Linear Social Learning in Networks with Rational Agents*

Krishna Dasaratha† Kevin He‡

First version: November 21, 2019
This version: February 11, 2020

Abstract

We consider a sequential social-learning environment with rational agents and Gaussian private signals, focusing on how the observation network affects the speed of learning. Agents learn about a binary state and take turns choosing actions based on own signals and network neighbors' behavior. Equilibrium learning may be slow when agents do not observe all predecessors, as agents compromise between incorporating the signals of the observed neighbors and not over-counting the confounding signals of the unobserved early movers. We show that on any network, equilibrium actions are a log-linear function of observations and each agent's accuracy admits a signal-counting interpretation. Adding links to the observation network can harm agents even without introducing new confounds. We then consider a network structure where agents move in generations and observe some members of the previous generation. When this observation structure is sufficiently symmetric, the additional information aggregated by each generation is asymptotically equivalent to fewer than two independent signals, even when generations are arbitrarily large. When agents observe all predecessors from the previous generation, social learning aggregates no more than three signals per generation starting from the third generation, and the long-run learning rate is slower when generations are larger.

*We thank Drew Fudenberg, Ben Golub, Matthew Rabin, Philipp Strack, Omer Tamuz, and various seminar participants for useful comments.

†Harvard University. Email: krishnadasaratha@gmail.com

‡California Institute of Technology and 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 agents take actions based on their private signals and their observations of social neighbors, the process of social learning gradually aggregates decentralized information into a group consensus.

How does the underlying social network influence how quickly this aggregation happens? This question about the speed of social learning carries important welfare implications. Even if two social networks both lead to the correct group consensus in the long run, agents might quickly achieve high confidence in the correct state with high probability in one network, but remain almost fully uncertain about the state for a very long time in the other.

The economic theory literature contains a large body of work on Bayesian models of sequential social learning, where privately informed agents move in turn and draw rational inferences from their observations. These papers have largely focused on long-run learning outcomes, and less is known about how the social network affects the rate of learning. As Golub and Sadler (2016)'s recent survey points out:

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

The present paper investigates the role of the social network on the rate of rational sequential learning. We avoid some of the "complexity" that Golub and Sadler (2016) mention by working with a tractable model that allows us to study the speed of rational learning in different networks but abstracts away from other "obstructions" to efficient learning. We assume the state is binary and agents have Gaussian private signals about the state. We also suppose that agents have rich actions, so players infer their neighbors' beliefs perfectly through their behavior. This rich-signals, rich-actions world strips away other sources of learning rate inefficiency[1] and isolates the role of the social network.

In general, the observation network generates informational confounds for social learning even with rich action spaces. Suppose an agent 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 the two neighbors' behavior, as the observation network makes it impossible to fully incorporate the two neighbors' private information without over-weighting the early mover's private information. Rational agents

[1]See the literature review below for papers dealing with how these other mechanisms affect the speed of learning.

solve an optimal signal-extraction problem to decide how to aggregate their observations and signals. Networks differ in the severity of such informational confounds, and thus lead to different rates of Bayesian social learning.

We show that 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 the finite-agent accuracy of social learning. The equilibrium action of each agent is distributed as if she sees some number of independent private signals. This signal-equivalence property characterizes action distributions up to a single parameter measuring the accuracy of beliefs, which simplifies the analysis.

We then apply our general results to study the speed of learning in the "maximal generations network" where agents are sequentially arranged into generations of size K, with each agent in generation t only observing the actions of every predecessor in generation $t-1$. Society learns completely in the long run for every K, but the speed of learning is eventually slower with larger K. This implies that when private signals are imprecise, accuracy thresholds (e.g., the first agent who is at least 90% confident about the true state at least 90% of the time) are achieved earlier in networks with smaller generations. We also show that no matter the size of the generations, social learning aggregates no more than three signals per generation starting with the third generation, and no more than two signals per generation asymptotically.

Actions are very confounded under the maximal generations network, but in fact the asymptotic learning rate is bounded uniformly even under more general network structures that imply less confounding of actions. If agents are arranged into generations of arbitrarily large size K but observe some subset of their generation $t-1$ predecessors, then the same long-run bound of two signals per generations holds whenever the observation structure is weakly aperiodic and strongly regular (i.e., all agents observe the same number of neighbors and all pairs of distinct agents in the same generation share the same number of neighbors). This result highlights the confounding effect of equilibrium play, as there exists a feasible (but non-equilibrium) strategy profile that is eventually more accurate than aggregating K_0 signals per generation for every $K_0 < K$.

For all networks (not necessarily with the generations structure), we show that society learns completely in the long run if and only if late enough agents have arbitrarily long *observational paths*. As a result, adding links to the observation network can only (weakly) improve long-run learning outcome. But, the same is not true for the rate of learning. In a special class of networks without confounding, adding links speeds up learning and improves every agent's accuracy. In general, however, agents can become less accurate in networks with additional links, even when those new links do not introduce new *intransitivities* into the network. Extra observations can harm agents, even without creating any additional confounds.

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.

Harel, Mossel, Strack, and Tamuz (2019) study a setting where a fixed group of agents repeatedly receive signals and choose actions each period, learning from each others' past actions.[2] Like in our generations network, they find that the rate of learning can be equivalent to perfectly observing 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 would proceed at the efficient rate if actions were rich. 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 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 study network-based obstructions to achieving the efficient rate of learning and characterize this rate asymptotically in some examples, by making stronger assumptions on the informational environment.[3]

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. In a binary-actions model, they compare two specific network

[2] When agents receive signals only once but then act repeatedly, Gale and Kariv (2003) show that learning occurs in finite time. Indeed, the number of periods needed is quite small in several examples.

[3] Liang and Mu (2020) consider a different obstruction to the efficient rate of learning: myopic agents who choose from multiple informational sources, unlike our agents who get signals from a fixed distribution.

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 on arbitrary fixed networks.

Several papers calculate speed of learning under naive updating heuristics instead of rational learning, e.g., Ellison and Fudenberg (1993). 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.

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

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

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.

With the network M fixed, let $n_i := |N(i)|$ denote the number of i's neighbors. A *strategy* for agent i is a function $A_i : [0, 1]^{n_i} \times \mathbb{R} \to [0, 1]$, where $A_i(a_{j(1)}, ..., a_{j(n_i)}, s_i)$ specifies i's play after observing actions $a_{j(1)}, ..., a_{j(n_i)}$ from neighbors $N(i) = \{j(1), ..., j(n_i)\}$ and when own private signal is s_i.[4] Given a profile of strategies $(A_i)_{i \in \mathbb{N}_+}$, observation $(a_{j(1)}, ..., a_{j(n_i|)}, s_i)$ 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, $s_i \mapsto A_i^*(a_{j(1)}, ..., a_{j(n_i)}, s_i)$ is a surjective function onto $(0, 1)$ for all i and $a_{j(1)}, ..., a_{j(n_i)}$. So an observation is on-path in equilibrium if and

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

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 $(A_i^*)_{i\in\mathbb{N}_+}$.

3 Linearity of Equilibrium and Measure of Accuracy

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 will 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 $\tilde{A}_i^*(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_i)}, \tilde{s}_i)$ as 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 $r_i \leq i$ so that in equilibrium, (a_i, ω) is jointly distributed as-if i chooses a_i solely based on r_i independent private signals.[5] The constants r_i depend on the network and may be interpreted as the number of signals that social learning on M aggregates by agent i. This gives a convenient language to compare society's short-run accuracy on different networks.

In general, the behavior 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 carries the risk of over-counting the signals of 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 $N(i)$. 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 $N(i) = \{j(1), ..., j(n_i)\}$, there exist constants $(\beta_{i,j(k)})_{k=1}^{n_i}$ so that*

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

[5]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 vector of coefficients $\vec{\beta}_{i,\cdot}$ *is given by*

$$\vec{\beta}_{i,\cdot} = 2\left(\mathbb{E}[(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_i)}) \mid \omega = 1] \cdot \text{Cov}[\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_i)} \mid \omega = 1]^{-1}\right).$$

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, $\tilde{s}_i$ *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* $\text{Var}[\tilde{s}_i \mid \omega = 0] = \text{Var}[\tilde{s}_i \mid \omega = 1] = 4/\sigma^2$.

Proposition 1 implies that we may find weights $(w_{i,j})_{i \geq j}$ 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. The next result relates $\vec{\beta}_{i,\cdot}$ and W and shows that both are independent of the private signal precision.

Proposition 2. *Let* $\hat{W}$ *be the submatrix of* W *with rows* $N(i)$ *and columns* $\{1, ..., i-1\}$. *Then* $\vec{\beta}_{i,\cdot} = \vec{\mathbf{1}}'_{(i-1)} \cdot \hat{W}'(\hat{W}\hat{W}')^{-1}$ *and the i-th row of* W *is* $W_i = \left((\vec{\beta}'_{i,\cdot} \cdot \hat{W}), 1, 0, 0, ...\right)$. *In particular, neither* W *nor* $\vec{\beta}_{i,\cdot}$ *depends on* σ^2.

Proposition 2 provides an inductive procedure to compute the coefficients in the unique equilibrium profile and the matrix W. We start with the first row of W, $W_1 = (1, 0, 0, ...)$. Proposition 2 gives an expression for the equilibrium coefficients $\vec{\beta}_{2,\cdot}$ of agent 2 and the weight she puts on different log-signals, W_2. This in turn lets us calculate the coefficients and log-signal weights for agent 3, $\vec{\beta}_{3,\cdot}$ and W_3, and so forth. Equilibrium behavior and log-signal weights depend on the neighborhoods defined by the network M, but not on private signal precision.

We would like to evaluate networks in terms of their short-run social-learning accuracy, so as to compare the rates of Bayesian learning on different networks. Towards a measure of accuracy, imagine that agent i observes $n \in \mathbb{N}_+$ independent private signals, but gets no other information about ω. 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 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 conditional on $\omega = 1$ and $\omega = 0$ respectively. Our concept of the speed of learning exploits the fact that on any network, the equilibrium log-action distribution of every agent must follow a generalized version of this distribution with a possibly 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.

When agents use a non-equilibrium strategy profile, 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. Indeed, if this profile results in i putting weights $(w_{i,j})_{j \le i}$ on log-signals $(\tilde{s}_j)_{j \le i}$, then $\tilde{a}_i$ has conditional distributions of the desired form 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 this kind of signal-counting interpretation.

Proposition 3. *There exist* $(r_i)_{i \in \mathbb{N}_+}$ *so that social learning aggregates* r_i *signals by agent* i*. Here* $(r_i)_{i \in \mathbb{N}}$ *depend on the network* M*, but not on private signal precision.*

It is clear that we must have $r_i = \sum_{j=1}^{i} w_{i,j}$ where $w_{i,j}$ is the equilibrium weight that i puts on $\tilde{s}_j$. Since the matrix of weights W is independent of σ^2 by Proposition 2, we can use $(r_i)_{i \in \mathbb{N}_+}$ as a measure of how the network M affects the speed of rational information aggregation in our social-learning setting.

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 4. *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 4 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 $r_i \ge \overline{PL}(i) + 1$ 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 3.

The conditions that guarantee long-run learning are mild. The remainder of the paper studies the speed of learning by comparing $(r_i)_{i \in \mathbb{N}_+}$ on different networks.

Agent j belongs to the set of *indirect neighbors* of n, $\bar{N}(n) \subseteq \{1, ..., n-1\}$, if there exists a path of any length from n to j in the observation network. The next proposition studies the uniform lower-bound on r_n across all network structures, as a function of the size of n's

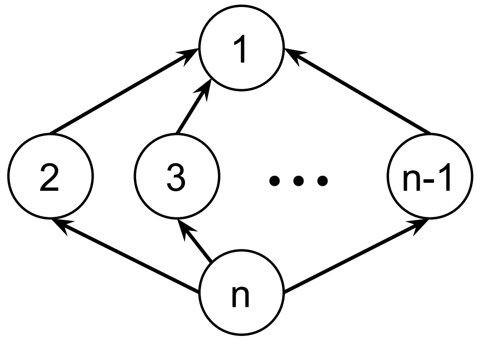

Figure 1: A network structure achieving the lower-bound on n's accuracy based on the size of her indirect neighborhood. For every n, $r_n < 5$.

indirect neighborhood. We show that an agent with arbitrarily many indirect neighbors must aggregate at least 5 signals in the limit. This bound is tight in that we exhibit a sequence of networks where $|\bar{N}(n)| = n - 1$ yet $r_n < 5$ for every $n \in \mathbb{N}$.

Proposition 5. *Let $\mathcal{M}_n$ be the set of all networks where $\bar{N}(n) = \{1, ..., n-1\}$, $r_n(M) = \mathbb{E}[\tilde{a}_n \mid \omega = 1]$ on network M, and $f(n) := \min_{M \in \mathcal{M}_n} r_n(M)$. Then $\lim_{n \to \infty} f(n) = 5$.*

On a network M where agent n has $n^{'} < n - 1$ indirect neighbors, we may apply Proposition 5 to the subnetwork given by the indirect neighborhood of n. This shows that on M, $r_n \geq f(n^{'})$. So in any network, agents with large enough indirect neighborhoods must aggregate at least close to 5 signals.

The proof of Proposition 5 makes use of the following example, which can be thought of the "worst" network for agent n in that it achieves the asymptotic lower-bound of signal aggregation as $n \to \infty$.

Example 1. In the network in Figure 1, agent n puts weight $\frac{2}{n-1}$ on each of $\tilde{a}_2, ..., \tilde{a}_{n-1}$ in equilibrium, thus social learning aggregates $r_n = 4 \cdot \frac{n-2}{n-1} + 1$ signals by agent n. We have $r_n < 5$ for all n.

4 Rate of Learning in the Generations Network

As an application of Section 3's general results, we study the speed of rational learning in the generations network. Agents are sequentially arranged into generations of size K. Agents in the first generation (i.e., $i = 1, ..., K$) have no neighbors. A collection of *inter-generational observation sets*, $\Psi_k \subseteq \{1, ..., K\}$ for $k = 1, ..., K$ define the network M for agents in later generations. For $i = (t-1)K + k$ where $t \geq 2$ and $1 \leq k \leq K$, network M

has $N(i) = \{(t-2)K + \psi : \psi \in \Psi_k\}$. So the k-th agent of generation $t \geq 2$ observes agents in Ψ_k from generation $t-1$ (and no agents from any earlier generations).[6]

4.1 Maximal 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.[7] The next result relates the generation size K to the speed of signal aggregation.

Proposition 6. *In the maximal generations network with any $K \geq 1$, society learns completely in the long run. We have $r_i = i \cdot \frac{(2K-1)}{K^2} + o(i)$, so social learning aggregates fewer signals per agent asymptotically with larger K, and no more than two signals per generation asymptotically for any K. Also, for any K and any agent i in generation $t \geq 3$, $r_i \leq K + 3t - 5$.*

Proposition 6 contains two parts. First, it shows that even though society eventually learns the true state with any K, the asymptotic rate of learning per agent is slower 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 comparison about the rate of learning hold asymptotically, but it also holds numerically for all agents $i \geq 16$ when comparing among $K \in \{2, 3, 4, 5\}$, as shown in Figure 2.

Second, Proposition 6 bounds the number of signals that social learning aggregates per generation in the maximal generations network. The proof of Proposition 4 shows $r_i \geq \overline{PL}(i) + 1$ in all networks and thus provides a lower bound of 1: each agent i in generation t has $\overline{PL}(i) = t - 1$. Proposition 6 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.

Suppose we are interested in a "confidence threshold" metric of learning efficiency — for instance, how many agents does it take to become at least 90% confident in the true state more than 90% of the time? The next corollary shows that societies with smaller generations reach any such threshold earlier when private signals are imprecise.

[6] 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.

[7] 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.

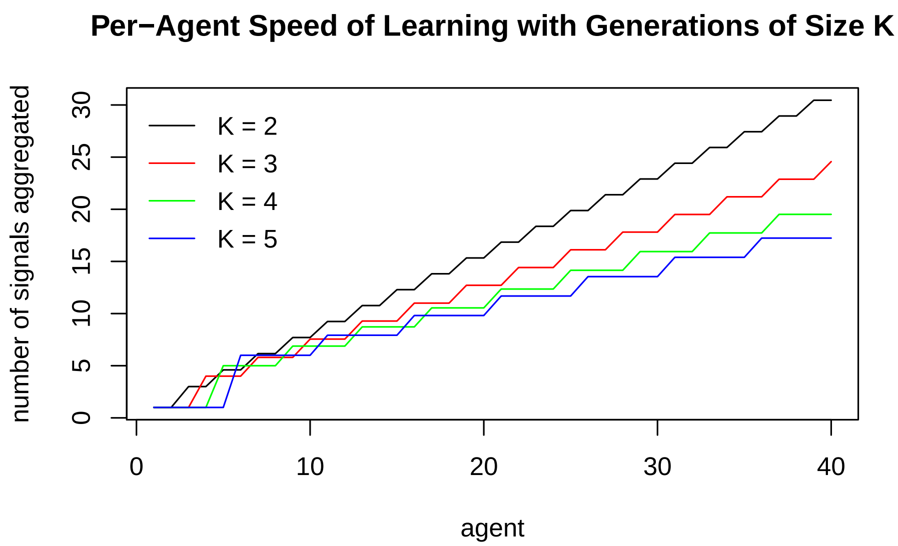

Figure 2: Number of signals aggregated by social learning in generations networks with different numbers of agents per generation, $K \in \{2, 3, 4, 5\}$.

Corollary 1. *For every* $0.5 < \bar{a} < 1$, $0 < \bar{p} < 1$, *and every pair of generation sizes* $K_1 < K_2$, *there exists a bound* $\tau > 0$ *on signal precision such that whenever* $0 < 1/\sigma^2 < \tau$, *the earliest* i *such that* $\mathbb{P}[a_i > \bar{a} \mid \omega = 1] > \bar{p}$ *is smaller with* K_1 *agents per generation than with* K_2 *agents per generation in the maximal generations network.*

Finally, we show that later agents put drastically different weights on the private signals from the first generation and the private signals from the immediate predecessor generation in forming their beliefs about ω. We emphasize that this asymmetric weighting arises in the equilibrium of a Bayesian model where agents engage in optimal signal extraction and make rational inferences about the state. Intuitively, this is because the informational confounds present in the maximal generations network lead to over-counting the earliest movers' signals even for rational players.

Corollary 2. *The total weight that generation* t *puts on the log-signals of generation* 1 *satisfies* $\sum_{j=1}^{K} w_{(t-1)K+1,j} \to \infty$ *as* $t \to \infty$, *while the total weight that generation* t *puts on the log-signals of generation* $t-1$ *satisfies* $\sum_{j=(t-2)K+1}^{(t-1)K} w_{(t-1)K+1,j} \to 1$ *as* $t \to \infty$.

4.2 Partial Observations and a General Learning-Rate Bound

Proposition 6 shows that social learning aggregates fewer than two signals per generation asymptotically on maximal generations networks with any K. We now provide an exact expression of the long-run rate of learning for a broad class of generations networks with more general inter-generational observation sets. In particular, this result will imply the same two signals per generation bound holds for all networks in this larger class.

We introduce two assumptions on $(\Psi_k)_k$.

Consider the directed graph on $\{1, ..., K\}$ induced by $(\Psi_k)_k$ where $k_1 \to k_2$ if and only if $k_2 \in \Psi_{k_1}$. This induced graph is *aperiodic at k* if there exists at least one cycle that starts at k, and no integer larger than 1 divides the length of every cycle that starts at k. The observation sets $(\Psi_k)_k$ are *weakly aperiodic* if every node is either aperiodic or has an in-degree of 0. This condition is weaker than the usual definition of aperiodicity and will be satisfied if $k \in \Psi_k$ for each k.

The observation sets are *strongly regular* if all agents observe $d \geq 1$ neighbors and all pairs of agents in the same generation share d_s common neighbors, i.e. $|N(i)| = d$ for all $i > K$ and $|N(i_1) \cap N(i_2)| = d_s$ 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. This adapts the usual notion of strongly regular networks to our directed networks with the generations structure.

To give a class of examples of weakly aperiodic and strongly regular networks, fix any non-empty subset $E \subseteq \{1, ..., K\}$, and let $(\Psi_k)_k$ be such that for all $1 \leq k \leq K$, $\Psi_k = E$. 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 represents the special case of $E = \{1, ..., K\}$. For another example, suppose $K \geq 2$ and each agent observes a different subset of $K-1$ predecessors from the previous generation. Specifically, $\Psi_k = \{1, ..., K\} \backslash \{k+1\}$ for $1 \leq k \leq K-1$, and $\Psi_K = \{1, ..., K\} \backslash \{1\}$. This network is weakly aperiodic and strongly regular with $d = K-1$ and $d_s = K-2$. Finally, there also exists a weakly aperiodic, strongly regular network with $K = 6$ agents, degrees $d = 3$, and pairwise intersections $d_s = 2$.

Theorem 1. *Suppose $(\Psi_k)_k$ are weakly aperiodic and strongly regular. Then*[8]

$$r_i = \left(1 + \frac{d^2 - d}{d^2 - d + d_s}\right) \frac{i}{K} + o(i).$$

Theorem 1 gives the exact asymptotic rate of learning for a more general class of generations networks. Provided $d_s \geq 1$, the number of signals aggregated per generation is strictly increasing in d and strictly decreasing in d_s, with the interpretation that more observations speed up the rate of learning per generation but more confounding slows it down, all else equal. Theorem 1 specializes to Proposition 6 by letting $d = d_s = K$.

[8] With the convention $0/0 = 0$.

In fact, the maximal generations network leads to the slowest per-generation rate of learning among all weakly aperiodic and strongly regular $(\Psi_k)_k$ where each agent observes K neighbors. Theorem 1 also implies that the speeds of learning are identical in a society where agents observe all predecessors from the past generation and another society where each agent observes a different subset of $K-1$ predecessors. The extra social observations exactly cancel out the reduced informational content of each observation, due to the more severe informational confounds in equilibrium.

Even though the maximal generations network leads to the slowest speed of signal aggregation conditional on the number of observations, Theorem 1 nevertheless provides a uniform learning-rate bound of two signals per generation across all weakly aperiodic and strongly regular generations networks, as $\frac{d^2-d}{d^2-d+d_s} \leq 1$. 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 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 the generations network for late enough t, each generation t agent's social observation constitutes a highly informative signal of the state (i.e., $n \to \infty$), but different agents can have different observations. This somewhat alleviates the informational confounding for generation $t+1$, but is limited by the fact that all agents' actions are approaching perfect correlation when $t \to \infty$. Even agents with very different neighborhoods end up observing highly correlated information in the long run, so no network in the class we consider aggregates more than two signals per generation asymptotically.

The uniformly slow speed of learning depends on each agent acting myopically and is not an inherent limitation of the generations structure. To illustrate this point, we show there exist feasible (but non-equilibrium) strategies so 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. 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 case where the k_2-th agent in every generation is always excluded from the indirect neighborhood of the k_1-th agent of every future generation, which would mean agents in position k_1 cannot aggregate more than $K-1$ signals per generation. The observation sets $(\Psi_k)_k$ are *aperiodic* if the induced graph on $\{1, ..., K\}$, which is given by $k_1 \to k_2$ if and only if $k_2 \in \Psi_{k_1}$, is aperiodic.

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 7. *Suppose the observation sets $(\Psi_k)_k$ are strongly connected and aperiodic. There is a 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 7 combine to say that in strongly connected and aperiodic generations networks, individuals only manage to aggregate an arbitrarily small fraction of the private signals that can be feasibly aggregated by a social planner.

5 When Does Adding Links Improve Accuracy?

For two observation networks M and $M^\bullet$, write $M^\bullet \geq M$ if $M^\bullet$ can be generated from M by adding links, that is $M^\bullet_{j,k} \geq M_{j,k}$ for all j, k. By Proposition 4, adding links leads to weakly better *asymptotic* learning outcomes — if the conditions for complete long-run learning are satisfied for M, then the same holds for $M^\bullet$. But, when does adding links improve *finite-time* accuracy?

We show that agent i is more accurate on network $M^\bullet$ than on network M if both networks are *transitive at i*: that is, whenever $j \in N(i)$ and $k \in N(j)$, we have $k \in N(i)$. We also highlight that *intransitivities* — that is, sequences of links $i_n \to i_{n-1}, i_{n-1} \to i_{n-2}, ..., i_2 \to i_1$ such that $i_n \not\to i_1$ — form a key obstacle to obtaining higher accuracy on denser networks. Accuracy may decrease for some agents if the new links create new intransitivities. Further, accuracy may decrease even in the absence of new intransitivities, if the baseline network M already contains intransitivities.

Proposition 8. *Suppose $M^\bullet \geq M$ and both networks are transitive at i. Then r_i is weakly higher on $M^\bullet$ than on M.*

The proof of Proposition 8 shows that for any network that is transitive at i, $W_{i,j} = 1$ for $j \in \bar{N}(i) \cup \{i\}$ and $W_{i,j} = 0$ otherwise — that is, i perfectly incorporates the private signals of all agents she indirectly observes. That is, r_i is equal to the number of agents indirectly observed by i. The denser network $M^\bullet$ improves i's accuracy because it expands i's indirect neighborhood.

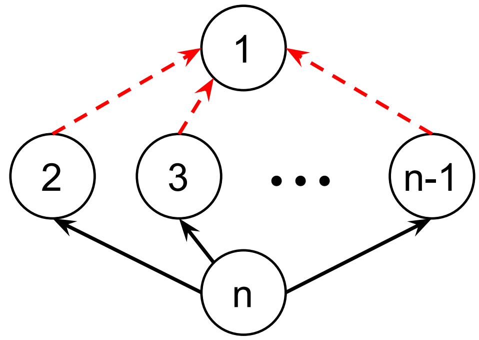

Figure 3: The black links define a transitive network M with n agents. For $k \in \{2, ..., n-1\}$, adding the $k-1$ red links from agents $2, ..., k$ to agent 1 creates a new network $M_k^\bullet$ that is no longer transitive for agent n. For $k \geq 5$, agent n has strictly lower accuracy on $M^\bullet$ than M.

We show by example that the same conclusion does not hold if $M^\bullet$ generates new intransitivities relative to M.

Example 2. Consider the networks $M, M_2^\bullet, ..., M_{n-1}^\bullet$ in Figure 3.

In network M, agent n perfectly incorporates the private signals of neighbors $2, ..., n-1$ and social learning aggregates $r_n = n-1$ signals. In network $M_k^\bullet$ for $2 \leq k \leq n-1$, the additional links expand n's indirect neighborhood relative to M, but also create informational confounds through intransitivities. In the new network, n cannot disentangle the private signals of her neighbors from the unobserved signal s_1 that serves as a common influence for her neighbors' behavior.

In the equilibrium on network $M_k^\bullet$, n puts weight $\frac{2}{k}$ on the log-actions of $2, ..., k$, thus social learning aggregates $r_n^{(k)} = 4 \cdot \frac{k-1}{k} + n - k$ signals by agent n. We have that $r_n^{(2)} > r_n$, so the first red link helps and allows n to incorporate all private signals. But $r_n^{(k)}$ is strictly decreasing in k. In particular, $r_n^{(k)}$ is strictly smaller than r_n whenever $k \geq 5$. Adding four or more red links to the original network M strictly harms n's welfare.

Suppose the baseline network M already contains some intransitivities. The next examples shows that adding links may decrease some agent's accuracy even if the new links do not create new intransitivities. In particular, links can harm agents without creating any new confounds simply by changing the weights on existing confounds.

Example 3. Consider the networks M and $M^\bullet$ in Figure 4.

Intransitivities exist both in the old network M defined by the black links, and in the new network $M^\bullet$ that adds the one red link. Even though the newly added link does not generate any additional intransitivities, we have $r_9 = \frac{3681}{533} \approx 6.91$ in the old network and

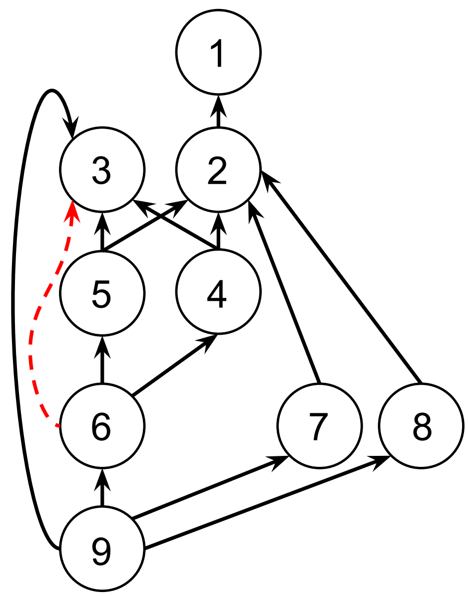

Figure 4: Adding the new link in red does not create new intransitivities, but nevertheless decreases agent 9's accuracy.

$r_9 = \frac{1977}{287} \approx 6.89$ in the new network, so socially learning aggregates fewer signals by agent 9 in $M^{\bullet}$.

Agent 9 becomes less accurate in $M^{\bullet}$ because agent 6's new link causes her to change her equilibrium play in a way that generates negative informational externality for agent 9. In both M and $M^{\bullet}$, agent 6 cannot fully incorporate the private signals of agents 4 and 5 without over-counting the private signals of agents 1 and 2. In network M, agent 6 puts weight $\frac{4}{7}$ on the log-actions of agents 4 and 5, thus weight $\frac{8}{7}$ on $\tilde{s}_1$ and $\tilde{s}_2$. In network $M^{\bullet}$, agent 6 puts a higher weight $\frac{3}{5}$ on the log-actions of agents 4 and 5, because she can now subtract off part of the informational confound using her observation of agent 3. This change in her equilibrium strategy means her over-weighting of $\tilde{s}_1$ and $\tilde{s}_2$ is exacerbated, with these log-signals each receiving weight $\frac{6}{5}$. At the same time, $\tilde{s}_1$ and $\tilde{s}_2$ also confound agent 9's inference about the private signals of agents 7 and 8. Agent 9 finds it harder to incorporate agent 6's private signal in $M^{\bullet}$, because $\tilde{a}_6$ now contains a more severe over-counting of $\tilde{s}_1$ and $\tilde{s}_2$. The change in agent 6's play on $M^{\bullet}$ does not taken into account the welfare of agent 9, who has a different signal-extraction problem that involves worse confounding by $\tilde{s}_1$ and $\tilde{s}_2$.

References

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

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.

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.

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.

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 (2019): "Rational Groupthink," *Working Paper*.

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.

Rosenberg, D. and N. Vieille (2019): "On the efficiency of social learning," *Econometrica*, 87, 2141–2168.

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.

Appendix

A 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 $\vec{\mu} = \mathbb{E}[(\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_i)})' \mid \omega = 1]$ and $\Sigma = \text{Cov}[\tilde{a}_{j(1)}, ..., \tilde{a}_{j(n_i)} \mid \omega = 1]$.

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 $\vec{\beta}_{I,\cdot}$. □

A.3 Proof of Proposition 2

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

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

For the last claim, $W_1 = (1, 0, 0, ...)$ 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.4 Proof of Proposition 3

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

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

$$\begin{aligned}\textsc{Var}\left[\sum_{k=1}^{n_i}\beta_{i,j(k)}\tilde{a}_{j(k)} \mid \omega = 1\right] &= \left(2\vec{\mu}'\Sigma^{-1}\right) \cdot \Sigma \cdot \left(2\vec{\mu}'\Sigma^{-1}\right)' \\ &= 4\vec{\mu}'\Sigma^{-1}\vec{\mu}\end{aligned}$$

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

A.5 Proof of Proposition 4

We first state and prove an auxiliary lemma.

Lemma 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 3, 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 4.

Proof. By Proposition 3, there exist $(r_i)_{i\in\mathbb{N}_+}$ 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 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 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 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 Example 1

Proof. We apply Proposition 2 to calculate $\vec{\beta}_{n,\cdot}$ in this network. The submatrix $\hat{W}$ of W with rows $\{2, ..., n-1\}$ and columns $\{1, ..., n-1\}$ is

$$\hat{W} = \begin{pmatrix} 1 & 1 & 0 & 0 & \dots \\ 1 & 0 & 1 & 0 & \dots \\ 1 & 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

So we get

$$(\hat{W}\hat{W}')^{-1} = \begin{pmatrix} 2 & 1 & 1 & \dots \\ 1 & 2 & 1 & \dots \\ 1 & 1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}^{-1} = I_{n-2} - \frac{1}{n-1}\text{Ones}_{n-2}$$

where I_{n-2} is the $(n-2)\times(n-2)$ identity matrix and Ones_{n-2} is the $(n-2)\times(n-2)$ matrix of all 1's. So,

$$\begin{aligned} \hat{W}'(\hat{W}\hat{W}')^{-1} &= \hat{W}' - \frac{1}{n-1}\hat{W}' \cdot \text{Ones}_{n-2} \\ &= \begin{pmatrix} 1 & 1 & 1 & \dots \\ 1 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} - \frac{1}{n-1}\begin{pmatrix} n-2 & n-2 & n-2 & \dots \\ 1 & 1 & 1 & \dots \\ 1 & 1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \end{aligned}$$

where the dimension of each matrix is $(n-1)\times(n-2)$. The sum of each column is $2-\frac{2n-4}{n-1} = \frac{2}{n-1}$, which is $\vec{\beta}_{n,j}$ for $j = 2, ..., n-1$. From here it is easy to calculate that $r_n = 4\cdot\frac{n-2}{n-1}+1$. □

A.7 Proof of Proposition 5

Proof. By the proof of Example 1, $f(n) < 5$ for all n. It suffices to show that $\liminf_n f(n) \geq 5$. We will use the following lemma.

Lemma 3. *In any network, suppose* $|\bar{N}(i)| \geq 4$. *Then* $r_i \geq 4$.

Proof. Suppose $r_i < 4$. There must not be any $j \in N(i)$ so that $r_j \geq 3$. For any agent j, we have $r_j \geq 3$ if $|\bar{N}(j)| \geq 2$. Therefore, for $j \in N(i)$, we must have $|\bar{N}(j)| \in \{0, 1\}$.

First, suppose $|\bar{N}(j)| = 0$ for all $j \in N(i)$. Since i then observes at least four agents with empty neighborhoods, $r_i \geq 5$.

Next, suppose there exists exactly one $j^* \in N(i)$ with $|\bar{N}(j^*)| = 1$. Then there exist two other $j_1 \neq j_2$ in $N(i)$ each with no observations, and neither can be observed by j^*. In this case again $r_i \geq 5$.

Finally, suppose that $|\bar{N}(j)| = 1$ for at least two agents observed by i. Consider any two such agents, $j_1 \neq j_2$. Neither of these agents can observe the other, because then one would have an indirect neighborhood of size at least 2. If these j_1 and j_2 observe distinct neighbors, then $r_i \geq 5$. Therefore, all $j \in N(i)$ with $|\bar{N}(j)| = 1$ must observe the same neighbor. If there are at least three such j, the computation in the upper bound in the proof of Example

1 shows that $r_i \geq 4$. Else, to account for the 4 agents in i's indirect neighborhood, there must be distinct $j_1, j_2 \in N(i)$ who observe the same neighbor, and an additional $j_3 \in N(i)$ who is not observed by either j_1 or j_2. Then $r_i \geq 5$ as well by the proof of Example 1, which completes the lemma. □

Now we return to the proof of Proposition 5. Suppose that for some $\epsilon > 0$, there is an increasing sequence of positive integers $(b_n)_{n\in\mathbb{N}}$ with $f(b_n) \leq 5-\epsilon$ for all n. For each $n \in \mathbb{N}$, choose a network M_n with $|\bar{N}(b_n)| = b_n - 1$ and $r_{b_n}(M_n) \leq 5-\epsilon$. If b_n observes some agent i with $|\bar{N}(i)| \geq 4$, then by Lemma 3, we would have $r_i \geq 4$ and therefore $r_{b_n} \geq 5$. Thus, for any agent $i \in N(b_n)$, we have $|\bar{N}(i)| \leq 3$. An immediate consequence is that the number of agents observed by agent b_n in M_n must converge to ∞ as $n \to \infty$.

Because there are finitely many networks with four agents, there are finitely many possible values of r_i for agents $i \in N(b_n)$. By the pigeonhole principle, there exists some r^* and an infinite subsequence of the networks $(M_n)_{n\in\mathbb{N}}$, so that the number of agents $i \in N(b_n)$ with $r_i = r^*$ in M_n converges to ∞ with n along the subsequence.

If $r^* = 1$, we have a contradiction since this shows $r_{b_n}(M_n) > 5$ for large n, since b_n observes an unbounded number of agents who have no social observations.

Now assume $r^* \geq 2$, which is the next smallest possible value. Identify $(b_n)_{n\in\mathbb{N}}, (M_n)_{n\in\mathbb{N}}$ with the subsequence and $d(n)$ is the number of agents $i \in N(b_n)$ with $r_i = r^*$ in M_n. We show that in each network M_n, it is feasible for b_n to construct an estimate from the log-actions of those agents i with $r_i = r^*$ and her own log-signal, which is at least as informative as five independent signals about the state. This leads to a contradiction as it shows $r_n(M_n) \geq 5$ for large n.

First consider the case where any two $i \neq j \in N(b_n)$ with $r_i = r_j = r^*$ have $\text{Cov}[\tilde{a}_i, \tilde{a}_j \mid \omega = 1] = (r^*-1)\frac{4}{\sigma^2}$. Suppose b_n places weight w on each of these agents and weight one on her private signal. Then the conditional mean of her action is $(w \cdot r^* \cdot d(n) + 1) \cdot \frac{2}{\sigma^2}$ and the conditional variance is

$$\frac{4}{\sigma^2}\left(w^2 d(n)^2 r^* - w^2 d(n)(d(n)-1) + 1\right).$$

These differ by a factor of 2 when

$$w = \frac{r^*}{d(n)r^* - d(n) + 1}.$$

Thus this estimate has the same joint distribution with the state as $\frac{(r^*)^2 d(n)}{d(n)r^*-d(n)+1} + 1$ independent private signals, showing that $r_{b_n} \geq \frac{(r^*)^2 d(n)}{d(n)r^*-d(n)+1} + 1$. Since $d(n) \to \infty$, this shows $\liminf_n r_{b_n}(M_n) \geq \frac{(r^*)^2}{r^*-1} + 1 \geq 5$ since $r^* \geq 2$. This contradicts the hypothesis that

$r_{b_n}(M_n) \leq 5 - \epsilon$ for all n.

Now we drop the assumption that $\mathrm{Cov}[\tilde{a}_i, \tilde{a}_j \mid \omega = 1] = (r^* - 1)\frac{4}{\sigma^2}$ for all $i \neq j \in N(b_n)$ with $r_i = r_j = r^*$. The conditional variance of the sum placing weight 1 on own signal and any weight $w > 0$ on the $d(n)$ neighbors is

$$\frac{4}{\sigma^2} + \sum_{i=1}^{d(n)} w^2 r^* \frac{4}{\sigma^2} + 2 \sum_{1 \leq i < j \leq d(n)} w^2 \mathrm{Cov}(\tilde{a}_i, \tilde{a}_j).$$

It is strictly increasing in each $\mathrm{Cov}(\tilde{a}_i, \tilde{a}_j)$ term. So if all i, j pairs have conditional covariance of larger than $(r^* - 1) \cdot \frac{4}{\sigma^2}$, the estimate constructed above will have the same conditional mean and smaller conditional variance than $\frac{(r^*)^2 d(n)}{d(n)r^* - d(n) + 1} + 1$ independent private signals, which again implies $r_{b_n} \geq \frac{(r^*)^2 d(n)}{d(n)r^* - d(n) + 1} + 1$.

Finally, we show that for two agents $i \neq j$ with $r_i = r_j = r^*$, $\mathrm{Cov}(\tilde{a}_i, \tilde{a}_j) \leq (r^* - 1) \cdot \frac{4}{\sigma^2}$. We have $\sum_k w_{i,k} = \sum_k w_{j,k} = r^*$. Also, i cannot be in j's indirect neighborhood and vice versa, else we cannot get $r_i = r_j$ as every agent is strictly more accurate than any agent in their indirect neighborhood. Thus we have $w_{i,i} = 1, w_{i,j} = 0, w_{j,i} = 0, w_{j,j} = 1$. So

$$\begin{aligned} \mathrm{Cov}(\tilde{a}_i, \tilde{a}_j) &= \mathrm{Cov}(\tilde{s}_i + \sum_{k \neq i,j} w_{i,k} \tilde{s}_k, \tilde{s}_j + \sum_{k \neq i,j} w_{j,k} \tilde{s}_k) \\ &= \frac{4}{\sigma^2} \sum_{k \neq i,j} w_{i,k} w_{j,k}. \end{aligned}$$

By the Cauchy–Schwarz inequality, $(\sum_{k \neq i,j} w_{i,k} w_{j,k})^2 \leq (\sum_{k \neq i,j} w_{i,k}^2) \cdot (\sum_{k \neq i,j} w_{j,k}^2)$. By Proposition 3, $\sum_{k \neq i,j} w_{i,k}^2 = \sum_{k \neq i,j} w_{j,k}^2 = r^* - 1$, thus $\sum_{k \neq i,j} w_{i,k} w_{j,k} \leq r^* - 1$. $\square$

A.8 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 $d_s = 0$, then we must have $d = 1$. From now on we assume $d \geq 2$ and $d_s \geq 1$.

We claim that for each generation t and each i, i' in generation t, $\mathrm{Var}[\tilde{a}_i \mid \omega = 1]$ and $\mathrm{Cov}[\tilde{a}_i, \tilde{a}_{i'} \mid \omega = 1]$ depend only on t and not on the identities of i or i', which we call Var_t and Cov_t, respectively. Similarly, for i in generation t and each $j \in N(i)$, the weight $\beta_{i,j}$ depends only on t, which we call β_t.

The claims hold by inductively applying the strong regularity 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 $\mathrm{Var}[\tilde{a}_j \mid \omega = 1]$ 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 3, 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

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

for all i in generation t, and

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

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 $\textsc{Var}_t$ and $\textsc{Cov}_t$ gives:

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

We now require two auxiliary lemmas.

Lemma 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 weakly aperiodic and strongly regular with $d_s \geq 1$. 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. Each closed communication class is irreducible and aperiodic when $(\Psi_k)_k$ is weakly aperiodic and $d \geq 1$, so 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 $i \neq j$, $\| p_i^\infty - p_j^\infty \|_{\max} \leq \left(\frac{d-d_s}{d}\right)^t$ for all $t \geq 1$. Since $d_s \geq 1$, this would show that in fact $p_i^\infty = p_j^\infty$ for all i, j.

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

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

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

so

$$\begin{aligned}\| p_i^\infty - p_j^\infty \|_{\max} &\leq \frac{1}{d}\sum_{k=d_s+1}^{d} \| p_{n_k}^\infty - p_{n_k'}^\infty \|_{\max} \\ &\leq \frac{d-d_s}{d}\cdot 1\end{aligned}$$

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 $\| p_{n_k}^\infty - p_{n_k'}^\infty \|_{\max} \leq \left(\frac{d-d_s}{d}\right)^{t-1}$ from the inductive hypothesis. □

Lemma 5. *$\beta_t \to 1/d$.*

Proof. By Proposition 3, we can compute that:

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

It is therefore sufficient to show that $\mathrm{VAR}_t/\mathrm{COV}_t \to 1$. 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 $\{1,\dots,K\}$ 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 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:

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

We want to show that

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

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. □

We return to the proof of Theorem 1. Fix small $\epsilon > 0$. By Lemma 5, we can choose T such that $\beta_t \leq \frac{1+\epsilon}{d}$ for all $t \geq T$. Therefore, $\beta_t^2 (d - d_s) \leq \frac{(1+\epsilon)^2}{d^2}(d - d_s)$ for $t \geq T$. Consider the contraction map $\varphi(x) = \frac{4}{\sigma^2} + \frac{(1+\epsilon)^2}{d^2}(d - d_s)x$. 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

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

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

At the same time, $\beta_t \geq \frac{1}{d}$ for all t. Consider the contraction map $\varphi(x) = \frac{4}{\sigma^2} + \frac{1}{d^2}(d - d_s)x$. 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

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

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

Using Proposition 3, we have $\textsc{Var}_{t+1} = 2(\beta_{t+1} d(\textsc{Var}_t/2) + 2/\sigma^2)$, so

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

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

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

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

A.9 Proof of Proposition 6

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 4, society learns completely in the long run. The expression for r_i comes from specializing Theorem 1 (whose proof does not depend on Proposition 6 or Corollaries 1 and 2) to the case of $d = d_s = K$. 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 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

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

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

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. □

Consider an agent i in generation t. From Proposition 3, 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 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 $x_{new} - x_{old} \leq 3 \cdot \frac{2}{\sigma^2}$. All log-actions from generation 1 are independent, so the sum of first generation actions has conditional mean and variance of $\mu_{\text{sum}} = 2K/\sigma^2$ and $\sigma_{\text{sum}}^2 = 4K/\sigma^2$. This shows agents in generation 3 have log-actions conditionally distributed as $\mathcal{N}(m_3, 2m_3)$ where $m_3 \leq (K+4) \cdot \frac{2}{\sigma^2}$. And so in general, the equilibrium log-actions of each agent in generation $t \geq 3$ has the conditional distributions $\mathcal{N}(\pm m_t, 2m_t)$ where $m_t \leq (K + 3t - 5)\frac{2}{\sigma^2}$. □

A.10 Proof of Corollary 1

Let $\bar{a} > 0.5$, $\bar{p} > 0$ be given.

Let $r_i^{(K_1)}, r_i^{(K_2)}$ denote the number of signals aggregated by social learning by agent i in the maximal generations networks with K_1 or K_2 agents per generation, respectively. Both $(r_i^{(K_1)})$ and $(r_i^{(K_2)})$ are increasing, unbounded sequences and Proposition 6 implies there exists some I so that $r_i^{(K_1)} > r_i^{(K_2)}$ for all $i \geq I$. Note that by Proposition 6, this I does not depend on the private signal variance σ^2.

Choose τ near enough to 0 so that $1 - \Phi\left(\frac{\tilde{\bar{a}} - r_I^{K_1}\frac{2}{\sigma^2}}{\sqrt{r_I^{K_1}\frac{2}{\sigma}}}\right) < \bar{p}$ whenever $0 < 1/\sigma^2 < \tau$. By Lemma 2, this ensures that for all $i \leq I$, $\mathbb{P}[a_i > \bar{a} \mid \omega = 1] < \bar{p}$ in both networks provided $1/\sigma^2 < \tau$. For $i > I$, if $\mathbb{P}[a_i > \bar{a} \mid \omega = 1] > \bar{p}$ with K_2 agents per generation, then by $r_i^{(K_1)} > r_i^{(K_2)}$, we also have the same condition satisfied with K_1 agents per generation, again by Lemma 2. Thus the earliest i where this condition holds occurs after $i = I$ and occurs earlier with K_1 agents per generation.

A.11 Proof of Corollary 2

Proof. That $\sum_{j=(t-2)K+1}^{(t-1)K} w_{(t-1)K+1,j} \to 1$ as $t \to \infty$ just comes from observing that this sum of weights is equal to the sum of entries in $\vec{\beta}_{(t-1)K+1,\cdot}$, which is $K \cdot \left(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}\right) / \left(K\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}\right)$ by Lemma 6, where μ_{sum} and σ_{sum}^2 are the conditional mean and variances of the sum of the log-actions played in generation $t-1$ in equilibrium. Since $\lim_{i\to\infty} r_i = \infty$, $\frac{\mu_{\text{sum}}}{\sigma_{\text{sum}}}$ must grow without bound in generations, so $K \cdot \left(\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}\right) / \left(K\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2} + \frac{1}{\sigma^2}\right) \to 1$.

Let $\bar{w}_{T,t} = w_{(T-1)K+1,(t-1)K+1}$, the weight that each agent in generation T puts on the log-signal of each agent in generation $t \leq T$. Let $\bar{\beta}_t$ be the weight that each agent in generation t puts on each log-action in generation $t-1$.

We have

$$\bar{w}_{T,t} = K^{T-t-1} \cdot \prod_{j=1}^{T-t} \beta_{T-j+1}$$

for $1 \leq t \leq T-1$. This is because we may imagine a weighted network where each agent in generation t is connected to each agent in generation $t-1$ with a link of weight $\bar{\beta}_t$, and $\bar{w}_{T,t}$ is the number of weighted paths between an agent in generation T and an agent in generation t.

Let $G_t = r_{(t-1)K+1}$, the number of signals that social learning aggregates by generation t. Consider all previous generations $t < T$ and all K predecessors in generation t, we have

$G_T = 1 + \sum_{t=1}^{T-1} K \cdot \bar{w}_{T,t}$. Therefore

$$G_{T+1} - G_T = K \cdot \bar{w}_{T+1,T} + \sum_{t=1}^{T-1} (K\bar{w}_{T+1,t} - K\bar{w}_{T,t})$$
$$= K\beta_{T+1} + (K\beta_{T+1} - 1) \sum_{t=1}^{T-1} (K\bar{w}_{T,t})$$

Since each β_t is of the form $\frac{\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2}+\frac{1}{\sigma^2}}{K\frac{\mu_{\text{sum}}^2}{\sigma_{\text{sum}}^2}+\frac{1}{\sigma^2}}$, we have $K\beta_t > 1$ for all t, hence $\bar{w}_{T,t}$ is positive and decreasing in t. So $G_{T+1}-G_T \leq K\beta_{T+1}+(T-1)\cdot(K\bar{w}_{T+1,1} - K\bar{w}_{T,1})$. For large T, Proposition 6 implies $G_{T+1} - G_T \approx \frac{2K-1}{K}$ and $K\beta_{T+1} \approx 1$, so $(K\bar{w}_{T+1,1} - K\bar{w}_{T,1}) \geq \frac{1}{T-1}(\frac{2K-1}{K} - 1)$ for large T. This shows $K\bar{w}_{T,1} \to \infty$ as $T \to \infty$ by comparison to the harmonic series. □

A.12 Proof of Proposition 7

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 and 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.13 Proof of Proposition 8

Proof. We first show that on any network transitive at i, the equilibrium strategy of i is such that $\tilde{a}_i = \sum_{j \in \bar{N}(i) \cup \{i\}} \tilde{s}_j$. Clearly, $\tilde{a}_i$ cannot put any weight on the log-signals of agents not in $\bar{N}(i) \cup \{i\}$, for information outside of the sub-network containing i cannot reach i. Also, if feasible, $\sum_{j \in \bar{N}(i) \cup \{i\}} \tilde{s}_j$ is the optimal signal aggregation for i. For every $j \in \bar{N}(i)$, we have $N(j) \subseteq N(i)$. Since i knows j's linear equilibrium strategy $\tilde{A}^*_j((\tilde{a}_k)_{k \in N(j)}, \tilde{s}_j)$, i can identify $\tilde{s}_j$ by calculating $\tilde{a}_j - \sum_{k \in N(j)} \beta_{j,k} \tilde{a}_k$. Therefore i can identify the sum $\sum_{j \in \bar{N}(i)} \tilde{s}_j$ using her neighbors' actions.

Combined with Proposition 3, this shows r_i on any network transitive at i is equal to the cardinality of $\bar{N}(i)$ plus one. Agent i must have a larger indirect neighborhood on $M^\bullet$ if $M^\bullet \geq M$. □

A.14 Proof of Example 2

Proof. In the equilibrium on network $M^\bullet_k$, agent n puts weight 1 on each of $\tilde{a}_{k+1}, ..., \tilde{a}_{n-1}$ and the same weights on $\tilde{a}_2, ..., \tilde{a}_k$ as in Example 1 when $n = k+1$. From the proof of Example 1, this weight is $\frac{2}{k}$.

We have $r^{(2)}_n = n > r_n$, while $\frac{d}{dk}(4\frac{k-1}{k} + n - k) = \frac{4}{k^2} - 1 < 0$ for $k > 2$. This shows $r^{(k)}_n$ is strictly decreasing in k for $k \geq 2$. Finally, $r_n - r^{(k)}_n = k^2 - 5k + 4$ is a convex quadratic function with zeroes at 1 and 4. So $r_n = r^{(4)}_n$ while $r_n > r^{(k)}_n$ for all $k \geq 5$. □

A.15 Proof of Example 3

Proof. For the network where agent 6 does not observe agent 3, by Proposition 2 we find

$$\vec{\beta}_{6,\cdot} = \vec{\mathbf{1}}'_{(5)} \cdot \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot \left(\begin{pmatrix} 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \right)^{-1} = \begin{pmatrix} \frac{4}{7} \\ \frac{4}{7} \end{pmatrix}.$$

Then, letting

$$\hat{W} := \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \frac{8}{7} & \frac{8}{7} & \frac{8}{7} & \frac{4}{7} & \frac{4}{7} & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix},$$

we get

$$\vec{\beta}_{9,\cdot} = \vec{\mathbf{1}}'_{(8)} \cdot \hat{W}'(\hat{W}\hat{W}')^{-1} = \vec{\mathbf{1}}'_{(8)} \cdot \hat{W}' \cdot \frac{1}{533} \begin{pmatrix} 853 & -280 & 128 & 128 \\ -280 & 245 & -112 & -112 \\ 128 & -112 & 371 & -162 \\ 128 & -112 & -162 & 371 \end{pmatrix} = \frac{1}{533} \begin{pmatrix} 61 \\ 413 \\ 131 \\ 131 \end{pmatrix},$$

and hence we can calculate the 9th row of W and r_9.

For the network where agent 6 observes agent 3, note that agent 6 can recover $\tilde{s}_1 + \tilde{s}_2 + \tilde{s}_4$ and $\tilde{s}_1 + \tilde{s}_2 + \tilde{s}_5$ using $\tilde{a}_4 - \tilde{a}_3$ and $\tilde{a}_5 - \tilde{a}_3$. Thus, the weights that agent 6 puts on $\tilde{a}_4$ and $\tilde{a}_5$ are the same as in a network where agents 4 and 5 only observe agent 2. This can be computed by Proposition 2:

$$\vec{\mathbf{1}}'_{(5)} \cdot \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot \left(\begin{pmatrix} 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \right)^{-1} = \begin{pmatrix} \frac{3}{5} \\ \frac{3}{5} \end{pmatrix},$$

which shows $\vec{\beta}_{6,\cdot} = (-\frac{1}{5}, \frac{3}{5}, \frac{3}{5})'$. Then, letting

$$\hat{W} := \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \frac{6}{5} & \frac{6}{5} & 1 & \frac{3}{5} & \frac{3}{5} & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix},$$

we get

$$\vec{\beta}_{9,\cdot} = \vec{\mathbf{1}}'_{(8)} \cdot \hat{W}'(\hat{W}\hat{W}')^{-1} = \vec{\mathbf{1}}'_{(8)} \cdot \hat{W}' \cdot \frac{1}{287} \begin{pmatrix} 412 & -125 & 60 & 60 \\ -125 & 125 & -60 & -60 \\ 60 & -60 & 201 & -86 \\ 60 & -60 & -86 & 201 \end{pmatrix} = \frac{1}{287} \begin{pmatrix} 72 \\ 215 \\ 69 \\ 69 \end{pmatrix},$$

and hence we can calculate the 9th row of W and r_9. □