

The coarseness of the action space serves as the primary obstruction to the efficient rate of social learning in several other papers. Rosenberg and Vieille (2019) consider rational sequential learning with binary actions and relate properties of the private signal distribution to whether the speed of learning achieves a particular benchmark. Hann-Caruthers, Martynov, and Tamuz (2018) compare the rates of learning from past binary actions versus past signals. By contrast, we remove any obstructions stemming from coarse communication and study network-based obstructions to achieving the efficient rate of learning and explicitly characterize this rate as a function of the network parameters in some examples, by making stronger assumptions on the informational environment.

Another group of papers point out that sequential social learning can be slow when information about the state comes from myopic agents' information-acquisition choices. In settings where agents pay for experiments and observe the actions but not the signal realizations of their predecessors, Burguet and Vives (2000) show that costly information acquisition

[4]Several 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.

slows down learning relative to exogenous signals, while Mueller-Frank and Pai (2016) and Lomys (2020) show that equilibrium learning is slower than the social planner's solution. Liang and Mu (2020) prove that slow learning obtains even in a setting where myopic agents see predecessors' signal realizations. 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 instead of the endogenous information-acquisition choices.

To the best of our knowledge, Lobel, Acemoglu, Dahleh, and Ozdaglar (2009) is the only other paper that considers how the rate of rational sequential learning varies with the observation network.[5] In a binary-actions model, they compare two specific network structures where each agent has one neighbor: either their immediate predecessor, or a random past agent drawn uniformly. We give an expression for the equilibrium accuracy of every agent in arbitrary fixed networks — in particular we allow for general neighborhood sizes. Informational confounds among social observations, the key obstacle to fast learning that we identify, only appear in networks where agents observe two or more neighbors.

Finally, Board and Meyer-ter-Vehn (2020) 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 confound — that is, multiple neighbors of i learning from a common source that i does not observe.

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*, $N(i) \subseteq \{1, ...i-1\}$. Agent i then chooses an *action* $a_i \in [0, 1]$ to maximize the expectation of

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

[5]In 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 group consensus and the speed of convergence to said consensus. In the sequential social-learning model we study, this trade-off does not appear since rational agents learn correctly in the long run in all networks satisfying mild conditions (Proposition 3).

given her belief about ω. So, she will choose 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.

Agents' neighbors are defined by a deterministic network with adjacency matrix M. We put $M_{i,j} = 1$ if $j \in N(i)$ and $M_{i,j} = 0$ otherwise. 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 $A_i : [0, 1]^{d_i} \times \mathbb{R} \to [0, 1]$, where $A_i(a_{j(1)}, ..., a_{j(d_i)}, s_i)$ specifies i's play after observing actions $a_{j(1)}, ..., a_{j(d_i)}$ from neighbors $N(i) = \{j(1), ..., j(d_i)\}$ and when own private signal is s_i.[6] Given a profile of strategies $(A_i)_{i \in \mathbb{N}_+}$, observation $(a_{j(1)}, ..., a_{j(d_i)}, s_i)$ is *on-path* if it has positive density under the profile. A perfect Bayesian equilibrium (*equilibrium* for short) is a strategy profile $(A_i^*)_{i \in \mathbb{N}_+}$ so that for all i and for all on-path observations of i, A_i^* maximizes the Bayesian expected utility given the (well-defined) posterior belief about ω. We will see that in any equilibrium, $s_i \mapsto A_i^*(a_{j(1)}, ..., a_{j(d_i)}, s_i)$ is a surjective function onto $(0, 1)$ for all i and $a_{j(1)}, ..., a_{j(d_i)}$. So an observation is on-path in equilibrium if and only if all observed actions are interior.

The sequential nature of the social-learning game implies there is a unique equilibrium. Agent 1 who has no social observations must use the same strategy $A_1^*(s_1)$ in all equilibria. So 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 the analysis is 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 $\hat{u}_i : \hat{A}_i \times \{0, 1\} \to \mathbb{R}$. Suppose $\hat{A}_i$ and $\hat{u}_i$ are such that $\hat{a}_i^*(p) := \text{argmax}_{\hat{a}_i \in \hat{A}_i} \mathbb{E}_p[\hat{u}_i(\hat{a}_i, \omega)]$ 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 $\hat{a}_i^* : [0, 1] \to \hat{A}_i$ 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.

[6] It is without loss for equilibrium analysis to focus on pure strategies, since agents are never indifferent between two actions in equilibrium.

3 Equilibrium

3.1 Linearity of Equilibrium

We will find it convenient to work with the following log-transformations of variables: $\tilde{s}_i := \ln\left(\frac{\mathbb{P}[\omega=1|s_i]}{\mathbb{P}[\omega=0|s_i]}\right)$, $\tilde{a}_i := \ln\left(\frac{a_i}{1-a_i}\right)$. We call $\tilde{s}_i$ the *log-signal* of i and $\tilde{a}_i$ the *log-action* of i. These changes are bijective, so it is without loss to use the log versions. Write $\tilde{A}_i^*(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)}, \tilde{s}_i)$ 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 $\tilde{A}_i^*$ is a linear function of its arguments, with coefficients that only depend on the network M and not on the precision of private signals. 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 asif i chose a_i solely based on r_i independent private signals.[7] 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, 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 $N(i) = \{j(1), ..., j(d_i)\}$, there exist constants $(\beta_{i,j(k)})_{k=1}^{d_i}$ so that i's equilibrium log-strategy is given by*

$$\tilde{A}_i^*(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)}, \tilde{s}_i) = \tilde{s}_i + \sum_{k=1}^{d_i} \beta_{i,j(k)}\tilde{a}_{j(k)}.$$

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

$$\vec{\beta}_i = 2\left(\mathbb{E}[(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)}) \mid \omega = 1] \times \mathrm{Cov}[\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)} \mid \omega = 1]^{-1}\right).$$

These coefficients do not depend on the conditional variance of the private signals $1/\sigma^2$.

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

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.

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

Proposition 1 implies that we may find weights $(w_{i,j})_{j \le i}$ so that the realizations of equilibrium log-actions are related to the realizations of log-signals by $\tilde{a}_i = \sum_{j=1}^{i} w_{i,j}\tilde{s}_j$. 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, $W_1 = (1, 0, 0, ...)$. 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 $\tilde{a}_{j(k)}$ puts on different log-signals, hence we can compute $\mathbb{E}[(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)}) \mid \omega = 1]$ and $\mathrm{COV}[\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)} \mid \omega = 1]$. 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 $\tilde{a}_i \sim \mathcal{N}\left(\pm n \cdot \frac{2}{\sigma^2}, n \cdot \frac{4}{\sigma^2}\right)$, 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.

Definition 1. Social learning *aggregates $r \in \mathbb{R}_+$ signals by agent i* if the equilibrium log-action $\tilde{a}_i$ 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 $\tilde{a}_i$ 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 profile is log-linear. Indeed, if this profile results in i putting weights $(w_{i,j})_{j \leq i}$ on log-signals $(\tilde{s}_j)_{j \leq i}$, then $\tilde{a}_i$ 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. *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 due to informational confounding.

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

The 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 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 $\beta_{i,0}\tilde{s}_i + \sum_{k=1}^{d_i} \beta_{i,j(k)}\tilde{a}_{j(k)}$ for any constants $(\beta_{i,j(k)})_{k=0}^{d_i}$ where $N(i) = \{j(1), ..., j(d_i)\}$. If agent I plays the best response against 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 results about finite-time accuracy, we develop two equivalent necessary and sufficient conditions for long-run learning in our setting. We say society *learns completely in the long run* if (a_i) converges to ω in probability. For a given network M, write $\overline{PL}(i) \in \mathbb{N}$ to refer to the length of the longest path in M originating from i (this length is 0 if $N(i) = \varnothing$).

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

Condition (1) of Proposition 3 says society learns completely in the long run if and only late enough agents have arbitrarily long observational paths. In fact, the proof of the result shows $r_i \geq \overline{PL}(i) + 1$ in all networks. 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 rich signals and binary actions. With continuous actions, the same result is a consequence of Proposition 2.

Proposition 3 tells us that whether society learns in the long run is not a useful criterion for comparing different networks in this setting, as the conditions that guarantee long-run learning are very mild. We will instead focus on comparing $(r_i)_{i\geq 1}$ and aggregative efficiency across different networks. These comparisons of aggregative efficiency also translate into welfare comparisons, as Section 5 will show.

4 Rate of Learning in Generations Networks

As an application of Section 3's characterization results, we study the speed of rational learning in *generations networks*.[8] We show that learning can be very inefficient in generations networks and 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., $i = 1, ..., K$) have no neighbors. A collection of *observation sets*, $\Psi_k \subseteq \{1, ..., K\}$ for $k = 1, ..., K$ 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

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

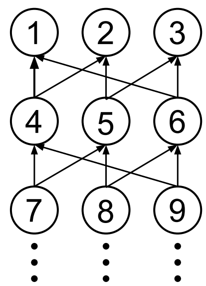

Figure 2: A generations network with $K = 3$ agents per generation and the observation sets $\Psi_1 = \{1, 2\}$, $\Psi_2 = \{2, 3\}$, and $\Psi_3 = \{1, 3\}$.

generation). That is, for $i = (t-1)K + k$ where $t \geq 2$ and $1 \leq k \leq K$, network M has $N(i) = \{(t-2)K + \psi : \psi \in \Psi_k\}$.[9] 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 $\Psi_k = \{1, ..., K\}$ for all k, so agents in generation t for $t \geq 2$ have all agents in generation $t-1$ as their neighbors.[10] The following result relates the generation size K to the speed of signal aggregation.

Proposition 4. *In the maximal generations network with any* $K \geq 1$, *society learns completely in the long run. We have* $\lim_{i\to\infty}(r_i/i) = \frac{(2K-1)}{K^2}$, *so aggregative efficiency is lower with larger* K, *and social learning aggregates no more than two signals per generation asymptotically for any* K. *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 $r_i \geq$

[9] Stolarczyk, 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.

[10] This 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.

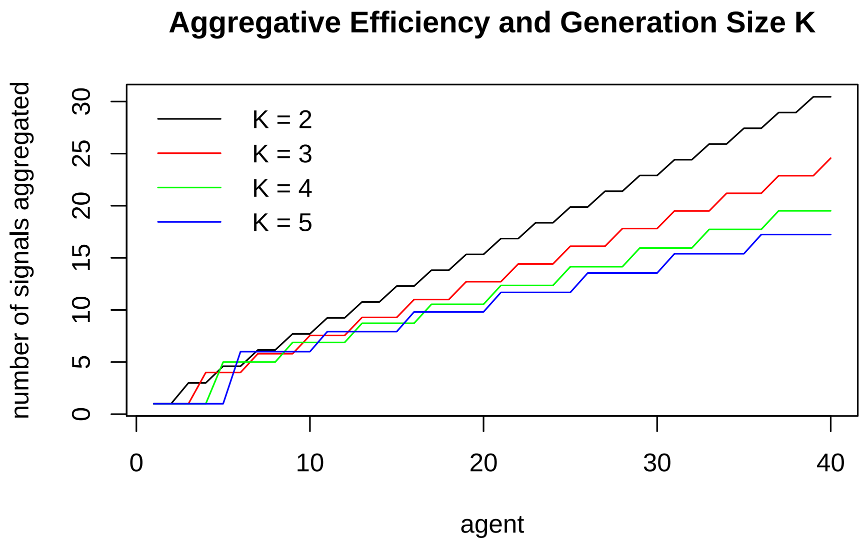

Figure 3: Number of signals aggregated by social learning in maximal generations networks with different generation sizes, $K \in \{2, 3, 4, 5\}$.

$\overline{PL}(i) + 1$ 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 $\overline{PL}(i) = t-1$. 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.

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. $|N(i)| = d$ for all $i > K$ and $|N(i_1) \cap N(i_2)| = c$ whenever $i_1 = (t-1)K + k_1$ and $i_2 = (t-1)K + k_2$ for some $t \geq 2$ and $1 \leq k_1 < k_2 \leq K$ distinct.

A generations network defined by symmetric observation sets is called a *symmetric network*. To give a class of examples symmetric networks, fix any non-empty subset $E \subseteq \{1, ..., K\}$, and let $(\Psi_k)_k$ be such that for all $1 \le k \le K$, $\Psi_k = E$. 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 $E = \{1, ..., K\}$. For another class of examples, suppose $K \ge 2$ and each agent observes a different subset of $K - 1$ predecessors from the previous generation. Specifically, $\Psi_k = \{1, ..., K\} \backslash \{k - 1\}$ for $2 \le k \le K$, and $\Psi_1 = \{1, ..., K - 1\}$. 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 $3 \le d \le 41$ and $1 \le c \le d - 2$ 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. *Suppose the observation sets $(\Psi_k)_k$ are symmetric, with every agent observing d neighbors and every pair of agents in the same generation sharing c common neighbors. Then, with the convention $0/0 = 0$,*

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

Theorem 1 gives the exact aggregative efficiency for a broader class of generations networks and quantifies the information loss due to confounding. Provided $c \ge 1$, the number of signals aggregated per generation (i.e., the term in the parenthesis in Theorem 1) is strictly increasing in d and strictly decreasing in c, with the interpretation 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 $n \to \infty$, 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 $n \to \infty$. But different agents may observe different predecessors. This somewhat alleviates the informational confounding for generation $t+1$, but is limited by the fact that in any symmetric network, actions in the same generation have a conditional correlation approaching 1 when $t \to \infty$. 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.

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

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 agents are asymptotically more accurate than aggregating K_0 signals per generation for every $K_0 < K$.

We consider a slightly more restricted class of networks.

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 $t_1 K + k_1$ is connected to $t_2 K + k_2$ 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.

We introduce a new measure of accuracy. Agent n's action is *more accurate than r signals* if $\mathbb{P}[a_n > 0.5 \mid \omega = 1] > \mathbb{P}[A_r > 0.5 \mid \omega = 1]$ and $\mathbb{P}[a_n < 0.5 \mid \omega = 0] > \mathbb{P}[A_r < 0.5 \mid \omega = 0]$, where the log transform of A_r has conditional distributions $\tilde{A}_r \sim \mathcal{N}(\pm r \cdot \frac{2}{\sigma^2}, r \cdot \frac{4}{\sigma^2})$ in the two states. That is, n's action is more likely to lean towards the correct state than the action of someone who observes r independent signals. While this definition applies even for non-equilibrium strategies that do not lead to $\tilde{a}_n$ having the conditional distributions $\mathcal{N}(\pm r_n \cdot \frac{2}{\sigma^2}, r_n \cdot \frac{4}{\sigma^2})$, if some such r_n existed then the definition would be equivalent to $r_n > r$.

Proposition 5. *Suppose the observation sets $(\Psi_k)_k$ are strongly connected and symmetric with $c \geq 1$. There is a log-linear 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 $1 \leq k \leq K$, the action of agent $(t-1)K + k$ is more accurate than $(t-1)K_0$ signals.*

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.[11]

[11] If 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.

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 any member of the previous generation. Then $r_i > i - K$ for all i, and therefore the aggregative efficiency is 1.*

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

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

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

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

If each mentor instead generates a new, independent private signal for their mentee, rather than sharing the realization of their 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*.[12] These silos have a number of causes: compensation structures that discourage collaboration between teams, different subunits storing information in mutually incompatible databases, or technical language barriers that stop ideas from flowing between specialties.[13] 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.[14]

We use a generations network to show that information silos may benefit the organization compared with fully transparent data sharing,[15] when the organization's success primarily depends on the actions of a few executives who can observe and process the behavior in all

[12] We thank Suraj Malladi for suggesting this application.

[13] Sethi 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.

[14] Arguments 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).

[15] Similar 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.

the silos.

Corollary 4. *In a generations network with $K \geq 2$ agents per generation, suppose $\{2, ..., K\}$ are partitioned into $N \geq 1$ silos $S_1, ..., S_N$ where $\Psi_k = S_k$ for $k \in S_k$, and $\Psi_1 = \{2, ..., K\}$. We have $\lim_{t\to\infty} \frac{r_{(t-1)K+1}}{t} = \sum_{n=1}^{N} \frac{2|S_n|-1}{|S_n|}$, while $\lim_{t\to\infty} \frac{r_{(t-1)K+k}}{t} = \frac{2|S_n|-1}{|S_n|} < \frac{2K-1}{K}$ for $1 \leq n \leq N$, $k \in S_n$.*

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.

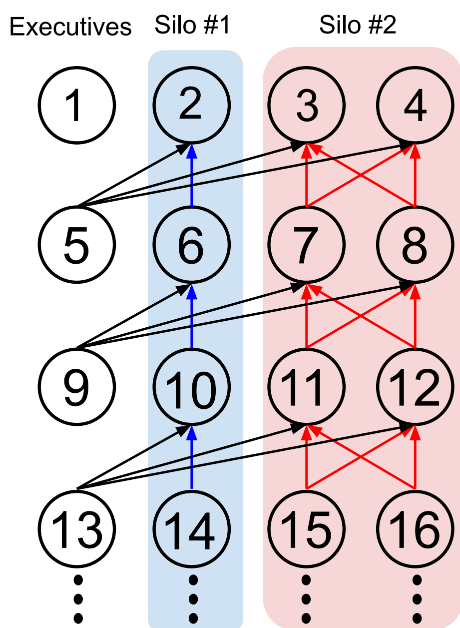

Figure 4: A generations network with executives in the first position and two information silos, $S_1 = \{2\}$ and $S_2 = \{3, 4\}$.

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

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

An important qualification is that if employees communicate private signals along with actions (perhaps as in the application in Section 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 and Welfare Comparisons Across Networks

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 $\lim_{i\to\infty} v_i = 0$. 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 $\underline{v}$).

The next result shows that when signals are not too precise, a network with a higher aggregative efficiency strongly attains any such utility threshold strictly earlier than a network with lower aggregative efficiency weakly attains the same threshold.

Proposition 6. *Suppose in networks M and M', social learning aggregates $(r_i)_{i\geq 1}$ and $(r_i')_{i\geq 1}$ signals by agent i, respectively, with $\lim_{i\to\infty}(r_i/i) > \lim_{i\to\infty}(r_i'/i) > 0$. For every utility threshold $\underline{v} \in (-0.25, 0)$, 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 Λ with $\Lambda(v) = \sum_{i=1}^{\infty} \lambda_i v_i + \lambda_\infty(\lim_{i\to\infty} v_i)$, where $\lambda_1, \lambda_2, ..., \lambda_\infty \geq 0$ is a summable sequence of welfare weights that combine utilities across agents. Here λ_∞ 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" Λ_∞ with the weights $\lambda_i = 0$ for all $i \in \mathbb{N}_+$ and $\lambda_\infty = 1$. A social welfare function Λ_T is called T*-patient* if $\lambda_i = 0$ for all $i < T$ and $\lambda_i > 0$ 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 $\lambda_i = \delta^{i-T}$ 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 7. *Suppose in networks* M *and* $M^{'}$, *social learning aggregates* $(r_i)_{i\geq 1}$ *and* $(r_i^{'})_{i\geq 1}$ *signals by agent* i, *respectively, with* $\lim_{i\to\infty}(r_i/i) > \lim_{i\to\infty}(r_i^{'}/i)$. *There exists a* $\underline{T} \in \mathbb{N}_+$ *such that for all* $T \geq \underline{T}$, *any* T*-patient social welfare function* Λ_T *is strictly higher on* M *than on* $M^{'}$.

6 Simulations

In this section, we provide simulation evidence on the robustness of our findings from Section 4. We first relax the assumption that agents only observe the previous generation, and then consider non-Gaussian signal structures.

6.1 Observing Multiple Past Generations

We first consider an extension to the maximal generations network, allowing observations of multiple previous generations. Suppose that each agents in generations t observes all members of generations $t-\tau, \ldots, t-1$ for some $\tau \geq 1$. When $\tau \geq 2$, the actions of generations $t-\tau, \ldots, t-2$ can be used to remove some of the confounds in generation $t-1$'s actions. A natural question is how much this extra information improves aggregative efficiency.

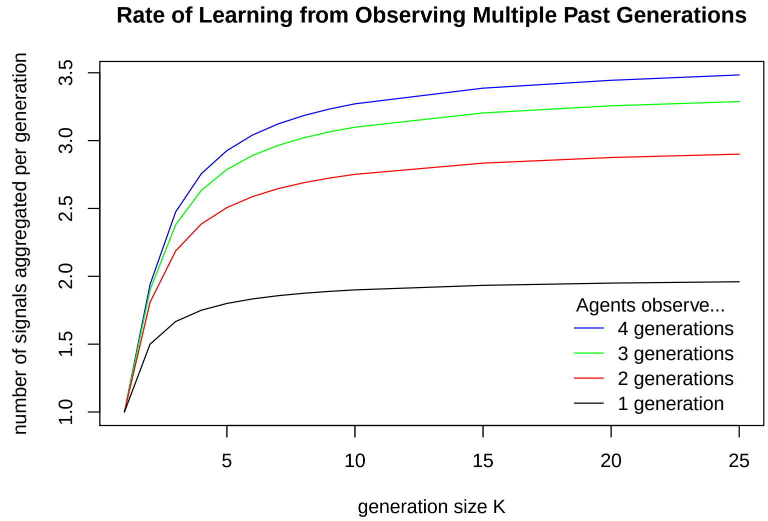

Figure 5: Number of signals aggregated by social learning in maximal generations networks when agents observe τ previous generations, $\tau \in \{1, 2, 3, 4\}$.

Figure 5 shows that the value of observing additional past generations is limited, and the marginal increase in aggregative efficiency for each extra generation of observations quickly diminishes. The figure shows simulations with 1850 generations and plots the rate of learning for the final generations.[16] The outcomes are plotted with number of generations observed $\tau \in \{1, 2, 3, 4\}$ and generation size K between 1 and 25.[17] While social learning can aggregate more than two signals per generation when $\tau \geq 2$, the rate of learning remains small relative to the generation size for larger values of K. While observing additional generations removes some confounds in the previous generations' actions, it also introduces new confounds. For instance, if $\tau = 2$, then generation 3 acts based on both generation 2 and generation 1's behavior and perfectly aggregates signals from all previous generations. But, all future generations $t \geq 4$ face confounding when interpreting the actions from generation $t-1$, since they do not observe behavior from generation $t-3$.

6.2 General Signal Structures

Next, we discuss a simulation that suggests our theoretical model involving agents with Gaussian private signals also provides useful insights about environments with other signal structures.

Consider a maximal generations network with two agents per generation, but let each

[16]When agents observe τ past generations, we report $\max_{1850-\tau+1 \leq t \leq 1850}\{r_{[t]} - r_{[t-1]}\}$ where $r_{[t]}$ is the number of signals aggregated in generation t. This is an upper bound on the per-generation improvement in accuracy for the final τ generations.

[17]The number of generations and the generation sizes are chosen based on computational constraints.

agent receive her private information as n i.i.d. signals. This corresponds to generations that gather information over a period of time, so each agent's private belief comes from combining a large number of imprecise private signals about the state.

We scale down the informativeness of each private signal as we increase the number of signals per agent. Each signal has four possible realizations: it may be strong or weak, and it has a value of 0 or 1. There is z_n chance of each signal being strong. If a signal is strong, then its value matches the state with probability $0.5 + x_n$. If a signal is weak, its value matches the state with probability $0.5 + y_n$. We choose the parameters x_n, y_n, z_n such that the conditional mean of the log-likelihood based on a single private signal is $0.2/n$ and the conditional variance of this log-likelihood is $0.4/n$. Thus, the conditional mean and conditional variance of each agent's overall log-likelihood ratio based on all private signals is independent of n.

Figure 6 shows that the distribution of log-actions in generation 4 quickly converges to a Gaussian distribution as n increases from 1 to 4. Moreover, the conditional mean and conditional variance converge to the corresponding values under our theoretical analysis with Gaussian signals (see Table 1). This suggests that as the log-likelihood based one each agents' private information converges to the Gaussian log-likelihood (by the central limit theorem), actions under social learning converge to the actions with Gaussian signals. In particular, as $n \to \infty$, the action distribution for agent i is approximately the same as $\mathcal{N}\left(\pm r_i \cdot \frac{2}{\sigma^2}, r_i \cdot \frac{4}{\sigma^2}\right)$, where r_i is the number of signals aggregated by social learning in our framework with Gaussian signals.

Further simulations suggest that even without increasing the number of signals per agent, the actions of later agents approach a Gaussian distribution as social learning aggregates many conditionally i.i.d. signals from different agents. For example, a sample of 1000 generation 4 log-actions when each person has one signal (i.e., sample drawn from the first histogram in Figure 6) will be rejected by the Shapiro-Wilk normality test at the $p < 0.05$ level with probability close to 1. But, a sample of 1000 generation 10 log-actions in the same environment is only rejected at the $p < 0.05$ level 6% of the time. This indicates that a straightforward modification of our techniques may give a useful approximation for behavior after the first few generations, even when the log-signals and hence the initial log-actions are far from Gaussian.

1 signal		10 signals		100 signals		1000 signals		Gaussian theory	
Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
1.315	1.629	1.201	1.543	1.207	1.550	1.224	1.575	1.232	1.570

Table 1: Mean and standard deviation of log-actions in generation 4 with $1, 10, 100$, and 1000 finitely-supported signals (simulations) and Gaussian signals (theoretical analysis).

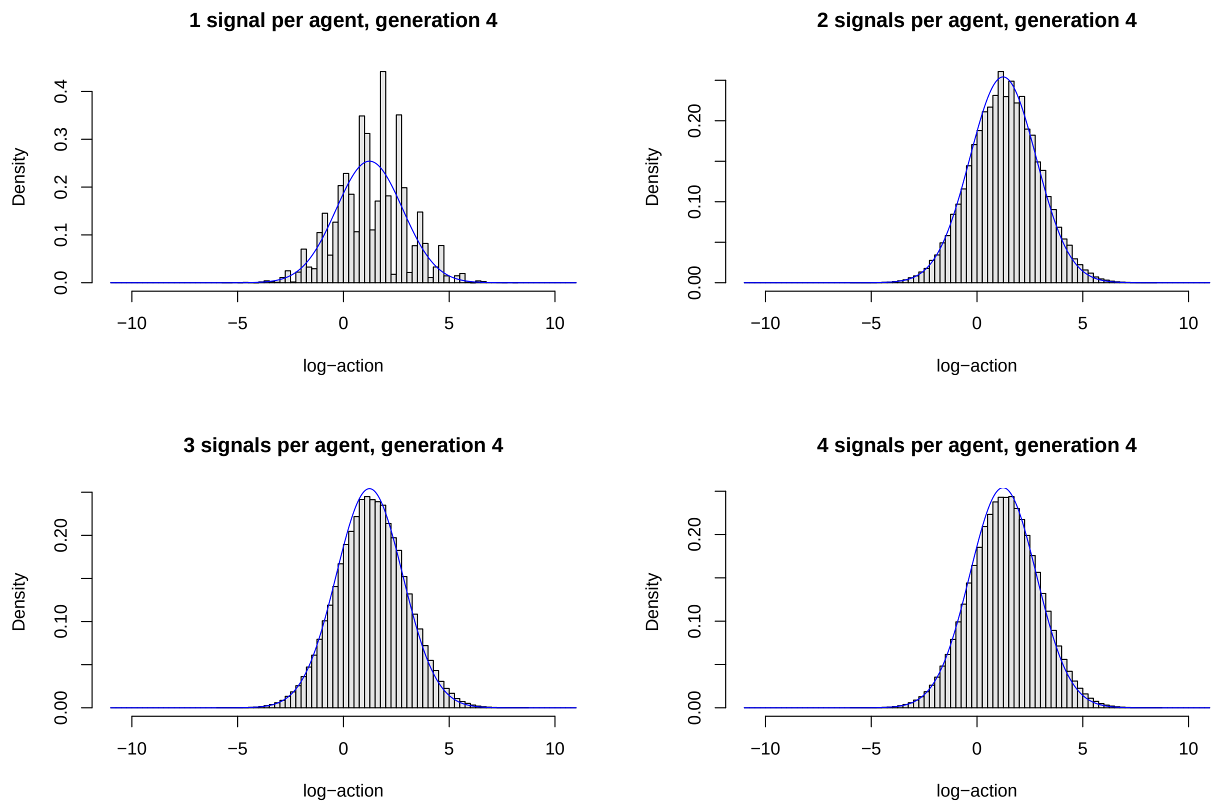

Figure 6: Probability mass functions of log-action conditional on $\omega = 1$ in generation 4 when each agent receives 1, 2, 3, or 4 private signals, each with four possible realizations. Blue curves show the distributions of log-actions when agents have Gaussian private signals with $\sigma^2 = 10$.

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. For a class of symmetric networks where agents move in generations, 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.

We derive an analytic expression of the aggregative efficiency in all such networks and quantify the information loss due to confounding. We then show 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.

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 (e.g., Ali (2018)), 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.

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 (2020): "Learning dynamics in social networks," *Working Paper*.

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.

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 (2020): "Rational groupthink," *Quarterly Journal of Economics, forthcoming.*

Immorlica, N., J. Mao, A. Slivkins, and Z. S. Wu (2020): "Incentivizing Exploration with Selective Data Disclosure," in *Proceedings of the 21st ACM Conference on Economics and Computation*, New York, NY, USA: Association for Computing Machinery, EC '20, 647–648.

Liang, A. and X. Mu (2020): "Complementary information and learning traps," *Quarterly Journal of Economics*, 135, 389–448.

Lobel, I., D. Acemoglu, M. Dahleh, and A. Ozdaglar (2009): "Rate of convergence of learning in social networks," *Proceedings of the American Control Conference*.

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

Lomys, N. (2020): "Collective search in networks," *Working Paper*.

Mathon, R. and A. Rosa (1985): "Tables of parameters of BIBDs with r<=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 Soc.

Sethi, R. and M. Yildiz (2016): "Communication with unknown perspectives," *Econometrica*, 84, 2029–2069.

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.

Wolitzky, A. (2018): "Learning from others' outcomes," *American Economic Review*, 108, 2763–2801.

A Appendix: Proofs

A.1 Proof of Lemma 1

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

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

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

A.2 Proof of Proposition 1

Proof. Agent 1 does not observe any predecessors, so clearly $\tilde{A}_1^*(\tilde{s}_1) = \tilde{s}_1$. Suppose by way of induction that the equilibrium strategies of all agents $j \leq I - 1$ are linear. Then each $\tilde{a}_j$ for $j \leq I - 1$ is a linear combination of $(\tilde{s}_\ell)_{\ell=1}^I$, 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 $(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_I)})$ have a conditional joint Gaussian distribution with $(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_I)}) \sim \mathcal{N}(\vec{\mu}, \Sigma)$ conditional on $\omega = 1$, and t $(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_I)}) \sim \mathcal{N}(-\vec{\mu}, \Sigma)$ conditional on $\omega = 0$, where $\vec{\mu} = \mathbb{E}[(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)})' \mid \omega = 1]$ and $\Sigma = \text{Cov}[\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)} \mid \omega = 1]$.

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

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

which is $2\left(\vec{\mu}'\Sigma^{-1}\right) \cdot (\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_I)})'$ 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 $\{1, ..., i-1\}$. Then $\vec{\beta}_i = \mathbf{1}'_{(i-1)} \times \hat{W}'(\hat{W}\hat{W}')^{-1}$ and the i-th row of W is $W_i = \left((\vec{\beta}'_{i,\cdot} \times \hat{W}), 1, 0, 0, ...\right)$.*

Proof. Suppose $N(i) = \{j(1), ..., j(d_i)\}$ with $j(1) < ... < j(d_i)$. By Lemma 1 and construction of $\hat{W}$, we have $\mathbb{E}[\tilde{a}_{j(k)} \mid \omega = 1] = \frac{2}{\sigma^2}\sum_{\ell=1}^{i-1}\hat{W}_{k,\ell}$. So, $\mathbb{E}[(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)}) \mid \omega = 1] = \frac{2}{\sigma^2}(\hat{W} \cdot \mathbf{1}_{(i-1)})' = \frac{2}{\sigma^2}\mathbf{1}'_{(i-1)}\hat{W}'$. Also, again by Lemma 1 and construction of $\hat{W}$, we can calculate that for $1 \le k_1 \le k_2 \le d_i$, $\text{Cov}[\tilde{a}_{j(k_1)}, \tilde{a}_{j(k_2)} \mid \omega = 1] = \frac{4}{\sigma^2}\sum_{\ell=1}^{i-1}(\hat{W}_{k_1,\ell}\hat{W}_{k_2,\ell})$, meaning $\text{Cov}[\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)} \mid \omega = 1] = \frac{4}{\sigma^2}\hat{W}\hat{W}'$. It then follows from what we have shown above that $\vec{\beta}_{i,\cdot} = 2 \cdot \frac{2}{\sigma^2}\mathbf{1}'_{(i-1)}\hat{W}' \times \left[\frac{4}{\sigma^2}\hat{W}\hat{W}'\right]^{-1} = \vec{\mathbf{1}}'_{(i-1)} \times \hat{W}'(\hat{W}\hat{W}')^{-1}$.

Since i puts weight 1 on $\tilde{s}_i$ and weights $\vec{\beta}_{i,\cdot}$ on $(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(d_i)})' = \hat{W} \times (\tilde{s}_1, ..., \tilde{s}_{i-1})'$, 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. □

To prove the final statement of Proposition 1, $W_1 = (1, 0, 0, ...)$ 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 $i \le I$. If $\hat{W}$ is the submatrix of W with rows $N(I+1)$, then since $N(I+1) \subseteq \{1, ..., I\}$, by the inductive hypothesis $\hat{W}$ must be independent of σ^2. Thus the same independence also applies to $\vec{\beta}_{I+1,\cdot}$ since this vector only depends on $\hat{W}$ by the result just derived. In turn, since W_{I+1} is only a function of $\vec{\beta}'_{I+1,\cdot}$ and $\hat{W}$, and these terms are independent of σ^2 as argued before, same goes for W_{I+1}, completing the inductive step. □

A.3 Proof of Proposition 2

Proof. It suffices to show that $\mathbb{E}[\tilde{a}_i \mid \omega = 1] = \frac{1}{2}\text{Var}\left[\tilde{a}_i \mid \omega = 1\right]$. By Proposition 1, $\tilde{a}_i = \tilde{s}_i + \sum_{k=1}^{d_i}\beta_{i,j(k)}\tilde{a}_{j(k)}$. From Lemma 1, we have $\mathbb{E}[\tilde{s}_i \mid \omega = 1] = \frac{1}{2}\text{Var}\left[\tilde{s}_i \mid \omega = 1\right]$. Furthermore, $\tilde{s}_i$ is independent from $\sum_{k=1}^{d_i}\beta_{i,j(k)}\tilde{a}_{j(k)}$, as the latter term only depends on $\tilde{s}_1, ..., \tilde{s}_{i-1}$. So we need only show $\mathbb{E}[\sum_{k=1}^{d_i}\beta_{i,j(k)}\tilde{a}_{j(k)} \mid \omega = 1] = \frac{1}{2}\text{Var}\left[\sum_{k=1}^{d_i}\beta_{i,j(k)}\tilde{a}_{j(k)} \mid \omega = 1\right]$

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

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

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

A.4 Proof of Corollary 1

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

A.5 Proof of Proposition 3

We first state and prove an auxiliary lemma.

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

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

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

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

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 society learns completely in the long run if and only if $\lim_{i\to\infty} r_i = \infty$. Let $\epsilon^{'} > 0$ be given and suppose $\lim_{i\to\infty} r_i = \infty$. Putting $\epsilon = \min(\epsilon^{'}, 0.4)$, we get that $\mathbb{P}[|a_i - \omega| < \epsilon \mid \omega = 1] \to 1$ and $\mathbb{P}[|a_i - \omega| < \epsilon \mid \omega = 0] \to 1$ since the two expressions in Lemma A.2 increase in r_i and approach 1, hence also $\mathbb{P}[|a_i - \omega| < \epsilon^{'}] \to 1$. So society learns completely in the long run. Conversely, if $r_i < K < \infty$ for infinitely many i, then 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 necessary and sufficient conditions for $\lim_{i\to\infty} r_i = \infty$.

Condition (1): $\lim_{i\to\infty}\overline{PL}(i) = \infty$.

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

Sufficiency: First note if $j \in N(i)$, then $r_i \geq r_j + 1$. This is because in equilibrium, $\tilde{a}_j \sim \mathcal{N}\left(\pm r_j \cdot \frac{2}{\sigma^2}, r_j \cdot \frac{4}{\sigma^2}\right)$ conditional on the two states, and furthermore $\tilde{a}_j$ is conditionally independent of s_i. So, $\tilde{a}_j + \tilde{s}_i$ 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 $\tilde{a}_i = \tilde{a}_j + \tilde{s}_i$ instead, since it follows from Lemma A.2 that a log-action that aggregates more signals leads to higher expected payoffs.

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

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

Sufficiency: Construct an increasing sequence $C_1 \leq C_2 \leq ...$ 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, $\overline{PL}(i) \geq 1$ for all $i \geq C_1$. Suppose $C_1 \leq ... \leq C_n$ are constructed with the property that $\overline{PL}(i) \geq k$ for all $i \geq C_k$, $k = 1, ..., n$. 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 $\overline{PL}(j) \geq n$ by the inductive hypothesis, all $i \geq C_{n+1}$ must have $\overline{PL}(i) \geq n+1$, completing the inductive step. This shows $\lim_{i\to\infty} \overline{PL}(i) = \infty$. By the sufficiency of Condition (1) for $\lim_{i\to\infty} r_i = \infty$, we see that Condition (2) implies the same. □

A.6 Proof of Theorem 1

Proof. If $d = 1$, then exactly one signal is aggregated per generation so $r_i/K \to 1$ 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,* $\textsc{Var}[\tilde{a}_i \mid \omega = 1]$ *and* $\textsc{Cov}[\tilde{a}_i, \tilde{a}_{i'} \mid \omega = 1]$ *depend only on t and not on the identities of i or i', which we call* $\textsc{Var}_t$ *and* $\textsc{Cov}_t$*, 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 $\textsc{Var}[\tilde{a}_j \mid \omega = 1]$ is the same for all $j \in N(i)$,$\mathbb{E}[\tilde{a}_j \mid \omega = 1]$ is the same for all $j \in N(i)$ (by using Proposition 2, and all pairs $j, j' \in N(i)$ with $j \neq j'$ have the same conditional covariance. Thus by Proposition 1, i places the same weight, say β_t, on all neighbors. □

So we have

$$\text{Var}[\tilde{a}_i \mid \omega = 1] = \frac{4}{\sigma^2} + \beta_t^2(d\text{Var}_{t-1} + (d^2 - d)\text{Cov}_{t-1})$$

for all i in generation t, and

$$\text{Cov}[\tilde{a}_i, \tilde{a}_{i'} \mid \omega = 1] = \beta_t^2(c\text{Var}_{t-1} + (d^2 - c)\text{Cov}_{t-1})$$

for all agents $i \neq i^{'}$ in generation t. This shows the claims for $t = T + 1$.

Taking the difference of the two expressions for Var_t and Cov_t gives:

$$\text{Var}_t - \text{Cov}_t = \frac{4}{\sigma^2} + \beta_t^2(d - c)(\text{Var}_{t-1} - \text{Cov}_{t-1}). \tag{1}$$

We now require two auxiliary lemmas.

Lemma A.4. *Consider the Markov chain on $\{1, ..., K\}$ with state transition matrix p, with $p_{i,j} = \mathbb{P}[i \to j] = 1/d$ if $j \in \Psi_i$, 0 otherwise. Suppose $(\Psi_k)_k$ is symmetric with $c \geq 1$. Then $p_i^\infty := \lim_{t\to\infty}(p^t)_i \in [0, 1]^K$ 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, $C_1, ..., C_L \subseteq \{1, ..., K\}$. 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 $i \in C_\ell$ for $1 \leq \ell \leq L^{'}$. Let $\Psi_i = \{j_1, ..., j_d\}$. If $i \in \Psi_i$, then i's periodicity is 1. Otherwise, $\Psi_i \subseteq C_\ell$ since C_ℓ 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 $i \to j_h \to ... \to i$. 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, $i \to j_1 \to j_{h^*} \to ... \to i$. 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 $\nu_\ell^*, 1 \leq \ell \leq L^{'}$, so that $\lim_{t\to\infty}(p^t)_i = \nu_\ell^*$ whenever $i \in C_\ell$. If $i \notin \cup_{1\leq\ell\leq L^{'}} C_\ell$, then starting the process at i, almost surely the process enters one of the closed communication classes eventually. This shows $\lim_{t\to\infty}(p^t)_i$ exists and is equal to $\sum_{\ell=1}^{L^{'}} q_\ell \nu_\ell^*$, where q_ℓ is the probability that the process started at i enters C_ℓ 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 $\Psi_i = \{n_1, ..., n_c, n_{c+1}, ..., n_d\}$, $\Psi_j = \{n_1, ..., n_c, n^{'}_{c+1}, ..., n^{'}_d\}$

where all $n_1, ..., n_d, n'_{c+1}, ..., n'_d \in \{1, ..., K\}$ 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. *$\beta_t \to 1/d$.*

Proof. By Proposition 2, we can compute that:

$$\beta_{t+1} = \frac{\textsc{Var}_t}{\textsc{Var}_t + (d-1)\textsc{Cov}_t} \geq \frac{1}{d}.$$

It is therefore sufficient to show that $\textsc{Var}_t/\textsc{Cov}_t \to 1$. 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 $\{1, \ldots, K\}$ and state transition probabilities $\mathbb{P}[k_1 \to k_2] = 1/d$ if $k_2 \in \Psi_{k_1}$, $\mathbb{P}[k_1 \to k_2] = 0$. 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:

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

We want to show that

$$\textsc{Var}_t/\textsc{Cov}_t = \frac{1+\sum_{\tau=1}^{t-1}\sum_{k=1}^{K} w_{i,(t-\tau)K+k}^2}{\sum_{\tau=1}^{t-1}\sum_{k=1}^{K} w_{i,(t-\tau)K+k}w_{i',(t-\tau)K+k}} \to 1.$$

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 $t-\tau_0+1, \ldots, t$ is bounded and therefore lower order. Thus,

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

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 $\textsc{Var}_t - \textsc{Cov}_t \leq \varphi^{(t-T)}(\textsc{Var}_T - \textsc{Cov}_T)$, so this shows

$$\limsup_{t\to\infty} (\textsc{Var}_t - \textsc{Cov}_t) \leq \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2 - (1+\epsilon)^2 d + (1+\epsilon)^2 c}$$

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

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

$$\liminf_{t\to\infty} (\textsc{Var}_t - \textsc{Cov}_t) \geq \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2 - d + c}$$

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

Using Proposition 2, we have $\textsc{Var}_{t+1} = 2(\beta_{t+1}d(\textsc{Var}_t/2) + 2/\sigma^2)$, so

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

Using $\lim_{t\to\infty}(\textsc{Var}_t - \textsc{Cov}_t) = \frac{4}{\sigma^2} \cdot \frac{d^2}{d^2-d+c}$, we conclude

$$\lim_{t\to\infty} (\textsc{Var}_{t+1} - \textsc{Var}_t) = \lim_{t\to\infty} \left(\frac{\textsc{Var}_t}{\textsc{Var}_t - \frac{4}{\sigma^2}\frac{d^2-d}{d^2-d+c}} - 1 \right) \textsc{Var}_t + \frac{4}{\sigma^2}.$$

Since $\textsc{Var}_t \to \infty$, we get $\lim_{t\to\infty} \left(\frac{\textsc{Var}_t}{\textsc{Var}_t - \frac{4}{\sigma^2}\frac{d^2-d}{d^2-d+c}} - 1 \right) \textsc{Var}_t = \frac{4}{\sigma^2}\frac{d^2-d}{d^2-d+c}$ using Taylor expansion. So $\lim_{t\to\infty} (\textsc{Var}_{t+1} - \textsc{Var}_t) = \frac{4}{\sigma^2}\left(\frac{d^2-d}{d^2-d+c} + 1\right)$, implying $r_i = \left(1 + \frac{d^2-d}{d^2-d+c}\right)\frac{i}{K} + o(i)$. So $\lim_{i\to\infty}(r_i/i) = \left(1 + \frac{d^2-d}{d^2-d+c}\right)\frac{1}{K}$. $\square$

A.7 Proof of Proposition 4

Proof. Regardless of K, for each agent i in generation t, $\overline{PL}(i) = t-1$, so $\lim_{i\to\infty} \overline{PL}(i) = \infty$. 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 $\tilde{a}_{\text{sum}}$ 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 $\mu_{\text{sum}}, \sigma^2_{\text{sum}} > 0$ so that the conditional distributions of $\tilde{a}_{\text{sum}}$ in the two states are $\mathcal{N}(\pm\mu_{\text{sum}}, \sigma^2_{\text{sum}})$.

Lemma A.6. *Each element in $\vec{\beta}_{i,\cdot}$ is $\left(\frac{\mu^2_{sum}}{\sigma^2_{sum}} + \frac{1}{\sigma^2}\right) / \left(K\frac{\mu^2_{sum}}{\sigma^2_{sum}} + \frac{1}{\sigma^2}\right)$.*

Proof. An application of Proposition 1 shows each agent j in generation t aggregates $\tilde{a}_{\text{sum}}$ and own private signal $\tilde{s}_j$ according to $\tilde{a}_j = 2 \cdot \frac{\mu_{\text{sum}}}{\sigma^2_{\text{sum}}}\tilde{a}_{\text{sum}} + \tilde{s}_j$.

Next, consider the problem of someone in generation $t+1$ who observes the log-actions $\tilde{a}_j$ 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

$$\mathbb{E}\left[\sum_{k=1}^{K} \tilde{a}_{(t-1)K+k} \mid \omega = 1\right] = 2K\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + 2K\frac{1}{\sigma^2}$$

$$\text{Var}\left[\sum_{k=1}^{K} \tilde{a}_{(t-1)K+k} \mid \omega = 1\right] = K \cdot \left(4 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + 4 \cdot \frac{1}{\sigma^2}\right) + K \cdot (K-1) \cdot 4 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2}$$

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 $x_{old} > 0$ so that $\tilde{a}_i \sim \mathcal{N}(\pm x_{old}, 2x_{old})$ conditional on the two states. In fact, from Proposition 1, $x_{old} = 2 \cdot \frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{2}{\sigma^2}$. 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 $x_{new} = \frac{2K(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2})^2}{K\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}} + \frac{2}{\sigma^2}$.

A hypothetical agent who observes $\tilde{a}_{\text{sum}}$ (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

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

We must have $\mathbb{P}[\tilde{a}_{\text{sum}} > 0 \mid \omega = 1] \geq \mathbb{P}[\tilde{s}_1 > 0 \mid \omega = 1]$, 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 $x_{new} - x_{old} \leq 3 \cdot \frac{2}{\sigma^2}$. $\square$

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, $\lim_{i\to\infty}(r_i/i) = \left(1 + \frac{d^2-d}{d^2-d+c}\right)\frac{1}{K} = \left(2 - \frac{c}{d^2-d+c}\right)\frac{1}{K} = (2 - \frac{1}{K}) \cdot \frac{1}{K}$. $\square$

A.9 Proof of Proposition 5

Proof. Consider a Markov process with states $\{1, ..., K\}$ and state transition probabilities $\mathbb{P}[k_1 \to k_2] = 1/|\Psi_{k_1}|$ if $k_2 \in \Psi_{k_1}$, $\mathbb{P}[k_1 \to k_2] = 0$ 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 $\lim_{\tau\to\infty}(M_\Psi)^\tau \vec{e}_k = \pi^*$ 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 where agent $[1, k]$ puts weight $1/\pi^*_k$ on her log-signal, while agent $[t, k]$ for $t \geq 2$ puts weight $1/|\Psi_k|$ on each observed log-action and weight $1/\pi^*_k$ on her log-signal. 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 $\max_{k,k'} |c_{\tau,k,k'} - 1| \to 0$ as $\tau \to \infty$, since $\lim_{\tau\to\infty}(M_\Psi)^\tau \vec{e}_k = \pi^*$ for every k.

We show that under this strategy profile, $\tilde{a}_i$ 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 and we show for all large enough $i = [t, k]$, $|\mathbb{E}[\tilde{a}_i \mid \omega = 1]/(2/\sigma^2) - ((t-1)K)| < \epsilon i$. 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

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

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,

$$|\mathrm{Var}[\tilde{a}_i \mid \omega = 1]/(4/\sigma^2) - ((t-1)K)| \leq (\epsilon/2 + \epsilon^2/16)(t-1-T)K + \max_{k,k',\tau<T} |c^2_{\tau,k,k'} - 1| \cdot (TK) + (1/\pi^*_k)^2.$$

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

Let $K_0 < K$ be given. If $\tilde{A}$ has the conditional distributions $\mathcal{N}(\pm(t-1)K_0 \cdot \frac{2}{\sigma^2}, (t-1)K_0 \cdot \frac{4}{\sigma^2})$ in the two states, then $\mathbb{P}[A > \frac{1}{2} \mid \omega = 1] = \Phi(\sqrt{(t-1)K_0}/\sigma)$. Pick some

$\epsilon > 0$ so that $\frac{K-\epsilon}{\sqrt{K+\epsilon}} > \sqrt{K_0}$. There corresponds T so that for for $i = [t,k]$ with $t \geq T$ and $1 \leq k \leq K$, $\mathbb{E}[\tilde{a}_i \mid \omega = 1] \geq (t-1)(K-\epsilon)\frac{2}{\sigma^2}$ and $\mathrm{Var}[\tilde{a}_i \mid \omega = 1] \leq (t-1)(K+\epsilon)\frac{4}{\sigma^2}$, so $\mathbb{P}[a_i > \frac{1}{2} \mid \omega = 1] \geq \Phi(\sqrt{(t-1)} \cdot (K-\epsilon)/(\sqrt{K+\epsilon}\sigma))$, so i is more accurate than $(t-1)K_0$ signals. □

A.10 Proof of Corollary 3

Proof. We claim that for any agent i in generation t, the action $\tilde{a}_i$ is equal to the sum of $\tilde{s}_i$ and $\tilde{s}_j$ for all agents j in generations $1, \ldots, t-1$. The proof is by induction on t. The claim holds for the first generation because all agents in the first generation choose $\tilde{a}_i = \tilde{s}_i$.

Consider an agent in generation t. By the inductive hypothesis, she observes neighbors' actions $\tilde{a}_j = \tilde{s}_j + \sum_{j' \leq (t-2)K} \tilde{s}_{j'}$ for all j in generation $t-1$ and observes s_j for one such j. Therefore, she can compute $\sum_{j' \leq (t-2)K} \tilde{s}_{j'}$ and $\tilde{s}_j$ 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 $\tilde{a}_i = \tilde{s}_i + \sum_{j \leq (t-1)K} \tilde{s}_j$. By induction, we have $r_i = K(t-1)+1 > i-K$ for all agents in generation t. □

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

A.12 Proof of Proposition 6

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 $\tilde{a}_j \sim \mathcal{N}(\pm r_i^{'} \cdot \frac{2}{\sigma^2}, r_i^{'} \cdot \frac{4}{\sigma^2})$. 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 $\tilde{a}_j^* \sim \mathcal{N}(\pm r_i \cdot \frac{2}{\sigma^2}, r_i \cdot \frac{4}{\sigma^2})$, 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 6.

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 $\tilde{a}_j \sim \mathcal{N}(\pm R \cdot \frac{2}{\sigma^2}, R \cdot \frac{4}{\sigma^2})$. Then $\mathbb{P}[\tilde{a}_j > \zeta \mid \omega = 1] \to 0$ as $1/\sigma^2 \to 0$, 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 $\sigma \to \infty$. But whenever $\mathbb{P}[\tilde{a}_j > \zeta \mid \omega = 1] \leq \epsilon$, 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 $\{\underline{I}+1, \underline{I}+2, ...\}$ 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.13 Proof of Proposition 7

Proof. As in the proof of Proposition 6, 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 $i \geq I$. $\square$