Social Learning in a Dynamic Environment*

Krishna Dasaratha Benjamin Golub Nir Hak†

May 27, 2018

Abstract

Agents learn about a state using private signals and the past actions of their neighbors. In contrast to most models of social learning in a network, the target being learned about is moving around. We ask: when can a group aggregate information quickly, keeping up with the changing environment? First, if private signal distributions are diverse enough across agents, then Bayesian learning achieves good information aggregation as long as individuals observe sufficiently many others. Second, without such diversity, Bayesian information aggregation can fall far short of good aggregation benchmarks, and can be Pareto-inefficient. Third, good aggregation requires anti-imitation; without it, agents' estimates are inefficiently confounded by "echoes." Our stationary equilibrium learning rules incorporate past information by taking linear combinations of other agents' past estimates (as in the simple DeGroot heuristic), and we characterize the coefficients in these linear combinations. The resulting tractability can facilitate structural estimation of equilibrium learning models and testing against behavioral alternatives, as well as the analysis of welfare and influence.

*We are grateful to Nageeb Ali, Drew Fudenberg, Kevin He, Matthew O. Jackson, Annie Liang, Eric Maskin, Margaret Meyer, Michael Ostrovsky, Matthew Rabin, Evan Sadler, Tomasz Strzalecki, Andrzej Skrzypacz, Omer Tamuz, Bob Wilson, and Jeff Zwiebel for valuable discussions.

†Department of Economics, Harvard University, Cambridge, U.S.A., dasarath@g.harvard.edu, bgolub@fas.harvard.edu, nirhak@g.harvard.edu.

1 Introduction

Consider a group learning, over a period of time, about an evolving fundamental state, such as future conditions in a market. For instance, analysts form opinions about the future demand for some product. In addition to making use of public information, individuals learn from their own private information and also from the estimates of others. Analysts have their own data and calculations, but may also access the reports of some other analysts, or simply talk to others.[1] A crucial feature of these environments is that without a central information aggregation device, aggregation of information occurs locally and estimates may differ across a population. This type of problem studied in various settings, including in the development literature focused on situations where the flow of information crucial to production decisions tends to be encumbered by geographic and social distance(see, e.g., Jensen, 2007; Srinivasan and Burrell, 2013). There, individuals may be trying to learn the yield of a new crop, or the wages of working at a certain job.

Given that the fundamental state in question is changing over time, a key question is: When can the group respond to the environment quickly, aggregating dispersed information efficiently in real time? In contrast, when are estimates of present conditions confounded? From a design perspective, when can a network that is not performing well be improved by changing aspects of it, e.g., facilitating interaction between some people, and when is a different approach needed—such as a centralized poll, exchange, or prediction market?

The question of whether decentralized communication can facilitate efficient adaptation to a changing world is a fundamental one in economic theory, related to questions raised by Hayek (1945).[2] Nevertheless, there is relatively little modeling addressing it: the large literatures on social learning and information aggregation in networks typically do not consider dynamic states, though there are notable exceptions,as we discuss in Section 6. Our first contribution is to define and study equilibria in a dynamic environment that captures

[1] Alternatively, as in some recent models of over-the-counter markets, agents learn about each other's valuations when they negotiate or trade with each other (Duffie and Manso, 2007; Duffie, Malamud, and Manso, 2009).

[2] "If we can agree that the economic problem of society is mainly one of rapid adaptation to changes in the particular circumstances of time and place...we must solve it by some form of decentralization. ... There still remains the problem of communicating to [each individual] such further information as he needs to fit his decisions into the whole pattern of changes of the larger economic system. How much knowledge does he need to do so successfully? Which of the events which happen beyond the horizon of his immediate knowledge are of relevance to his immediate decision, and how much of them need he know?" Hayek was especially concerned with the function of market prices in relation to these questions, but very similar issues are relevant when we consider information aggregation more generally.

two essential dimensions emphasized above: there is a changing state of the world, and communication is decentralized, occurring in an arbitrary network. Our second contribution is to derive conditions under which decentralized information aggregation works well. We find that the dynamic state changes the answers considerably relative to models where the state is static: In contrast to previous work, we show that whether signal endowments are the same (in distribution) across agents matters critically for information aggregation. Our main substantive finding is that in large populations decentralized learning *can* approach optimal benchmarks, as long as (i) agents have diverse signal endowments and (ii) their learning rules are Bayesian, responding to correlations in a sophisticated way. When signal endowments are not diverse, then social learning can be inefficiently confounded far from optimal, even when each agent has access to an unbounded amount of information. When, instead of using Bayesian learning rules, agents use reasonable heuristics that are known to work well in static-state benchmarks, learning outcomes are also far from efficient benchmarks.

We now describe the model in more detail. The state, θ_t, drifts around according to a stationary, discrete-time AR(1) process given by $\theta_{t+1} = \rho\theta_t + \nu_{t+1}$, and agents receive conditionally independent Gaussian signals of its current value. The population consists of overlapping generations of decision-makers (agents), located in a network. In each period, each agent observes an independent, normal signal $s_{i,t} \sim \mathcal{N}(\theta_t, \sigma_i^2)$ of the current state andsome past estimates of her neighbors in an arbitrary network, which (because they depend on recent states) are relevant for estimating θ_t .[3] The agent's action is an estimate $a_{i,t}$: she sets it to her expectation, given all her information, of the current state θ_t. Her estimate is then used by her neighbors in the next round in the same way. We vary three features of the environment:

1. The distributions of individuals' information: in particular, the precisions σ_i^{-2} of private signals and how these precisions vary across the population.

2. The structure of the network: the sizes and compositions of individuals' neighborhoods.

3. How agents update: the baseline model is that agents take correct Bayesian conditional expectations of the state of interest; an important alternative is that agents do not optimally account for redundancies in their observations.

[3] The network can be directed or undirected; that is, observation opportunities need not be reciprocal.

A helpful feature of the model is that, when agents are Bayesian, stationary equilibrium learning rules take a simple, time-invariant form: agents form their next-period estimates by taking linear combinations of each other's earlier estimates and their own private signals. The form of these equilibrium learning rules and the DeGroot (1974) updating rule are related, in that DeGroot agents also incorporate information from the past by taking weighted averages of their neighbors' actions. Our model is one in which such updating arises from Bayesian behavior. We also compare Bayesian updating to behavioral rules in which players use weights that are not optimal, but arise from heuristics such as weighting others in proportion to the quality of their signals.

To derive our basic results about equilibrium, we study the problem of all agents simultaneously seeking to extract data about the underlying state from neighbors' estimates. Technically, the problem facing each agent is a very standard one, of estimation in a linear statistical model. However, the linear combinations everyone observes of the underlying fundamentals depend on others' strategies, i.e., updating rules. Agents' (linear) strategies in a given period determine a variance-covariance structure on the distribution of their estimates, and thus determine each agent's weights on neighbors and signals about the fundamental in the next period. That, in turn, determines the variance-covariance structure of next period's estimates, and each stationary equilibrium corresponds to a variance-covariance structure that is invariant under this process. Thus, the question of whether states of interest are identified is a subtle one, and depends on the whole equilibrium, as well as the network. Equilibrium weights satisfy a system of polynomial equations, but this system is usually complex and has high degree.

The key findings on the efficiency of information aggregation can be summarized as follows. First, consider agents who are Bayesian. Suppose there is sufficient diversity of private information: there are at least two possible private signal precisions, and each individual is exposed to sufficiently many neighbors with each kind of signal. Then, under some technical conditions, information aggregation is as good as it can possibly be: each agent can figure out an arbitrarily good estimate of the previous period's state (which is the best information that she could hope to extract from others' actions), and then combine it with her own current private signal. Without sufficient diversity of private information—e.g., if all agents have the same kind of private signals—good information aggregation may fail. This can occur, as we explain below, even if each individual has access to very many neighbors' estimates, each of which is a noisy estimate of the recent state.

We can describe some important forces behind information aggregation and its failures at an intuitive level. Take a period-$(t+1)$ agent $(i, t+1)$, whose social information consists of the observed actions $a_{j,t}$ of very many agents (j, t). Imagine that each of these actions i observes, $a_{j,t} = \theta_t + \epsilon_{j,t} + \xi_t$, includes an idiosyncratic shock $\epsilon_{j,t}$ but is also confounded by a common piece of noise ξ_t. Even if the $\epsilon_{j,t}$ are uncorrelated or otherwise conducive to good learning, there is nothing that i can do to clean the correlated shock ξ_t out of her social information. Thus, there will only be a limited amount she can learn from her social information, no matter how many neighbors she observes.

We show that this obstruction arises when i's neighbors have symmetric private signal distributions. The role of the correlated shock ξ_t is played by the common dependence of everyone's action on the past before period t, which agent i would ideally like to ignore. When the neighbors of i have diverse signal distributions, however, there will no longer be a common shock ξ_t to all actions, so the recent state can be identified. The dependence of these conditions on signal endowments distinguishes the dynamic learning environment from the case of an unchanging state $\theta_t = \theta$, where good aggregation is the prediction independent of signal structure (DeMarzo et al., 2003). We give a more complete version of the ideas we have sketched in Section 4.

In addition to sufficient diversity of private signal distributions, something else is important to good aggregation: Bayesian behavior by individuals. In particular, even when the fundamental identification problem just described can be avoided,it is key that agents understand the correlation structure of their neighbors' estimates, and use this understanding to form an estimate that is not confounded by old values of the state. This involves subtracting some observations from others in order to "cancel out" the confounding. To make the point that this is necessary, we study the condition that agents do *not* use negative weights in their updating rules. This is the case in the well-known non-Bayesian heuristic learning rules in network models, most prominently the DeGroot model and its elaborations (Golub and Jackson, 2010; Jadbabaie et al., 2012). Under this condition, information aggregation is, in realistic networks, guaranteed to fall short of good aggregation benchmarks for all agents. This distinguishes our dynamic setting from models with an unchanging state; there, the papers we have just cited show that, even without much sophistication, agents are able to learn well.

Many of our results are asymptotic—the points are clearest where small numbers of signals are not an obstacle to aggregation. But the negative result we have just stated—that without anti-imitation it is not possible to learn well—has a counterpart in small networks, too. There, we can say that in any stationary equilibrium (or steady state of a more general form) on virtually any connected network, assuming agents put positive weights on all neighbors, the learning strategies agents use are necessarily Pareto inefficient.

Technical contributions of the model

At a technical level, several features of the model are important in our derivation of the results above: because of its Gaussian, stationary structure, it involves linear learning rules. We prove the existence of a stationary Bayesian learning equilibrium, and find that the coefficients agents place on their sources of information are a solution to a system of (nonlinear) polynomial equations in $O(n^2)$ variables (where n is the number of agents). When an agent uses her information to learn about the state, she must think about how precise and how correlated her signals about it are. Thus, it is very tractable numerically to analyze how these coefficients depend on the network and other parameters of the model. Indeed, we were able to conjecture our theorems about outcomes of learning because we can quickly compute equilibrium outcomes in networks of several thousand nodes.[4] Theoretical characterizations are more challenging, and require insights that cannot be read off from this system of equations. Nevertheless, the numerical tractability of the model allows us to do many computations to test the robustness of those theorems to relaxing technical assumptions (e.g., on network structure) that are needed for the proof techniques, but do not seem essential.

The numerical tractability of the model stands in contrast to the typical situation with Bayesian learning models in networks, and can facilitate several other kinds of exercises that are usually challenging:

1. Given suitable observations of agents' behavior, parameters (e.g., coefficients agents place on others) can be estimated readily from observed behavior. More precisely, an econometrician who observes a panel consisting only of the estimates of the agents and the realized states can use these data and a VAR model to recover the network structure and the precisions of the underlying signals available to agents in the network.

2. Testing against behavioral alternatives. Equilibrium is characterized by simple equations relating each agent's weights, and therefore her precision, to the precisions and

[4] The only computational limitation is matrix inversion.

correlations of her neighbors' estimates. Thus, the assumption of equilibrium weights can be tested against alternatives, such as naive rules.

3. Standard approaches to analyzing welfare: since agents' preferences for minimizing error are explicitly modeled and reflected in their behavior, the model is suited to standard welfare analysis based on revealed preferences.

4. Counterfactual analysis: since agents are maximizing their utilities, we can analyze how their learning rules will react to changes in the environment.

Outline

Section 2 sets up the basic model and discusses its interpretaiton. Section 3 defines our equilibrium concept and shows that equilibria exist. In Section 4, we give our main results on the quality of learning and information aggregation. In Section 5, we discuss learning outcomes with naive agents and more generally without anti-imitation. Section 6 relates our model and results to social learning literatures. In Section 7, we discuss structural estimation of our model, an extension to include endogenous information acquisition, and the role of Gaussian signals.

2 Model

2.1 Description

State of the world At each discrete instant (also called period) of time,

$$t \in \{\dots, -2, -1, 0, 1, 2, \dots\},$$

there is a state of the world, a random variable θ_t taking values in $\mathbb{R}$. This state evolves as an AR(1) stochastic process. That is,

$$\theta_{t+1} = \rho\theta_t + \nu_{t+1},$$

where $\rho \in (0, 1]$ is a constant and $\nu_{t+1} \sim \mathcal{N}(0, \sigma_\nu^2)$ are independent innovations. We can write explicitly

$$\theta_t = \sum_{\ell=0}^{\infty} \rho^\ell \nu_{t-\ell},$$

and thus $\theta_t \sim \mathcal{N}\left(0, \frac{\sigma_\nu^2}{1-\rho^2}\right)$. We make the normalization $\sigma_\nu = 1$ throughout.

Information and observations The set of *nodes* is $N = \{1, 2, \ldots, n\}$. Each node has a set $N_i \subseteq N$ of other nodes that i can observe, called its *neighborhood*.

Each node is populated by a sequence of *agents* in overlapping generations. At each time t, there is a node-i agent, labeled (i, t), who takes that node's action $a_{i,t}$. When taking her action, the agent (i, t) can observe the actions in her node's neighborhood in the m periods leading up to her decision. That is, she observes $a_{j,t-\ell}$ for all nodes $j \in N_i$ and "lags" $\ell \in \{1, 2, \ldots, m\}$. (One interpretation is that the agent (i, t) is born at time $t - m$ and has m periods to observe the actions taken around her before she acts.) She also sees a private signal,

$$s_{i,t} = \theta_t + \eta_{i,t},$$

where $\eta_{i,t} \sim \mathcal{N}(0, \sigma_i^2)$ has a variance $\sigma_i^2 > 0$ that depends on the agent but not on the time period. All the $\eta_{i,t}$ and ν_t are independent of each other. A vector of all of agent (i, t)'s observations—$s_{i,t}$ and the neighbors' past actions—defines her information. An important special case will be $m = 1$, where there is one period of memory, so that the agent's information is $(s_{i,t}, (a_{j,t-1})_{j \in N_i})$. The observation structure is common knowledge, and we will sometimes use the network G to mean the set of nodes N together with the set of *links* E, defined as the subset of pairs $(i, j) \in N \times N$ such that $j \in N_i$.

Preferences and best responses As stated above, in each period t, agent (i, t) at each node i chooses an action $a_{i,t} \in \mathbb{R}$. Utility is given by

$$u_{i,t}(a_{i,t}) = -(a_{i,t} - \theta_t)^2.$$

The agent makes the optimal choice for the current period given her information—i.e., does not seek to affect future actions.[5] By a standard fact about squared-error loss functions, given the distribution of $(\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m$, she sets:

$$a_{i,t} = \mathbb{E}[\theta_t \mid s_{i,t}, (\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m]. \tag{1}$$

Here the notation $\boldsymbol{a}_{N_i,t}$ refers to the vector $(a_{j,t})_{j \in N_i}$.Thus, an action can be interpreted as her estimate of the state, and we will sometimes use this terminology.

[5] A simple motivation for this assumption is that the agent makes only one decision during her lifetime. In Section 2.2 we discuss this asumption and how it relates to applications.

2.2 Interpretation

The agents are fully Bayesian given the information they have access to. Much of our analysis is done for arbitrary finite m; we view the restriction to finite memory as an assumption that avoids technical complications but, because m can be arbitrarily large, has little substantive content. The model generalizes "Bayesian without Recall" agents from the engineering and computer science literatures (e.g., Rahimian and Jadbabaie, 2017), which, within our notation, is the case of $m = 1$. Even when m is small, observed actions will indirectly incorporate signals from further in the past, and so can convey a great deal of information .

Note that an agent does not have access to the past private signals observed either at her own node or at neighboring ones. This is not a critical choice—our main results are robust to changing this assumption—but it is worth explaining. Whereas $a_{i,t}$ is an observable choice, such as a published evaluation of an asset or a mix of inputs actually used by an agent in production, the private signals are not shareable. Though we model the signals for convenience as real numbers, a more realistic interpretation of these is an aggregation of all of an agent's experiences, impressions, etc., and these may be difficult to summarize or convey.

Finally, our agents act once and do not consider future payoffs. These assumptions are made so that an agent's equilibrium action reflects her best current guess about the state, and turn off the possibility that she may try to strategically manipulate the future path of social learning (which could, in principle, help her successors). Substantively, like Gale and Kariv (2003) and Harel, Mossel, Strack, and Tamuz (2017)[6], we view these types of assumptions as a clean way of capturing that in our applications, such strategic considerations—if present at all—are likely to be secondary to matching the state. Equivalently, we could simply assume that agents sincerely announce their subjective expectations of the state, as in Geanakoplos and Polemarchakis (1982).

3 Equilibrium

In this section we present the substance of our notion of equilibrium and the basic existence result. Because time in this game is doubly infinite, there are some subtleties

[6]See also Manea (2011) and Talamàs (2017) in the bargaining literature.

3.1 Equilibrium in linear strategies

A strategy of an agent is *linear* if the action taken is a linear combination of the random variables in her information set (and a constant). We will focus on *stationary equilibria in linear strategies*—ones in which all agents' strategies are linear with time-invariant coefficients—though, of course, we will allow agents to consider deviating at each time to arbitrary strategies. Once we establish the existence of such equilibria, we will refer to them simply as *equilibria* for the rest of the paper.

We first argue that in studying agents' best responses to stationary linear strategies, we may restrict attention to linear strategies. If stationary linear strategies have been played up to time t, we can express each action up until time t as a summation of past signals.[7] Because all innovations ν_t and signal errors $\eta_{i,t}$ are independent and Gaussian, it follows the joint distribution of any finite random vector of the past errors $(a_{i,t-\ell'} - \theta_t)_{i\in N,\ell'\geq 1}$ is multivariate Gaussian. Thus, $\mathbb{E}[\theta_t \mid s_{i,t}, (\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m]$ is a linear function of $s_{i,t}$ and $(\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m$ (Kay, 1993, Example 4.4). It follows that solving for equilibrium can be reduced to searching for the weights agents place on the variables in their information set.

3.2 Covariance matrices

Having reduced the problem to one of choosing weights, our next observation is that the optimal weights for an agent to place on her sources of information depend on the accuracy and correlation of these sources, as we explain in this subsection and the next.

Let $\boldsymbol{V}_t$ be the $nm \times nm$ covariance matrix of the vector $(\rho^\ell a_{i,t-\ell} - \theta_t)_{\substack{i\in N\\ 0\leq\ell\leq m-1}}$. The entries of this vector are the distances between the actions in the past m periods (or, more precisely, the best predictors of θ_t given those actions) and the current state of the world. Denote the space of such covariance matrices by $\mathcal{V}$. In the case $m = 1$, the typical element is simply the covariance matrix $\boldsymbol{V}_t = \mathrm{Cov}(a_{i,t} - \theta_t)$. The matrix $\boldsymbol{V}_t$ records covariances of action errors: diagonal entries measure the accuracy of each action, while off-diagonal entries indicate how correlated the two agent's action errors are. The entries of $\boldsymbol{V}_t$ are $V_{ij,t}$.

[7]To ensure this series is almost surely convergent, note in any best response of agent (i, t), the random variable $a_{i,t} - \theta_t$, has a finite variance: each player seeks to minimize the variance of this error and always has the option of relying on her own private signal, in which case her error has finite variance.

3.3 Best-response weights

A strategy profile is an equilibrium if weights agents place on the variables in their information set minimize their posterior precision. We now characterize these in terms of the covariance matrices we have defined. Let $\boldsymbol{V}_{N_i,t-1}$ be a sub-matrix of $\boldsymbol{V}_{t-1}$ that contains only the rows and columns corresponding to neighbors of i[8] and let

$$\boldsymbol{C}_{i,t-1} = \begin{pmatrix} & & & 0 \\ & \boldsymbol{V}_{N_i,t-1} & & 0 \\ & & & \vdots \\ 0 & 0 & \ldots & \sigma_i^2 \end{pmatrix}.$$

Conditional on observations $(\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m$ and $s_{i,t}$, the state θ_t is normally distributed with mean

$$\frac{\mathbf{1}^T \boldsymbol{C}_{i,t-1}^{-1}}{\mathbf{1}^T \boldsymbol{C}_{i,t-1}^{-1} \mathbf{1}} \cdot \begin{pmatrix} \rho \boldsymbol{a}_{N_i,t-1} \\ \vdots \\ \rho^m \boldsymbol{a}_{N_i,t-m} \\ s_{i,t+1} \end{pmatrix} \tag{2}$$

(see Example 4.4 of Kay (1993)).[9] This gives $\mathbb{E}[\theta_t \mid s_{i,t}, (\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m]$ (recall that this is the $a_{i,t}$ the agent will play). Expression (2) is a linear combination of the agent's signal and the observed actions; the coefficients in this linear combination depend on the matrix $\boldsymbol{V}_{t-1}$ (but not on realizations of any random variables).

We denote by (W_t, w_t^s) the weights agents use in period t, with $w_t^s \in \mathbb{R}^n$ being the weights agents place on their private signals and W_t recording the weights they place on their other information. We do not describe the indexing of coefficients in W_t explicitly in general, but when $m = 1$, we refer to the weight agent i places on $a_{j,t-1}$, j's action yesterday, as $W_{ij,t}$ and the weight on $s_{i,t}$, her private signal, as $w_{i,t}^s$.

In view of the formula (2) for the optimal weights, we can compute the resulting next-period covariance matrix $\boldsymbol{V}_t$ from the previous covariance matrix. This defines a map $\Phi : \mathcal{V} \to \mathcal{V}$, given by

$$\Phi : \boldsymbol{V}_{t-1} \mapsto \boldsymbol{V}_t \tag{3}$$

[8] Explicitly, $\boldsymbol{V}_{N_i,t-1}$ are the covariances of $(\rho^\ell a_{j,t-\ell} - \theta_t)$ for all $j \in N_i$ and $\ell \in \{1, ..., m\}$.

[9] We implicitly assume here that agents' prior beliefs about the state are an improper distribution giving all states equal weight. With $\rho < 1$, agents' priors could instead be equal to the stationary distribution of the state and the analysis would not change substantially.

which we will study, for instance, in characterizing equilibria.

In the case of $m = 1$, we will write out the map explicitly using our notation for weights above:

$$\boldsymbol{V}_{ii,t} = (w^s_{i,t})^2\sigma_i^2 + \sum W_{ik,t}W_{ik',t}(\rho^2\boldsymbol{V}_{kk',t-1}+1) \text{ and } \boldsymbol{V}_{ij,t} = \sum W_{ik,t}W_{i'k',t}(\rho^2\boldsymbol{V}_{kk',t-1}+1). \tag{4}$$

There is an analogous (but more cumbersome) expression for $m > 1$, which we omit.

A natural question is how agents would, in practice, form the beliefs over others' play that are needed to carry out the inferences discussed above. The simplest foundation is that there is common knowledge (perhaps passed down from generation to generation) of the equilibrium covariance matrix $\boldsymbol{V}_{t-1}$, which is sufficient for playing optimally. Indeed, even if the covariance matrix is changing over time, the covariance matrix $\boldsymbol{V}_{t-1}$ of the previous period's errors is sufficient to compute the next period's $\boldsymbol{V}_t$ when agents play optimally. Thus there is a manageable information that must be carried from period to period. Alternately, at equilibrium the empirical covariances over many periods converge to the entries of the analytic covariance matrix. Intuitively, after many periods of play agents likely develop good estimates of how accurate each player is and how correlated two neighbors are. These estimates are sufficient to determine best responses.

3.4 Equilibrium existence

Consider the map Φ defined in (3). Stationary equilibria in linear strategies correspond to fixed points of the map Φ. More generally, one can use this map to study how covariances of actions evolve given any initial distribution of play. Note that $\boldsymbol{V}_t$ and the map Φ are deterministic, and we can study the resulting dynamical system without considering the particular realizations of signals.

Our first result concerns the existence of equilibrium:

Proposition 1. *There is a matrix $\widehat{\boldsymbol{V}}$ of covariances such that $\Phi(\widehat{\boldsymbol{V}}) = \widehat{\boldsymbol{V}}$. Equivalently, a stationary equilibrium in linear strategies exists.*

At the stationary equilibrium, the covariance matrix and all agent strategies are time-invariant. As in the DeGroot model, agent actions are linear combinations of observations with stationary weights (which we refer to as $\widehat{W}_{ij,t}$ and $\widehat{w}^s_i$). We discuss the relationship between our model and DeGroot learning further in Section 6.

The main insight is that we can find equilibria by studying action covariances, and this insight applies to many extensions of our model. We give two examples: (1) We assume that agents observe neighbors perfectly, but the existence result extends to other observation structures. These could include observing actions with noise and/or observing some set of linear combinations of neighbors' actions. (2) We assume agents are Bayesian, but the same proof shows equilibria exist for naive agents (see Section 5.1) or other behavioral rules such that actions depend continuously on observations and action variances are bounded.

The idea of the argument is as follows. First, all variances are bounded when agents best-respond to any beliefs about prior play, because all agents' actions must be at least as precise in estimating θ_t as their private signals, and cannot be more precise than estimates given perfect knowledge of yesterday's state combined with the private signal. Because the Cauchy-Schwartz inequality bounds covariances in terms of the corresponding variances, all entries of the image of Φ are bounded. Once these bounds are established, the result follows from the Brouwer fixed point theorem.

When $m = 1$, the proof gives bounds $\widehat{V}_{ii} \in [\frac{1}{1+\sigma_i^{-2}}, \sigma_i^2]$ on equilibrium variances and $\widehat{V}_{ij} \in [-\sigma_i\sigma_j, \sigma_i\sigma_j]$ on equilibrium covariances.

The equilibrium variances and covariances $\widehat{\boldsymbol{V}}$ satisfy

$$\widehat{V}_{ii} = (\widehat{w}_i^s)^2\sigma_i^2 + \sum \widehat{W}_{ik}\widehat{W}_{ik'}(\rho^2\widehat{V}_{kk'} + 1) \text{ and } \widehat{V}_{ij} = \sum \widehat{W}_{ik}\widehat{W}_{jk'}(\rho^2\widehat{V}_{kk'} + 1)$$

(or corresponding equations with $m > 1$). Substituting equation (2) and rearranging terms, we find the entries of $\widehat{\boldsymbol{V}}$ are the solutions to a system of polynomial equations. These equations have large degree and cannot be solved analytically except in very simple cases, but can be used to numerically solve for equilibria.

4 Conditions for fast information aggregation

In this section, we consider whether the agents are able to use social learning to form estimates that keep up with the evolution of the state. Because agents cannot learn a moving state exactly, we must define what it means for agents to learn well. Our benchmark is the action variance an agent would obtain given her private signal and perfect knowledge of the state in the previous period. The state in the previous period is the maximum that an agent can hope to learn from neighbors' information, since social information arrives with a one-period delay.

Definition 1. An equilibrium achieves the *ε-perfect aggregation benchmark* if

$$\widehat{V}_{ii}(\sigma_i^{-2}+1) \leq 1+\varepsilon$$

for all i.

This says that all agents do as well as if each knew her private signal and yesterday's state. The same notion of achieving the perfect aggregation benchmark can be formulated for any steady state, which need not come from rational agents optimizing.[10] Note agents can never infer yesterday's state perfectly from observed actions in any finite network, and so we must have $\widehat{V}_{ii}(\sigma_i^{-2}+1) > 1$ for all i on any fixed network.

We give conditions under which ε-perfect aggregation is achieved for ε small on large networks. To make this formal, we fix ρ and consider a sequence of networks $\{G_n\}_{n=1}^{\infty}$, where G_n has n nodes.

Example 1. We use a very simple example to demonstrate that the perfect aggregation benchmark can be achieved. Suppose each G_n for $n \geq 2$ has a connected component with two agents, 1 and 2, with $\sigma_1^2 = 1$ and $\sigma_2^2 = 1/n$ in G_n. Then agent 2's weight on her own signal converges to 1 as $n \to \infty$. So $\widehat{V}_{ii}$ converges to $(\sigma_1^{-2}+1)^{-1} = \frac{1}{2}$ as $n \to \infty$. Thus, the learning benchmark is achieved (for all agents).

The environment we have devised is quite special: it exogenously gives someone else having precise information, which allows agent 1 to infer last period's state. A much more interesting question is whether anything similar can occur without anyone having extremely precise signals. In the next section we address this and show that asymptotic learning can be achieved by all agents simultaneously even without anyone having very precise signals.

4.1 Diverse signals

This section studies an environment in which there are many agents, each having one of several possible types of signals. We show that in large networks, as agents' observation neighborhoods grow, the perfect aggregation benchmark is achieved (for all agents). We need to specify a model of large networks to proceed, and choose to consider the following stochastic block model.

[10] One can also extend the definition to cover cases where there is not a steady state, for example by taking the limsup also over times t.

Let $\{G_n\}_{n=1}^{\infty}$ be a sequence of undirected random networks, with G_n having n nodes. The agents have finitely many possible network types: let the nodes in network n be a disjoint union of $G_n^1, G_n^2, \ldots, G_n^K$. We say the agents in G_n^k have network type k. Suppose that any agent of network type k has an edge to any agent of network type k' with probability $p_{kk'}$, and these link realizations are independent. Assume that each network type observes at least one other network type with positive probability.

Suppose there are finitely many possible private signal variances. We say a signal variance is represented in a network type if, as $n \to \infty$, at least a positive share of agents of that type have that signal variance.[11] We assume that, for each network type, there are at least two distinct signal variances that are represented in that type.

Theorem 1. *Let $\varepsilon > 0$. Under the assumptions in this subsection, for large enough n each G_n has a equilibrium where the ε-perfect aggregation benchmark is achieved with probability at least $1 - \varepsilon$.*

So on large networks, society is very likely to aggregate information as well as possible. Agents learn well by anti-imitating neighbors with worse private signals to subtract out correlated information. The uncertainty is over the network, as there is always a small probability of a realized network which obstructs learning (e.g. if an agent has no neighbors).

The assumptions of finitely many signal and network types are purely technical, and could likely be relaxed.

Outline of Argument. To give intuition for the result, we first describe why the theorem holds on the complete network. We then discuss the obstacles and basic technique on general stochastic block networks.

We define agent i's *social signal* $r_{i,t}$ to be the optimal estimate of θ_{t-1} based on the actions of agent i's neighborhood. On the complete network, all players have the same social signal, which we call r_t.

At any equilibrium, each agent's actions is a weighted average of her private signal and this social signal:

$$a_{i,t} = \widehat{w}_s^i s_{i,t} + (1 - \widehat{w}_s^i) r_t.$$

The weight $\widehat{w}_s^i$ depends only on the precision of agent i's signal. We can all the weights of the two network types $\widehat{w}_s^A$ and $\widehat{w}_s^B$.

[11] Formally, the lim inf of the share of agents with that signal variance is positive.

By the law of large numbers, averaging a large number of private signals gives a very accurate estimate of the state. So by averaging the actions of each type, we obtain:

$$\sum_{i:\sigma_i^2=\sigma_A^2} a_{i,t} \approx \widehat{w}_s^A \theta_t + (1-\widehat{w}_s^A) r_t$$

$$\sum_{i:\sigma_i^2=\sigma_B^2} a_{i,t} \approx \widehat{w}_s^B \theta_t + (1-\widehat{w}_s^B) r_t$$

Because the two weights are distinct, we can solve for θ_t as a linear combination of the average actions from each type (up to signal error). This expression will anti-imitate the B types: the weight on $\sum_{i:\sigma_i^2=\sigma_B^2} a_{i,t}$ is negative. Using this procedure, any agent in period $t+1$ can obtain a very accurate estimate of θ_t. Since we are at an equilibrium, this means all agents will be very close to the efficient benchmark.

To use the same approach in general, we need to show that each individual observes a large number of neighbors of each signal type with similar social signals. More precisely, the proof shows that agents with the same network type have highly correlated social signals.

Because the social signals at an equilibrium are endogenous, it is difficult to show that two agents have similar social signals.

The key insight is that the number of paths of length two between any two agents is close to deterministic. While any two agents of the same network type have very different neighborhoods, their connections at distance two look similar. We can express any social signal as a combination of private signals and social signals from two periods earlier, and we can study the expression using this insight. Using this expression, we show that if agents of the same network type have similar social signals two periods ago, the same will hold in the current period. $\square$

We assume that agents know the signal types of their neighbors exactly, but this assumption could be relaxed. For example, if each agent were instead to receive only a noisy signal about each of her neighbor's signal types, she could solve her estimation problems in a similar way. By conditioning on the observable correlate of signal type, an agent could form enough distinct linear combinations reflecting the previous state and the current state to infer what she would like to know. Of course, in finite populations the precision of this inference would depend on the details.

Finally, the random graphs we study in this subsection have expected degrees that grow linearly in the population size, which may not be the desired asymptotic model.

While it is important to have neighborhoods "large enough" (i.e., growing in n) to permit the application of laws of large numbers, their rate of growth can be considerably slower than linear: for example, the our proof extends to degrees that scale as n^{α} for any $\alpha > 0$.[12]

4.2 Non-diverse signals

It is essential to the argument from the previous subsection that different agents have different signal precisions. From the perspective of an agent i learning at time $t+1$, the fact that type A and type B neighbors place different weights on the social signal allows i to back out the social signal used by his neighbors, and then to remove this confound. We now give examples where without diversity in signal quality, information aggregation is much slower.

We first define a class of very structured networks and show that for this class there is a unique equilibrium at which good aggregation is not achieved.

Definition 2. A network G has *symmetric neighbors* if $N_j = N_{j'}$ for any $j, j' \in N_i$.

In the undirected case, the graphs with symmetric neighbors are the complete network and complete bipartite networks (which are both special cases of our stochastic block model from Section 4.1). For directed graphs, the condition allows a larger variety of networks.

Consider a sequence $\{G_n\}_{n=1}^{\infty}$ of strongly connected graphs with symmetric neighbors. Assume that all signal qualities are the same, equal to σ^2, and that $m = 1$.

Proposition 2. *Under the assumptions in the previous paragraph, each G_n has a unique equilibrium. There exists $\varepsilon > 0$ such that the ε-perfect aggregation benchmark is not achieved at this equilibrium for any n.*

Indeed, all agents are bounded away from our learning benchmark at the unique equilibrium. So all agents learn poorly compared to the diverse signals case.

When the informational environment is symmetric, agents no longer avoid including correlated old information in their actions. Although the effect of signal errors $\eta_{i,t}$ vanishes as n grows large, the correlated error from past changes in the state ν_t prevents perfect aggregation. This is easiest to see on the complete graph, where all actions are exchangeable and so all social signals are an unweighted average of actions from the previous period. So the social signals place substantial weight on information from many periods ago.

[12] Instead of studying Φ^2 and second-order neighborhoods, we apply the same analysis to Φ^k and k^{th}-order neighborhoods for k larger than $1/\alpha$.

To give intuition for the result more generally, suppose that all agents were very close to the benchmark at equilibrium. Then all agents' actions are nearly exchangeable, so all actions would be close to unweighted averages of observations. As in the complete graph case, information from many periods in the past would contaminate estimates.

If the degrees of agents converge to infinity uniformly, then we can explicitly characterize the limit action variances and covariances. Define constants V^∞ and Cov^∞ by

$$V^\infty = \frac{1}{\sigma^{-2} + (\rho^2 Cov^\infty + 1)^{-1}}, Cov^\infty = \frac{(\rho^2 Cov^\infty + 1)^{-1}}{[\sigma^{-2} + (\rho^2 Cov^\infty + 1)^{-1}]^2}. \tag{5}$$

One way to interpret V^∞ and Cov^∞ is as the limit as n grows large of the variances and covariances, respectively, at the symmetric equilibrium of a complete graph with $n$agents.

The same argument as in the proof of Proposition 3 shows that with symmetric neighbors and growing degrees, the variances of all agents converge to V^∞ and the covariances of all pairs of agents converge to Cov^∞. This implies that the equilibrium action distributions are close to symmetric. These actions are equal to an appropriately discounted sum of past states, up to signal error terms which vanish asymptotically.

As a consequence of the theorem, we can give an example where making one agent's private information less accurate helps all agents.

Corollary 1. *There exists a network G and an agent $i \in G$ such that increasing σ_i^2 gives a Pareto improvement in equilibrium variances.*

To prove the corollary, we consider the complete graph with homogeneous signals and n large. By Theorem 2, all agents do substantially worse than perfect aggregation. If we give agent 1 a very uninformative signal, all players can anti-imitate agent 1 and almost achieve perfect aggregation. When the initial signals are not too accurate, this gives a Pareto improvement.

We next show that on Erdos-Renyi random networks, there is an equilibrium with the same learning outcomes when signals precisions are homogeneous. Let $\{G_n\}_{n=1}^\infty$ be a sequence of undirected random networks, with G_n having n nodes, with any pair of distinct nodes linked (i.i.d.) with probability p. We continue to assume all signal variances are equal to σ^2 and $m = 1$.

Proposition 3. *Under the assumptions in the previous paragraph, for large enough n there exists an equilibrium on G_n where the perfect aggregation benchmark is not achieved with probability at least $1 - \varepsilon$.*

The equilibrium covariances again converge to V^∞ and Cov^∞ (for any value of p). So we obtain the same learning outcomes asymptotically on a variety of networks.

4.3 The benchmark in finite populations

The results of the previous section can be summarized as saying that to achieve the aggregation benchmark of essentially knowing the previous period's state there need to be at least two different private signal variances in the network. Formally, this is a knife-edge result: as long as private signal variances differ by at least some amount, as $n \to \infty$, perfect aggregation is achieved; with homogeneous signal endowments, agents' variance is much higher. In this section, we show numerically that for fixed values of n, the transition from the first regime to the second is actually gradual, with error well above the perfect aggregation benchmark when signal qualities differ slightly.

In Figure 1, we study the complete network. The private signal variance of agents of network type 1 is fixed at $\sigma_A^2 = 2$. We then vary the private signal variance of agents of type 2 (the horizontal axis), while measuring the prediction quality of agents of type 1 (plotted on the vertical axis). The variance of type 1 agents at the benchmark is $2/3$. We note several relevant features. First, the change in aggregation is continuous and, indeed, reasonably gradual for n in the hundreds as we vary σ_2. Second, as n increases, we can see that the curve is moving toward the theoretical limit: a spike at $\sigma_B^2 = 2$. Third, there are nevertheless considerable gains to increasing n, the number of agents: going from $n = 200$ to $n = 600$ results in a gain of in precision when $\sigma_B^2 = 3$.

5 The importance of anti-imitation

In the proof of our positive result on achieving the perfect aggregation benchmark, Theorem 1, a key step involved agents placing negative weights on some neighbors' estimates to solve their signal extraction problem. In this section, we demonstrate that this anti-imitation is indeed essential for perfect aggregation.

We first consider a particular model of naive agents who do not account for correlation and show such agents fall far short of perfect aggregation. We then formalize the idea that anti-imitation is crucial to reaching the benchmark by documenting a more general failure of learning when all weights are positive, whether agents are rational or naive. Finally, we show that even in fixed finite networks any positive weights chosen by optimizing agents

Figure 1: Distinct Variances Result in Learning

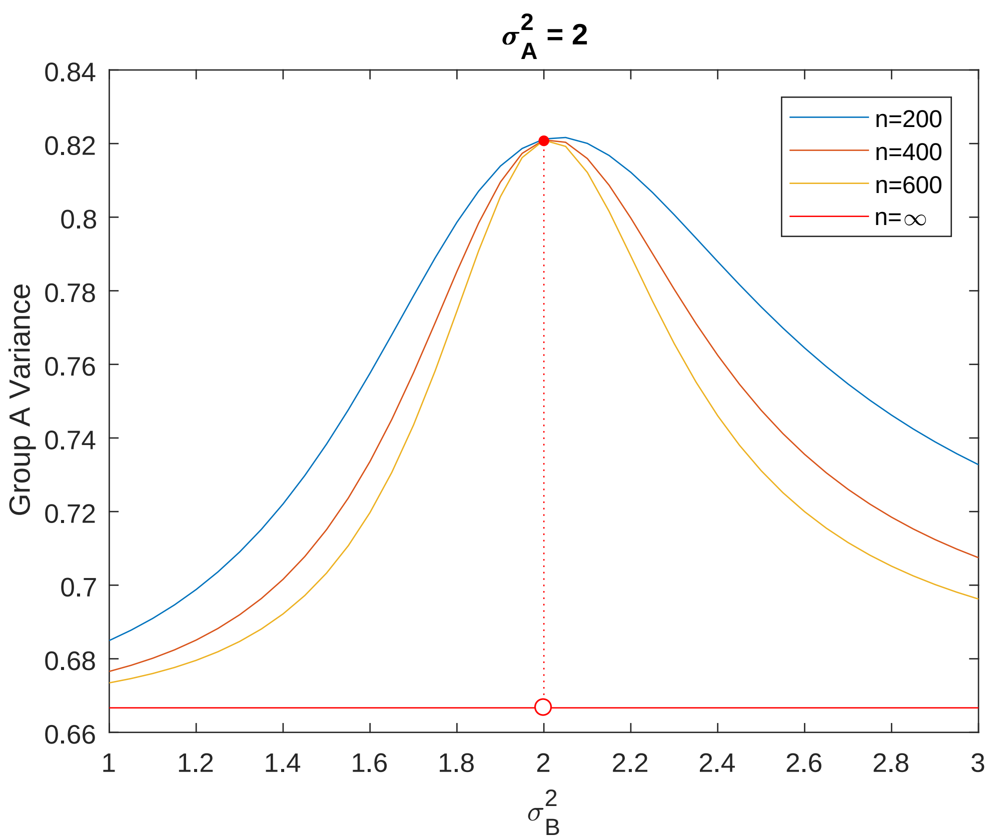

will be Pareto-dominated.

5.1 Naive agents

In this part we introduce agents who observe and perceive their neighbors' actions, but do not process those observations in a fully Bayesian manner. We consider a particular misspecification that simplifies solving for equilibria analytically. There are a number of possible variants of our behavioral assumption, and it is straightforward to numerically study alternative specifications of behavior in our model (Alatas et al. 2016 consider one such variant).

Our notion of naiveté follows:

Definition 3. We call an agent *naive* if she believes that all neighbors choose actions equal to their private signals and maximizes her expected utility given these incorrect beliefs.

Equivalently, a naive agent believes her neighbors all have empty neighborhoods. This is the analogue of best-response trailing naive inference (Eyster and Rabin, 2010) in our dynamic model.[13] So naive agents understand that their neighbors' actions from the previous period are estimates of θ_{t-1}, but think these estimates are independent and incorrectly perceive their variances.

With multiple signal qualities, learning outcomes depend more fundamentally on the network: even if she observes a large number of agents of each signal type, a naive agent's learning outcome will depend substantially on the relative numbers of neighbors of each signal type.

We will describe outcomes with two signal types σ_A^2 and σ_B^2. We use the same random network model as in Section 4.1 and assume each network type contains equal shares of agents with each signal type.

We can define variances

$$V_A^\infty = \frac{\kappa_t^2 + \sigma_A^{-2}}{\left(1+\sigma_A^{-2}\right)^2}, V_B^\infty = \frac{\kappa_t^2 + \sigma_B^{-2}}{\left(1+\sigma_B^{-2}\right)^2} \tag{6}$$

where

[13] An alternate interpretation of best-response trailing naive inference might be that a naive agent treats all actions as independent signals about today's state. This implies the degenerate behavior on large networks: an agent's action converges to distribution to zero as her degree grows large.

$$\kappa_t^{-2} = 1 - \frac{\rho^2}{(\sigma_A^{-2} + \sigma_B^{-2})} \left(\frac{\sigma_A^{-2}}{1+\sigma_A^{-2}} + \frac{\sigma_B^{-2}}{1+\sigma_B^{-2}} \right).$$

Naive agents' equilibrium variances converge to these values.

Proposition 4. *Let $\varepsilon > 0$. Under the assumptions in this subsection, there is a unique equilibrium on G_n. For large enough n all agents' equilibrium variances are within ε of V_A^∞ and V_B^∞ with probability at least $1-\varepsilon$. Then the perfect aggregation benchmark is not achieved, and when $\sigma_A^2 = \sigma_B^2$ all agents' variances are larger than V^∞.*

Aggregating information well requires a sophisticated response to the correlations in observed actions. Because naive agents completely ignore these correlations, their learning outcomes are poor. In particular their variances are larger than at the equilibria we discussed in the Bayesian case, even when that equilibrium is inefficient ($\sigma_A^2 = \sigma_B^2$).

When signal qualities are homogeneous ($\sigma_A^2 = \sigma_B^2$), we obtain the same limit on any network with enough observations. That is, on any sequence of (deterministic) networks G_n with agents' degrees converging to infinity uniformly and any sequence of equilibria on G_n, the equilibrium action variances of all agents converge to V_A^∞.

In Figure 2, we compare Bayesian and naive learning outcomes. As in Figure 1, we consider a complete network where half of agents have signal variance $\sigma_A^2 = 3$ and we vary the signal variance σ_B^2 of the remaining agents. We observe that naive agents learn substantially worse than rational agents, even when signals are not diverse. Our explicit formulas for variances under Bayesian and naive learning also allow for more general comparisons in the limit as $n \to \infty$.

5.2 Anti-imitation is essential for reaching the benchmark

Naive agents fail to anti-imitate some neighbors to account for correlation in observations, and as a result do not learn well. We now show more generally that anti-imitation is necessary for good learning: given a sequence of undirected networks along with a (rational or naive) equilibrium on each network where where weights are positive, we show that these equilibria are bounded away from perfect aggregation.

Proposition 5. *Consider a sequence of undirected networks $G_1, G_2, \ldots$ with n agents in G_n and assume that all private signal variances are bounded below by $\underline{\sigma}^2 > 0$.*

1. *Suppose all agents are Bayesian. Consider any sequence of equilibria on G_n in which all agents are using positive weights.*

Figure 2: Bayesian and Naive Learning

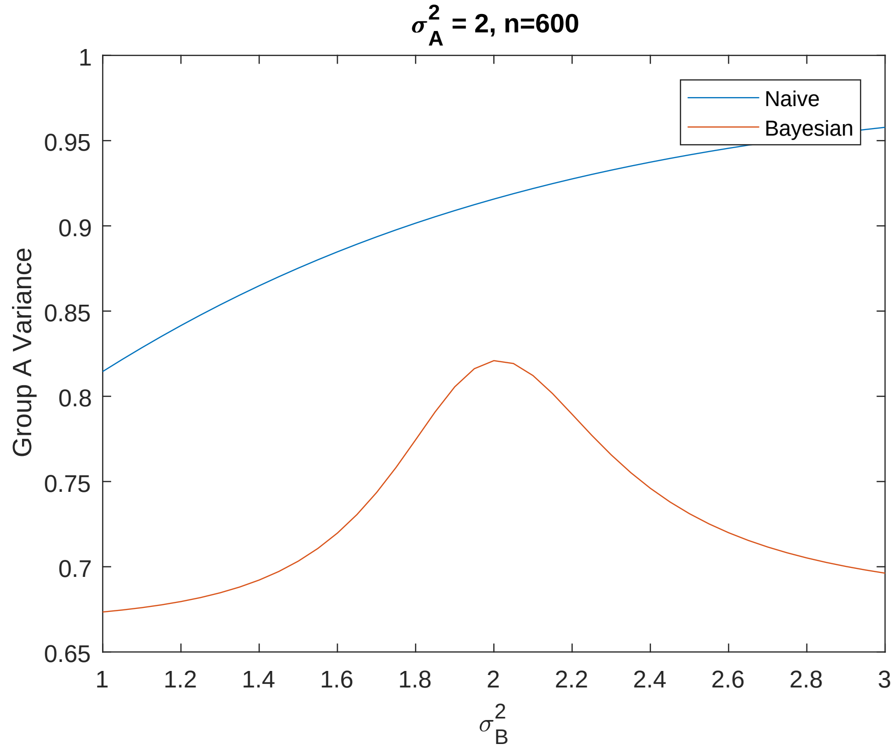

2. Suppose all agents have are naive, and consider any sequence of equilibria on G_n.

In either case, there is an $\epsilon > 0$ such that, for all n, the ϵ-aggregation benchmark is not achieved.

The essential idea is that at time $t+1$ observed time-t actions all put weight on actions from period $t-1$, which causes θ_{t-1} to have a (positive weight) contribution to all observed actions. Agents do not know θ_{t-1} and, with positive weights, cannot take any linear combination that would recover it. Even with a very large number of observations, this confound prevents agents from learning yesterday's state precisely.

To make the argument more precise, assume toward a contradiction that agent i achieves the ϵ-perfect aggregation benchmark for an arbitrarily small ϵ. Because of the confounding discussed in the last paragraph, she would have to observe many neighbors who place almost all of their weight on their private signals. Because the network is undirected, though, these neighbors themselves see i. Since i's actions in this hypothetical reflects the state very accurately, the neighbors would do better by placing substantial weight on agent i and *not* just on their private signals. So we cannot have such an agent i.

In summary, bidirectional observation presents a fundamental obstruction to attaining the best possible benchmark of aggregation. This is related to a basic observation about learning from multivariate Gaussian signals about a parameter: if the signals (here, social observations), conditional on the state of interest (θ_t) are all correlated and the correlation is bounded below, away from zero, (here this occurs because all involve some indirect weight on θ_{t-2}) then the amount one can learn from these signals is bounded, even if there are infinitely many of them. A related observation plays a role in Harel et al. (2017)—we discuss this more in Section 6.

5.3 Without anti-imitation, outcomes are Pareto-inefficient

The previous section argued that anti-imitation is critical to achieving the perfect aggregation benchmark. We now show that even in small networks, where that benchmark is not relevant, any equilibrium without anti-imitation is Pareto-inefficient relative to another steady state. This result complements our asymptotic analysis by showing a different sense (relevant for small networks) in which anti-imitation is necessary to make the best use of information.

To show the result, we will define an alternative to equilibrium weights, and study the associated stationary outcome. To make this more formal, we make the following definition:

Definition 4. The *steady state* associated with weights $\boldsymbol{W}$ and $\boldsymbol{w}^s$ is the (unique) covariance matrix $\boldsymbol{V}^*$ such that if actions have a variance-covariance matrix given by $\boldsymbol{V}_t = \boldsymbol{V}^*$ and next-period actions are set using weights $(\boldsymbol{W}, \boldsymbol{w}^s)$, then $\boldsymbol{V}_{t+1} = \boldsymbol{V}^*$ as well.

In this definition of steady state, instead of optimizing (as at equilibrium) agents use fixed weights in all periods. By a straightforward application of the contraction mapping theorem, any non-negative weights under which covariances remain bounded at all times determine a unique steady state.

Theorem 2. *Suppose the network G is strongly connected. Consider weights $\boldsymbol{W}$ and $\boldsymbol{w}^s$ and suppose they are all positive, with an associated steady state $\boldsymbol{V}_t$. Suppose either*

(1) there is an agent i whose weights are a Bayesian best response to $\boldsymbol{V}_t$, and some agent observes that agent and at least one other neighbor; or

(2) there is an agent whose weights are a naive best response to $\boldsymbol{V}_t$, and who observes multiple neighbors.

Then the steady state $\boldsymbol{V}_t$ is Pareto-dominated by another steady state.

We state the result under relatively weak hypotheses on behavior, but clearly it implies that under mild conditions under the network, either a Bayesian or naive equilibrium will be Pareto-inefficient, as long as weights are all positive. Thus:

Corollary 2. *Suppose the network G is strongly connected and some agent has more than one neighbor. If all weights are positive at a either a Bayesian or naive equilibrium, then the variances at that equilibrium are Pareto-dominated by variances at another steady state.*

The basic argument behind Theorem 2 is that if agents place marginally more weight on their private signals, this introduces more independent information that eventually benefits everyone. Special cases include equilibria, where all agents are rational, or naive equililbria.

In a review of sequential learning experiments, Weizsäcker (2010) finds that subjects weight their private signals more heavily than is optimal (given the empirical behavior of others they observe). Theorem 2 implies that in our environment with optimizing agents, it is actually welfare-improving for individuals to "overweight" their own information relative to best-response behavior.

Discussion of conditions in the theorem We next briefly discuss the sufficient conditions in the theorem statement. It is clear that some condition on neighborhoods is needed: If every agent has exactly one neighbor and updates rationally or naively, there are no externalities and the equilibrium weights are Pareto optimal. In fact, the same result (with the same proof) applies to a larger class of networks: it is sufficient that, starting at each agent, there are two paths of some length k to a rational agent and another distinct agent. Finally, the condition on equilibrium weights says that no agent anti-imitates any of her neighbors. This assumption makes the analysis tractable, but we believe the basic intuition carries through in finite networks with some anti-imitation.

Proof sketch The idea of the proof of the rational case is to begin at the steady state and then marginally shift the rational agent's weights toward their private signal. By the envelope theorem, this means agents' actions are less correlated but not significantly worse in the next period. We show that if all agents continue using these new weights, the decreased correlation eventually benefits everyone. In the last step, we use the absence of anti-imitation, which implies that the updating function associated with agents using fixed weights (instead of an optimization procedure) is monotonic. To first order, some covariances decrease while others do not change after one period under the new weights.

Monotonicity of the updating function and strong connectedness imply that eventually all agents' variances decrease.

The proof in the naive case is simpler. Here a naive agent is overconfident about the quality of her social information, so she would benefit from shifting some weight from her social information to her signal. This deviation also reduces her correlation with other agents, so it is Pareto-improving.

An illustration An example to illustrates the phenomenon:

Example 2. Consider $n = 100$ agents in an undirected circle—i.e., each agent observes the agent to her left and the agent to her right. Let $\sigma_i^2 = \sigma^2$ be equal for all agents and $\rho = .9$. The equilibrium strategies place weight $\widehat{w}^s$ on private signals and weight $\frac{1}{2}(1 - \widehat{w}^s)$ on each observed action.

When $\sigma^2 = 10$, the equilibrium weight is $\widehat{w}^s = 0.192$ while the welfare-maximizing symmetric weights have $w^s = 0.234$. That is, weighting private signals substantially more is Pareto improving. When $\sigma^2 = 1$, the equilibrium weight is $\widehat{w}^s = 0.570$ while the welfare maximizing symmetric weights have $w^s = 0.586$. The inefficiency persists, but the equilibrium strategy is now closer to the optimal strategy.

6 Related literature

We now put our contribution in the context of the extensive literatures on social learning and learning in networks.[14]

6.1 DeGroot learning

Play in the stationary linear equilibria of our model closely resembles behavior in the DeGroot (1974) heuristic, where social learning occurs by linearly aggregating neighbors' past estimates, with constant weights on neighbors over time.[15] This may be viewed as a simple way of aggregating information in a complex world, without considering various possible realizations that might underlie neighbors' behavior. The studies that introduced such heuristics did not posit an informational environment in the standard game-theoretic

[14] Golub and Sadler (2016) survey parts of this literature.

[15] It is also closely related to the model of Friedkin and Johnsen (1997), who allow each individual to place weight on a "personal" estimate of his own. A more flexible specification, which allows for information arriving over time, is studied by Jadbabaie et al. (2012) and Molavi et al. (2018).

sense, nor did they analyze agents' rules in terms of their performance relative to an objective—but studies since then have done both (French (1956); Lehrer and Wagner (1981)).

DeMarzo, Vayanos, and Zweibel (2003) provided foundations for the heuristic assuming agents have an oversimplified model of their environment. In their model, the state is drawn once and for all at time zero, and each agent receives one signal about it; thenagents repeatedly observe each other's conditional expectations of the state and form estimates. At time zero, assuming all randomness is Gaussian, the Bayesian estimation rule is linear with certain weights. DeMarzo, Vayanos, and Zweibel (2003) made the behavioral assumption that in subsequent periods, agents treat the informational environment as being identical to that of the first period (even though past learning has, in fact, induced redundancies and correlations). In that case, the agents behave according to the DeGroot rule, using the same weights over time. We give an alternative, Bayesian microfoundation for the same sort of rule by studying a different environment. Our foundation relies on the fact that the environment is stationary and so, in fact, the joint distribution of the random variables in the model (neighbors' estimates and the state of interest) is actually stationary.[16]

In terms of performance of the DeGroot learning rule, DeMarzo, Vayanos, and Zweibel (2003) emphasized that in their model, the stationary rule could be far from optimal: Truly Bayesian agents with a correct model of the environment would be able to infer the state exactly, achieving the best possible payoff for all but finitely many periods, while DeGroot agents could be far from achieving this. Golub and Jackson (2010) asked whether, in the same informational environment, DeGroot agents can nevertheless do reasonably well in large networks due to laws of large numbers. They show that as long as no agent has too prominent a network position, good aggregation (negligible welfare loss for each agent) is guaranteed after sufficient learning, even if agents are naive. Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012) look at a variant of the DeGroot rule in which agents receive new signals in each period. Because the state remains fixed in their model, stationary strategies are not Bayesian best responses. Nevertheless, they find that under fairly mild conditions, their variant of the DeGroot rule leads to long-run learning aessentially as good as if agents knew all signals.

To summarize, in the fixed-state environments, certain simple heuristics, requiring no sophistication about correlations between neighbors' behavior, can perform quite well: at least in large networks, they allow agents to guess the state quite precisely. The success of

[16]Indeed, agents behaving according to the DeGroot heuristic even when it is not appropriate may have to do with their experiences in stationary environments where it is closer to optimal.

these rules relies several features of the environments in which they are derived. Our results highlight that an unchanging state is an important ingredient. Making the state dynamic causes many simple heuristics to fail. As we show in Theorem 1, to learn well agents need a sophisticated response to the correlation in neighbors' estimates that arises from those neighbors' past learning. Moreover, even Bayesian agents are not guaranteed to be able to aggregate information well, as we show in the homogeneous-signal environment of Section 4.2. Overall, the message for learning, especially non-Bayesian learning, is more complex than in existing models and more sensitive to details of the environment. The tractability of the learning rules makes the predictions particularly amenable to testing as we discuss below in Section 7.1.

6.2 Recent models with evolving states

Several recent papers in computer science and engineering study environments similar to ours. Frongillo, Schoenebeck, and Tamuz (2011) study (in our notation) a θ_t that follows a random walk ($\rho = 1$). They examine agents who learn using fixed (exogenous) weights on arbitrary networks, where they characterize the steady-state distribution of behavior with arbitrary (non-equilibrium) fixed weights on any network. They also examine best-response (equilibrium) weights on a complete network, where all agents observe all of yesterday's actions. Their main result concerning these is to show that the equilibrium weights can be inefficient. This is generalized by our Theorem 2 on Pareto-inefficiency on an arbitrary graph. Our existence result (Proposition 1) generalizes the construction in their paper from the symmetric case of the complete network to arbitrary networks.

The stochastic process and information structure in Shahrampour, Rakhlin, and Jadbabaie (2013) are also the same as ours, though their analysis does not consider optimizing agents. The authors consider a class of fixed weights and study heuristics, computing or bounding various measures of welfare. When we study Pareto inefficiency, we compare welfare under such fixed exogenous weights with the welfare obtained by optimizing agents at equilibrium. Because our model endogenously selects weights, we can consider how it responds endogenously to changes in the environment (e.g. network). We also give conditions for good learning when agents are optimizing for themselves, rather than for a global objective. In economics, the model in Alatas, Banerjee, Chandrasekhar, Hanna, and Olken (2016) most closely resembles ours. There, agents are not fully Bayesian, ignoring the correlation between social observations. The model is estimated using data on social learning in Indonesian villages, where the state variables are the wealths of villagers. As we show, how rational agents are in their inferences plays a major role in the accuracy of aggregation processes. Our model provides foundations for structural estimation with Bayesian behavior as well as testing of the Bayesian model against behavioral alternatives.

Our agents are selecting optimal weights for their own learning from Gaussian signals, but they, in turn, affect the behavior of the process since others observe them. Thus, there is a tight connection to a large engineering literature on the Kalman filter. However, that literature does not consider nodes who are optimizing in view of each other's behavior—i.e., solution concepts such as our stationary equilibrium.

6.3 Sequential social learning models

A canonical model of social learning involves infinitely many agents choosing, in sequence, from finitely many (often two) actions to match a fixed state,with access to predecessors' actions (Bikhchandani, Hirshleifer, and Welch (1992); Banerjee (1992); Smith and Sørensen (2000); Eyster and Rabin (2010)). The first models were worked out with observation of *all* predecessors, but recent papers have developed analyses where some *subset* of predecessors seen by each agent Acemoglu et al. (2011); Eyster and Rabin (2014); Lobel and Sadler (2015a,b), which can be thought of as a "unidirectional" network structure, where directed links correspond to observation opportunities.

A major question in this literature is how information aggregation can stop after some finite time due to inference problems. The discreteness of individuals' actions often plays an important role. Our focus is quite different in terms of the questions: we study a moving continuous state and continuous actions, and ask how well agents aggregate information, in steady state, about the relatively recent past. These modeling differences allow new insights to emerge: for example, heterogeneity of signal endowments turns out to be critical to good aggregation in the Bayesian case, which is very different from the kinds of conditions that play a role in Smith and Sørensen (2000); Acemoglu et al. (2011); Lobel and Sadler (2015a).[17] Another difference concerns the modeling of the network: Our agents are at the nodes of a finite, unchanging network, and there is bidirectional observation over time: node A learns from node B, which then learns from node A. Thus, interesting "feedback" considerations can emerge in our stationary linear equilibria, which are absent in standard sequential models.

[17] Those conditions require either that some signals are very informative, or that some people see many independent signals—neither of which we assume in our main results.

Despite the modeling differences, the inefficiencies that drive information cascades are related to the ones identified by our Theorem 2: both occur because agents place less weight on their private signals than is socially optimal. Because the learning weights are described by fixed points, we can study structural features of the informational environment that affect the magnitude of the inefficiency, as illustrated by Corollary 1.

Another point of contact concerns the modeling of changing states: Moscarini, Ottaviani, and Smith (1998) (see also van Oosten (2016)) study learning models where the binary state evolves as a two-state Markov chain. Their results focus largely on the frequency and dynamics of cascades: changes in the state can break cascades/herds and renew learning. Our main focus is on the aggregation properties.

A robust aspect of rational learning in sequential networks is anti-imitation. Eyster and Rabin (2014) give general conditions for fully Bayesian agents to anti-imitate in the sequential model. We find that anti-imitation also is a robust feature in our dynamic model, and in our context is crucial for good learning. Despite this similarlity, thre is a sharp contrast between our findings and standard sequential models. In those models, while rational agents *do* prefer to anti-imitate, in many cases individuals agents as well as society as a whole could obtain good outcomes using heuristics without any anti-imitation: for instance, by combining the information that can be inferred from one neighbor with one's own private signal, as in Acemoglu, Dahleh, Lobel, and Ozdaglar (2011). Our dynamic learning problem is fundamentally different, as shown in Theorem 1: agents must respond in a sophisticated way, with anti-imitation, to their neighbors' (correlated) learning in order to reach aggregation benchmarks.

7 Discussion and extensions

7.1 Identification and testable implications

One of the main advantages of the parametrization we have studied is that standard methods can eaisly be applied to estimate the model and test hypotheses within it. The key feature making the model econometrically well-behaved is that, in the solutions we focus on, agents' actions are linear functions of the random variables they observe. Moreover, the evolution of the state and arrival of information creates exogenous variation. We briefly sketch how these features can be used for estimation and testing.

Assume the following. The analyst obtains noisy measurements $\overline{a}_{i,t} = a_{i,t} + \xi_{i,t}$ of agent's actions (where $\xi_{i,t}$ are i.i.d., mean-zero error terms). He knows the parameter ρ governing the stochastic process, but may not know the network structure or the qualities of private signals $(\sigma_i)_{i=1}^n$. Suppose also that the analyst observes the state θ_t ex post (perhaps with a long delay).[18]

Now, consider *any* steady state in which agents put constant weights W_{ij} on their neighbors and w_i^s on their private signals over time. We will discuss the case of $m = 1$ to save on notation, though all the statements here generalize readily to arbitrary m.

We first consider how to estimate the weights agents are using, and to back out the structural parameters our model when it applies. The strategy does not rely on uniqueness of equilibrium. We can identify the weights agents are using through standard vector autoregression methods. In steady state,

$$\overline{a}_{i,t} = \sum_j W_{ij}\rho\overline{a}_{j,t-1} + w_i^s\theta_t + \zeta_{i,t}, \tag{7}$$

where $\zeta_{i,t} = w_i^s\eta_{i,t} - \sum_j W_{ij}\rho\xi_{j,t-1} + \xi_{i,t}$ are error terms i.i.d. across time. The first term of this expression for $\zeta_{i,t}$ is the error of the signal that agent i receives at time t. The summation combines the measurement errors from the observations $\widetilde{a}_{j,t}$ the previous period.[19] Thus, we can obtain consistent estimators $\widetilde{W}_{ij}$ and $\widetilde{w}_i^s$ for W_{ij} and w_i^s, respectively.

We now turn to the case in which agents are using *equilibrium* weights. First, and most simply, our estimates of agents' equilibrium weights allow us to recover the network structure. If the weight $\widehat{W}_{ij}$ is non-zero for any i and j, then agent i observes agent j. Generically the converse is true: if i observes j then the weight $\widehat{W}_{ij}$ is non-zero. Thus, tests of whether the recovered social weights are nonzero generically identify network links. For such tests, and generally, the standard errors in the estimators can be obtained by standard techniques.[20]

Now we examine the more interesting question of how structural parameters can be identified assuming an equilibrium is played, and also how to test the assumption of equilibrium.

The first step is to compute the empirical covariances of action errors from observed data; we call these $\widetilde{V}_{ij}$. Under the assumption of equilibrium, we now show how to deter-

[18] We can instead assume that the analyst observes (a proxy for) the private signal $s_{i,t}$ of agent i; we mention how below.

[19] This system defines a VAR(1) process (or generally VAR(m) for memory length m).

[20] Methods involving regularization may be practically useful in identifying links in the network. Manresa (2013) proposes a regularization (LASSO) technique for identifying such links (peer effects). In a dynamic setting such as ours, with serial correlation, such techniques will generally be more complicated.

mine the signal variances using the fact that equilibrium is characterized by $\Phi(\widehat{\boldsymbol{V}}) = \widehat{\boldsymbol{V}}$ and recalling the explicit formula (4) for Φ. In view of this formula, the signal variances σ_i^2 are uniquely determined by the other variables:

$$\widehat{V}_{ii} = \sum_j \sum_k \widehat{W}_{ij} \widehat{W}_{ik} (\rho^2 \widehat{V}_{jk} + 1) + (\widehat{w}_i^s)^2 \sigma_i^2. \tag{8}$$

Replacing the model parameters other than σ_i^2 by their empirical analogues, we obtain a consistent estimate $\widetilde{\sigma}_i^2$ of σ_i. This estimate could be directly useful—for example, to an analyst who wants to choose an "expert" from the network and ask about her private signals directly.

Note that our basic VAR for recovering the weights relies only on constant linear strategies and does not assume that agents are playing any particular strategy within this class. Thus, if agents are using some other behavioral rule (e.g., optimizing in a misspecified model) we can replace (8) by a suitable analogue that reflects the bounded rationality in agents' inference. If that such steady state exists, and using the results in this section, one can create an econometric test that is suitable for testing how agents are behaving. For instance, we can test the hypothesis that they are Bayesian against the naive alternative of our section .

7.2 Signal acquisition

In this section we analyze what would happen if agents were to choose in every period their precision $\sigma_{i,t}^{-2}$. We assume that there is a convex cost function in precisions $c(\sigma_{i,t}^{-2})$.

Corollary. *Let G be a complete network, and with each agent maximizing*

$$u_{i,t}(a_{i,t}, \sigma_{i,t}) = -(a_{i,t} - \theta_t)^2 - c(\sigma_{i,t}^{-2}).$$

There is a unique equilibrium in which $\sigma_{i,t}^2 = \sigma_{j,t}^2$ for all i, j, and t, and there is no asymptotic learning.

The proof is simple–since the network is complete all agents have the same social signal $r_{i,t} = r_{j,t}$. Therefore, their maximization is identical and they will choose the same $\sigma_{i,t}^{-2}$, we also know that any complete network has a unique steady state hence the equilibrium is unique.

This lemma shows that while the result of non-diverse signals may seem like an edge-knife case, it is endogenously selected as the equilibrium of a signal acquisition game.

7.3 General distributions

Our analysis of stationary linear learning rules relied crucially on the assumptions that the innovations ν_t and signal errors $\eta_{i,t}$ are Gaussian random variables. However, we believe the basic logic of our result about good aggregation with signal diversity (Theorem 1) does not depend on this particular distributional assumption, or the exact functional form of the AR(1) process.

Suppose we have

$$\theta_t = \mathbb{T}(\theta_{t-1}, \nu_t) \quad \text{and} \quad s_{i,t} = \mathbb{S}(\theta_t, \eta_t)$$

and consider more general distributions of innovations ν_t and signal errors η_t. For simplicity, consider the complete graph and $m = 1$. Because θ_{t-1} is still a sufficient condition for the past, an agent's action in period t will still be a function of her subjective distribution over θ_{t-1} and her private signal. An agent with type τ (which is observable) who believes θ_{t-1} is distributed according to $\mathcal{D}$ takes an action equal to $f(\tau, \mathcal{D}, s_{i,t})$. Here, τ could reflect the distribution of agent i's signal, but also perhaps her preferences. We no longer assume that an agent's action is her posterior mean of the random variable: it might be some other statistic, and might be multi-dimensional.

This framework gives an abstract identification condition: agents can learn well if, for any feasible distribution $\mathcal{D}$ over θ_{t-1}, the state θ_t can be inferred from the observed distributions of actions, i.e., distribution of $(\tau, f(\tau, \mathcal{D}, s_{i,t}))$, which each agent would essentially know given enough observations.

Now consider a time-t agent i. Suppose now that any possible distribution that time-$(t-1)$ agents might have over θ_{t-2} can be fully described by a finite tuple of parameters $d \in \mathbb{R}^p$ (e.g., a finite number of moments). For each type τ of $t-1$ agent that i observes, distribution of $f(\tau, d, s_{i,t})$ gives an agent a different measurement of d, a summary of beliefs about θ_{t-2}, and of θ_{t-1}. Assuming there is not too much "collinearity," each of these measurements of the finitely many parameters of interest should, in fact, provide new information about θ_{t-1}. Thus, as long as the set of signal types τ is sufficiently rich, we would expect the identification condition to hold.

References

Acemoglu, D., M. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian Learning in Social Networks," *The Review of Economic Studies*, 78, 1201–1236.

Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016): "Network structure and the aggregation of information: Theory and evidence from Indonesia," *The American Economic Review*, 106, 1663–1704.

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.

DeGroot, M. H. (1974): "Reaching a consensus," *Journal of the American Statistical Association*, 69, 118–121.

DeMarzo, P., D. Vayanos, and J. Zweibel (2003): "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, 118, 909–968.

Duffie, D., S. Malamud, and G. Manso (2009): "Information percolation with equilibrium search dynamics," *Econometrica*, 77, 1513–1574.

Duffie, D. and G. Manso (2007): "Information percolation in large markets," *American Economic Review*, 97, 203–209.

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

——— (2014): "Extensive Imitation is Irrational and Harmful," *Quarterly Journal of Economics*, 129, 1861–1898.

French, J. R. (1956): "A formal theory of social power." *Psychological review*, 63, 181.

Friedkin, N. E. and E. C. Johnsen (1997): "Social positions in influence networks," *Social Networks*, 19, 209–222.

Frongillo, R., G. Schoenebeck, and O. Tamuz (2011): "Social Learning in a Changing World," *Internet and Network Economics*, 146–157.

Gale, D. and S. Kariv (2003): "Bayesian learning in social networks," *Games and Economic Behavior*, 45, 329–346.

Geanakoplos, J. D. and H. M. Polemarchakis (1982): "We can't disagree forever," *Journal of Economic Theory*, 28, 192–200.

Golub, B. and M. Jackson (2010): "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2, 112–149.

Golub, B. and E. Sadler (2016): "Learning in Social Networks," in *The Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, B. Rogers, and B. Rogers, Oxford University Press, chap. 19, 504–542.

Harel, M., E. Mossel, P. Strack, and O. Tamuz (2017): "The Speed of Social Learning," *arXiv preprint arXiv:1412.7172*.

Hayek, F. A. (1945): "The use of knowledge in society," *The American economic review*, 35, 519–530.

Jadbabaie, A., P. Molavi, A. Sandroni, and A. Tahbaz-Salehi (2012): "Non-Bayesian social learning," *Games and Economic Behavior*, 76, 210–225.

Jensen, R. (2007): "The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector," *The quarterly journal of economics*, 122, 879–924.

Kay, S. M. (1993): *Fundamentals of statistical signal processing*, Prentice Hall PTR.

Lehrer, K. and C. Wagner (1981): *Rational consensus in science and society: A philosophical and mathematical study*, vol. 24, Springer Science & Business Media.

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

——— (2015b): "Preferences, homophily, and social learning," *Operations Research*, 64, 564–584.

Manea, M. (2011): "Bargaining in stationary networks," *The American Economic Review*, 101, 2042–2080.

Manresa, E. (2013): "Estimating the structure of social interactions using panel data," *Unpublished Manuscript. CEMFI, Madrid.*

Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2018): "A Theory of Non-Bayesian Social Learning," *Econometrica*, 86, 445–490.

Moscarini, G., M. Ottaviani, and L. Smith (1998): "Social Learning in a Changing World," *Economic Theory*, 11, 657–665.

Rahimian, M. A. and A. Jadbabaie (2017): "Bayesian learning without recall," *IEEE Transactions on Signal and Information Processing over Networks.*

Shahrampour, S., S. Rakhlin, and A. Jadbabaie (2013): "Online learning of dynamic parameters in social networks," *Advances in Neural Information Processing Systems.*

Smith, L. and P. Sørensen (2000): "Pathological outcomes of observational learning," *Econometrica*, 68, 371–398.

Srinivasan, J. and J. Burrell (2013): "Revisiting the fishers of Kerala, India," in *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development: Full Papers-Volume 1*, ACM, 56–66.

Talamàs, E. (2017): "Prices and Efficiency in Networked Markets," *Available at SSRN: https://ssrn.com/abstract=2882377.*

van Oosten, R. (2016): "Learning from Neighbors in a Changing World," *Master's Thesis.*

Weizsäcker, G. (2010): "Do we follow others when we should? A simple test of rational expectations," *The American Economic Review*, 100, 2340–2360.

A Details of definitions (for online publication)

A.1 Exogenous random variables

Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $(\nu_t, \eta_{i,t})_{t\in\mathbb{Z}, i\in N}$ be standard normal, mutually independent random variables. Also take a stochastic process $(\theta_t)_{t\in\mathbb{Z}}$, such that for each $t \in \mathbb{Z}$ and $i \in N$, we have (for $0 < \rho \leq 1$)

$$\theta_t = \rho\theta_{t-1} + \nu_t$$

Such a stochastic process exists by standard constructions of AR(1) process or, in the case of $\rho = 1$, the Gaussian random walk on a doubly infinite time domain. Define $s_{i,t} = \theta_t + \eta_{i,t}$.

A.2 Formal definition of game and stationary linear equilibria

Players and strategies The set of players (or agents) in the game is $\mathcal{A} = \{(i,t) : i \in N, t \in \mathbb{Z}\}$. The set of (pure) *responses* of an agent (i,t) is defined to be the set of all Borel-measurable functions $\sigma_{(i,t)} : \mathbb{R} \times (\mathbb{R}^{|N(i)|})^m \to \mathbb{R}$, mapping her own signal and her neighborhood's actions, $(s_{i,t}, (\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m)$, to a real-valued action $a_{i,t}$. We call the set of these functions $\widetilde{\Sigma}_{(i,t)}$. Let $\widetilde{\Sigma} = \prod_{(i,t)\in\mathcal{A}} \widetilde{\Sigma}_{(i,t)}$ be the set of response profiles. We now define the set of *(unambiguous) strategy profiles*, $\Sigma \subset \widetilde{\Sigma}$. We say that a response profile $\sigma \in \widetilde{\Sigma}$ is a strategy profile if the following two conditions hold

1. There is a tuple of real-valued random variables $(a_{i,t})_{i\in N, t\in\mathbb{Z}}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ such that for each $(i,t) \in \mathcal{A}$, we have

$$a_{i,t} = \sigma_{(i,t)}\left(s_{i,t}, (\boldsymbol{a}_{N_i,t-\ell})_{\ell=1}^m\right).$$

2. Any two tuples tuple of real-valued random variables $(a_{i,t})_{i\in N, t\in\mathbb{Z}}$ satisfying Condition 1 are equal almost surely.

That is, a response profile is a strategy profile if there is an esesntially unique specification of behavior that is consistent with the responses: i.e., if the responses uniquely determine the behavior of the population, and hence payoffs.[21] Note that if $\sigma \in \Sigma$, then

[21]Condition 1 is necessarily to rule out response profiles such as the one given by $\sigma_{i,t}(s_{i,t}, a_{i,t-1}) = |a_{i,t-1}| + 1$. This profile, despite consisting of well-behaved functions, does not correspond to any specification of behavior for the whole population (because time extends infinitely backward). Condition 2 is

$\widetilde{\sigma} = (\sigma'_{(i,t)}, \sigma_{-(i,t)}) \in \Sigma$ whenever $\sigma'_{(i,t)} \in \widetilde{\Sigma}_{(i,t)}$. This is because any Borel-measurable function of a random variable is itself a well-defined random variable. Thus, if we start with a strategy profile and consider agent (i,t)'s deviations, they are unrestricted: she may consider any response.

Payoffs The payoff of an agent (i,t) under any strategy profile $\sigma \in \Sigma$ is

$$u_{i,t}(\sigma) = -\mathbb{E}\left[(a_{i,t} - \theta_t)^2\right] \in [-\infty, 0],$$

where the actions $a_{i,t}$ are taken according to $\sigma_{(i,t)}$ and the expectation is taken in the probability space we have described. This expectation is well-defined because inside the expectation there is a nonnegative, measurable random variable, for which an expectation is always defined, though it may be infinite.

Equilibria A (Nash) *equilibrium* is defined to be a strategy profile $\sigma \in \Sigma$ such that, for each $(i,t) \in \mathcal{A}$ and each $\widetilde{\sigma} \in \Sigma$ such that $\widetilde{\sigma} = (\sigma'_{(i,t)}, \sigma_{-(i,t)})$ for some $\sigma'_{(i,t)} \in \Sigma_{(i,t)}$, we have

$$u_{i,t}(\widetilde{\sigma}) \leq u_{i,t}(\sigma).$$

For $p \in \mathbb{Z}$, we define the shift operator $\mathfrak{T}_p$ to translate variables to time indices shifted p steps forward. This definition may be applied, for example, to Σ.[22] A strategy profile $\sigma \in \Sigma$ is *stationary* if, for all $p \in \mathbb{Z}$, we have $\mathfrak{T}_p \sigma = \sigma$.

We say $\sigma \in \Sigma$ is a *linear* strategy profile if each σ_i is a linear function. Our analysis focuses on *stationary, linear equilibria.*

B Proof of Theorem 1 (for online publication)

B.1 Notation and key notions

Let $\mathbb{S}$ be the (by assumption finite) set of all possible signal variances, and let $\overline{\sigma}^2$ be the largest of them. The proof will focus on the covariances of errors in social signals. Recall

necessary to rule out response profiles such has the one given by $\sigma_{i,t}(s_{i,t}, a_{i,t-1}) = a_{i,t-1}$, which have many satisfying action paths, leaving payoffs undetermined.

[22] I.e., $\sigma' = \mathfrak{T}_p \sigma$ is defined by $\sigma_{(i,t)} = \sigma_{(i,t-p)}$.

that both $r_{i,t}$ and $r_{j,t}$ have mean θ_{t-1}, because each is an unbiased estimate[23] of θ_{t-1}; we will thus focus on the errors $r_{i,t} - \theta_{t-1}$. Let A_t denote the variance-covariance matrix $(\mathrm{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}))_{i,j}$ and let $\mathcal{W}$ be the subset of such covariance matrices. For all i, j note that $\mathrm{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}) \in [-\overline{\sigma}^2, \overline{\sigma}^2]$ using the Cauchy-Schwarz inequality and the fact that $\mathrm{Var}(r_{i,t} - \theta_{t-1}) \in [0, \overline{\sigma}^2]$ for all i. This fact about variances says that no social signal is worse than putting all weight on an agent who follows only her private signal. Thus the best-response map Φ is well-defined and induces a map $\widetilde{\Phi}$ on $\mathcal{W}$.

Next, for any $\delta, \zeta > 0$ we will define the subset $\mathcal{W}_{\delta,\zeta} \subset \mathcal{W}$ of covariance matrices such that both of the following hold:

1. for any pair of distinct agents $i \in G_n^k$ and $j \in G_n^{k'}$,

$$\mathrm{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}) = \delta_{kk'} + \zeta_{ij}$$

   where (i) $\delta_{kk'}$ depends only on the network types of the two agents (k and k', which may be the same); (ii) $|\delta_{kk'}| < \delta$; and (iii) $|\zeta_{ij}| < \zeta$;

2. for any single agent $i \in G_n^k$,

$$\mathrm{Var}(r_{i,t} - \theta_{t-1}) = \delta_k + \zeta_{ii}$$

   where (i) δ_k only depends on the network type of the agent; (ii) $|\delta_k| < \delta$, and (iii) $|\zeta_{ii}| < \zeta$.

This is the space of covariance-matrices such that each covariance is split into two parts. Considering (1) first, $\delta_{kk'}$ is an effect that depends only on the agents' network types, while ζ_{ij} adjusts for the individual-level heterogeneity arising from different link realizations. The description of the decomposition in (2) is analogous.

B.2 Proof strategy

B.2.1 A set $\mathcal{W}_{\overline{\delta},\overline{\zeta}}$ of outcomes with good learning

We will take $\overline{\delta}$ and $\overline{\zeta}$ to be arbitrarily small numbers and show that for large enough n, with high probability (which we will abbreviate "asymptotically almost surely" or "a.a.s.") the

[23]This is because it is a linear combination, with coefficients summing to 1, of unbiased estimates of θ_{t-1}.

equilibrium outcome has a social error covariance matrix A_t in the set $\mathcal{W}_{\bar{\delta},\bar{\zeta}}$. In particular, $\mathrm{Var}(r_{i,t} - \theta_{t-1})$ becomes arbitrarily small in this limit, and that will readily imply that individuals learn very well. In our constructions, the ζ_{ij} (resp., ζ_i) terms will be set to much smaller values than the $\delta_{kk'}$ (resp., δ_k) terms, because group-level covariances are more predictable and less sensitive to idiosyncratic realizations.

B.2.2 Approach to showing that $\mathcal{W}_{\bar{\delta},\bar{\zeta}}$ contains an equilibrium

To show that the equilibrium outcome has (a.a.s.) a social error covariance matrix A_t in the set $\mathcal{W}_{\bar{\delta},\bar{\zeta}}$, the plan is to construct a set so that (a.a.s.) $\overline{\mathcal{W}} \subset \mathcal{W}_{\bar{\delta},\bar{\zeta}}$ and $\widetilde{\Phi}(\overline{\mathcal{W}}) \subset \overline{\mathcal{W}}$. This set will contain an equilibrium by the Brouwer fixed point theorem, and therefore so will $\mathcal{W}_{\bar{\delta},\bar{\zeta}}$.

To construct the set $\overline{\mathcal{W}}$, we will fix a positive constant β (to be determined later), and define

$$\overline{\mathcal{W}} = \mathcal{W}_{\frac{\beta}{n},\frac{1}{n}} \cup \widetilde{\Phi}(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}).$$

We will then prove that, for large enough n, (i) $\widetilde{\Phi}(\overline{\mathcal{W}}) \subseteq \overline{\mathcal{W}}$ and (ii) for another suitable positive constant λ,

$$\overline{\mathcal{W}} \subset \mathcal{W}_{\frac{\beta}{n},\frac{\lambda}{n}}.$$

This will allow us to establish the claims made in the first sentence of the paragraph, with $\bar{\delta}$ and $\bar{\zeta}$ being arbitrarily small numbers.

The following two lemmas will allow us to deduce (as we do immediately after stating them) properties (i) and (ii) of $\overline{\mathcal{W}}$.

Lemma 1. *For all large enough β and all $\lambda \geq \underline{\lambda}(\beta)$, with probability at least $1 - \frac{1}{n}$, we have $\widetilde{\Phi}(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n},\frac{\lambda}{n}}$.*

Lemma 2. *For all large enough β, with probability at least $1 - \frac{1}{n}$, the set $\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}$ is invariant under*[24] *$\widetilde{\Phi}^2$, i.e. $\widetilde{\Phi}^2(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}$.*

Putting these lemmas together, a.a.s. we have,

$$\widetilde{\Phi}^2(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n},\frac{1}{n}} \quad \text{and} \quad \widetilde{\Phi}(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n},\frac{\lambda}{n}}.$$

From this it follows that $\overline{\mathcal{W}} = \mathcal{W}_{\frac{\beta}{n},\frac{1}{n}} \cup \widetilde{\Phi}(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}})$ is invariant under $\widetilde{\Phi}$ and contained in $\mathcal{W}_{\frac{\beta}{n},\frac{\lambda}{n}}$, as claimed.

[24] The notation $\widetilde{\Phi}^2$ means the operator $\widetilde{\Phi}$ applied twice.

B.2.3 Proving the lemmas by analyzing how $\widetilde{\Phi}$ and $\widetilde{\Phi}^2$ act on sets $\mathcal{W}_{\delta,\zeta}$

The lemmas are about how $\widetilde{\Phi}$ and $\widetilde{\Phi}^2$ act on the covariance matrix A_t, assuming it is in a certain set $\mathcal{W}_{\delta,\zeta}$, to yield new covariance matrices. Thus, we will prove these lemmas by studying two periods of updating. The analysis will come in five steps.

Step 1: No-large-deviations (NLD) networks and the high-probability event Step 1 concerns the "with high probability" part of the lemmas. In the entire argument, we condition on the event of a *no-large-deviations (NLD)* network realization, which says that certain realized statistics in the network (e.g., number of paths between two nodes) are close to their expectations. The expectations in question depend only on agents' types. Therefore, on the NLD realization, the realized statistics do not vary much based on which exact agents we focus on, but only on their types. Step 1 defines the NLD event E formally and shows that it has high probability. We use the structure of the NLD event throughout our subsequent steps, as we mention below.

Step 2: Weights in one step of updating are well-behaved We are interested in $\widetilde{\Phi}$ and $\widetilde{\Phi}^2$, which are about how the covariance matrix A_t of social signal errors changes under updating. How this works is determined by the "basic" updating map Φ, and so we begin by studying the weights involved in it and then make deductions about the matrix A_t.

The present step establishes that in one step of updating, the weight $W_{ij,t'}$ that a time-t' agent i (where $t' = t + 1$) places on the action of another agent j in period t, does not depend too much on the identities of i and j. It only depends on their (network and signal) types. This is established by using our explicit formula for weights in terms of covariances. We rely on (i) the fact that covariances are assumed to start out in a suitable $\mathcal{W}_{\delta,\zeta}$, and (ii) our conditioning on the NLD event E. The NLD event is designed so that the network quantities that go into determining the weights depend only on the types of i and j (because the NLD event forbids too much variation conditional on type). The restriction to $A_t \in \mathcal{W}_{\delta,\zeta}$ ensures that covariances in the initial period t did not depend too much on type, either.

Step 3: Lemma 1: $\widetilde{\Phi}(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n},\frac{\lambda}{n}}$ Once we have analyzed one step of updating, it is natural to ask what that does to the covariance matrix. Because we now have a bound on how much weights can vary after one step of updating, we can compute bounds on covariances. This step shows that the initial covariances A_t being in $\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}$ implies that after one step, covariances are in $\mathcal{W}_{\frac{\beta}{n},\frac{\lambda}{n}}$. Note that the introduction of another parameter λ on the right-hand side implies that this step might worsen our control on covariances somewhat, but in a bounded way. This establishes Lemma 1.

Step 4: Weights in two steps of updating are well-behaved The fourth step establishes that the statement made in Step 2 remains true when we replace t' by $t+2$. By the same sort of reasoning as in Step 2, an additional step of updating cannot create too much further idiosyncratic variation in weights. Proving this requires analyzing the covariance matrices of various social signals (i.e., the A_{t+1} that the updating induces), which is why we needed to do Step 3 first.

Step 5: Lemma 2: $\widetilde{\Phi}^2(\mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}$ Now we use our understanding of weights from the previous steps, along with additional structure, to show the key remaining fact. The control on weights we have obtained allows us to control the weight that a given agent's estimate at time $t+2$ places on the social signal of another agent at time t. This is Step 5(a). In the second part, Step 5(b), we use that to control the covariances in A_{t+2}. It is important in this part of the proof that different agents have very similar "second-order neighborhoods": the paths of length 2 beginning from an agent are very similar, in terms of their counts and what types of agents they go through. We carefully separate the variation (across agents) in covariances in A_t into three pieces and use our control of second-order neighborhoods to bound this variation such that $A_{t+2} \in \mathcal{W}_{\frac{\beta}{n},\frac{1}{n}}$.

B.3 Carrying out the steps

B.3.1 Step 1

Here we formally define the NLD event, which we call E. It is given by $E = \cap_{i=1}^{5} E_i$, where the events E_i will be defined next.

(E_1) Let $X^{(1)}_{i,\tau k}$ be the number of agents having signal type τ and network type k who are observed by i. The event E_1 is that this quantity is close to its expected value in the following sense, simultaneously for all possible values of the subscript:

$$(1-\zeta^2)\mathbb{E}[X^{(1)}_{i,\tau k}] \leq X^{(1)}_{i,\tau k} \leq (1+\zeta^2)\mathbb{E}[X^{(1)}_{i,\tau k}].$$

(E_2) Let $X^{(2)}_{ii',\tau k}$ be the number of agents having signal type τ and network type k who are observed by *both* i and i'. The event E_2 is that this quantity is close to its expected value in the following sense, simultaneously for all possible values of the subscript:

$$(1-\zeta^2)\mathbb{E}[X^{(2)}_{ii',\tau k}] \le X^{(2)}_{ii',\tau k} \le (1+\zeta^2)\mathbb{E}[X^{(2)}_{ii',\tau k}].$$

(E_3) Let $X^{(3)}_{i,\tau k,j}$ be the number of agents having signal type τ and network type k who are observed by agent i and who observe agent j. The event E_3 is that this quantity is close to its expected value in the following sense, simultaneously for all possible values of the subscript:

$$(1-\zeta^2)\mathbb{E}[X^{(3)}_{i,\tau k,j}] \le X^{(3)}_{i,\tau k,j} \le (1+\zeta^2)\mathbb{E}[X^{(3)}_{i,\tau k,j}].$$

(E_4) Let $X^{(4)}_{ii',\tau k,j}$ be the number of agents having signal type τ and network type k who are observed by both agent i and i' and who observe j. The event E_4 is that this quantity is close to its expected value in the following sense, simultaneously for all possible values of the subscript:

$$(1-\zeta^2)\mathbb{E}[X^{(4)}_{ii',\tau k',j}] \le X^{(4)}_{ii',\tau k',j} \le (1+\zeta^2)\mathbb{E}[X^{(4)}_{ii',\tau k',j}].$$

(E_5) Let $X^{(5)}_{i,\tau k,jj'}$ be the number of agents of signal type τ and network type k who are observed by agent i and who observe both j and j'. The event E_5 is that this quantity is close to its expected value in the following sense, simultaneously for all possible values of the subscript:

$$(1-\zeta^2)\mathbb{E}[X^{(5)}_{i,\tau k,jj'}] \le X^{(5)}_{i,\tau k,jj'} \le (1+\zeta^2)\mathbb{E}[X^{(5)}_{i,\tau k,jj'}].$$

We claim that the probability of the complement of the event E vanishes exponentially. We can check this by showing that the probability of each of the E_i vanishes exponentially. For E_1, for example, the bounds will hold unless at least one agent has degree outside the specified range. The probability of this is bounded above by the sum of the probabilities of each individual agent having degree outside the specified range. By the central limit theorem, the probability a given agent has degree outside this range vanishes exponentially. Because there are n agents in G_n, this sum vanishes exponentially as well. The other cases are similar.

For the rest of the proof, we condition on the event E.

B.3.2 Step 2

As a shorthand, let $\delta = \beta/n$ for a sufficiently large constant β, and let $\zeta = 1/n$.

Lemma 3. *Suppose that in period t the matrix $A = A_t$ of covariances of social signals satisfies $A \in \mathcal{W}_{\delta,\zeta}$ and all agents are optimizing in period $t+1$. Then there is a γ so that for all n sufficiently large,*

$$\frac{W_{ij,t+1}}{W_{i'j',t+1}} \in \left[1 - \frac{\gamma}{n}, 1 + \frac{\gamma}{n}\right].$$

whenever i and i' have the same network and signal types and j and j' have the same network and signal types.

To prove this lemma, we will use our weights formula:

$$W_{i,t+1} = \frac{\mathbf{1}^T \boldsymbol{C}_{i,t}^{-1}}{\mathbf{1}^T \boldsymbol{C}_{i,t}^{-1} \mathbf{1}}.$$

This says that in period $t+1$, agent i's weight on agent j is proportional to the sum of the entries of column j of $\boldsymbol{C}_{i,t}^{-1}$. We want to show that the change in weights is small as the covariances of observed social signals vary slightly. To do so we will use the Taylor expansion in an arbitrary direction of $f(A) = \boldsymbol{C}_{i,t}^{-1}$ around the covariance matrix $A(0)$ at which all $\delta_{kk'} = 0$, $\delta_k = 0$ and $\zeta_{ij} = 0$.

We begin with the first partial derivative of f at $A(0)$ in an arbitrary direction. Let $A(x)$ be any perturbation of A_0 in one parameter, i.e. $A(x) = A(0) + xM$ for some constant matrix M with entries in $[-1, 1]$. Let $\boldsymbol{C}_i(x)$ be the matrix of covariances of the actions observed by i given that the covariances of agents' social signals were $A(x)$. There exists a constant γ_1 depending only on the possible signal types such that each entry of $\boldsymbol{C}_i(x) - \boldsymbol{C}_i(x')$ has absolute value at most $\gamma_1(x - x')$ whenever both x and x' are small.

We will now show that the column sums of $\boldsymbol{C}_i(x)^{-1}$ are close to the column sums of $\boldsymbol{C}(0)_i^{-1}$. To do so, we will evaluate the formula

$$\frac{\partial \boldsymbol{C}_i(x)^{-1}}{\partial x} = \boldsymbol{C}_i(x)^{-1} \frac{\partial \boldsymbol{C}_i(x)}{\partial x} \boldsymbol{C}_i(x)^{-1} \tag{9}$$

at zero. If we can bound each column sum of this expression (evaluated at zero) by a constant (depending only on the signal types and the number of network types K), then the first derivative of f will also be bounded by a constant.

Recall that $\mathbb{S}$ is the set of signal types and let $S = |\mathbb{S}|$; index the signal types by numbers ranging from 1 to S. To bound the column sums of $\boldsymbol{C}_i(0)^{-1}$, suppose that the agent observes r_i agents from each signal type $1 \leq i \leq S$. Reordering so that all agents of each signal type are grouped together, we can write

$$\boldsymbol{C}_i(0) = \begin{pmatrix} a_{11}\mathbf{1}_{r_1\times r_1} + b_1 I_{r_1} & a_{12}\mathbf{1}_{r_1\times r_2} & & a_{S1}\mathbf{1}_{r_1\times r_S} \\ a_{12}\mathbf{1}_{r_2\times r_1} & a_{22}\mathbf{1}_{r_2\times r_2} + b_2 I_{r_2} & & \vdots \\ & & \ddots & \\ a_{1S}\mathbf{1}_{r_S\times r_1} & \cdots & & a_{SS}\mathbf{1}_{r_s\times r_s} + b_s I_{r_s} \end{pmatrix}$$

Therefore, $\boldsymbol{C}_i(0)$ can be written as a block matrix with blocks $a_{ij}1_{r_i\times r_j} + b_i\delta_{ij}I_{r_i}$ where $1 \leq i,j \leq S$ and $\delta_{ij} = 1$ for $i = j$ and 0 otherwise.

We now have the following important approximation of the inverse of this matrix.[25]

Lemma 4. *Let C be a matrix consisting of $S \times S$ blocks, with its (i,j) block given by*

$$a_{ij}1_{r_i\times r_j} + b_i\delta_{ij}I_{r_i}$$

and let $A = a_{ij}1_{r_i\times r_j}$ be an invertible matrix. As $n \to \infty$, then the (i,i) block of C^{-1} is equal to

$$\frac{1}{b_i}I_{r_i} - \frac{1}{b_i r_i}1_{r_i\times r_i} + O(1/n^2)$$

while the off-diagonal blocks are $O(1/n^2)$.

Proof. First note that the ij-block of C^{-1} has the form

$$c_{ij}1_{r_i\times r_j} + d_i\delta_{ij}I_{r_i}$$

for some real c_{ij} and d_i.

Therefore, CC^{-1} can be written in matrix form as

$$\begin{aligned} &\textstyle\sum_k (a_{ik}1_{r_i\times r_k} + b_i\delta_{ik}I_{r_i})(c_{kj}1_{r_k\times r_j} + d_k\delta_{kj}I_{r_k}) = \\ &a_{ij}d_j + \textstyle\sum_k (a_{ik}r_k + \delta_{ik}b_k)c_{kj}1_{r_i\times r_j} + b_i d_i\delta_{ij}I_{r_i}. \end{aligned} \tag{10}$$

Note that the last summand is the identity matrix.

Let D_d denote the diagonal matrix with d_i in the (i,i) diagonal entry, let $D_{1/b}$ denote the diagonal matrix with $1/b_i$ in the (i,i) diagonal entry, etc. Breaking up the previous

[25]We are very grateful to Iosif Pinelis for the proof of this lemma **?**.

display (10) into its diagonal and off-diagonal parts, we can write

$$AD_d + (AD_r + D_b)C = 0 \text{ and } D_d = D_{1/b}.$$

Hence,

$$\begin{aligned} C &= -(AD_r + D_b)^{-1}AD_d \\ &= -(I_q + D_r^{-1}A^{-1}D_b)^{-1}(AD_r)^{-1}AD_{1/b} \\ &= -(I_q + D_r^{-1}A^{-1}D_b)^{-1}D_{1/(br)} \\ &= -D_{1/(br)} + O(1/n^2) \end{aligned}$$

where $br := (b_1r_1, \ldots, b_qr_q)$. Therefore as $n \to \infty$ the off-diagonal blocks will be $O(1/n^2)$ while the diagonal blocks are

$$\frac{1}{b_i}I_{r_i} - \frac{1}{b_ir_i}1_{r_i\times r_i} + O(1/n^2)$$

as desired. $\square$

Using Lemma 4 we can compute the column sums of

$$\boldsymbol{C}_i(0)^{-1}M\boldsymbol{C}_i(0)^{-1}$$

and find that all terms containing only constant or order $O(1/n)$ terms from both matrices $\boldsymbol{C}_i(0)^{-1}$ cancel. So the column sums are $O(1/n)$ whenever all entries of M are in $[-1,1]$.

We can bound the higher-order terms in the Taylor expansion by the same technique: by differentiating equation 9 repeatedly in x, we obtain an expression for the k^{th} derivative in terms of $\boldsymbol{C}_i(0)^{-1}$ and M:

$$f^{(k)}(0) = k!\boldsymbol{C}_i(0)^{-1}M\boldsymbol{C}_i(0)^{-1}M\boldsymbol{C}_i(0)^{-1}\cdot\ldots\cdot M\boldsymbol{C}_i(0)^{-1},$$

where M appears k times in the product. By the same argument as above, we can show that the column sums of $\frac{f^{(k)}(0)}{k!}$ are bounded by a constant independent of n. The Taylor expansion is

$$f(A) = \sum_k \frac{f^{(k)}(0)}{k!}x^k.$$

Since we take $A \in \mathcal{W}_{\delta,\zeta}$, we can assume that x is $O(1/n)$. Because the column sums of each summand are bounded by a constant times x^k, the column sums of $f(A)$ are bounded by a constant.

Finally, because the variation in the column sums is $O(1/n)$ and the weights are proportional to the column sums, each weight varies by at most a multiplicative factor of γ_1/n for some γ_1. We find that the first part of the lemma, which bounded the ratios between weights $W_{ij,t+1}/W_{i'j',t+1}$, holds.

B.3.3 Step 3

We complete the proof of Lemma 1, which states that the covariance matrix of $r_{i,t+1}$ is in $\mathcal{W}_{\delta,\zeta'}$. Recall that $\zeta' = \lambda/n$ for some constant n, so we are showing that if the covariance matrix of the $r_{i,t}$ is in a neighborhood $\mathcal{W}_{\delta,\zeta}$, then the covariance matrix in the next period is in a somewhat larger neighborhood $\mathcal{W}_{\delta,\zeta'}$. The remainder of the argument then follows by the same arguments as in the proof of the first part of the lemma: we now bound the change in time-$(t+2)$ weights as we vary the covariances of time-$(t+1)$ social signals within this neighborhood.

Recall that we decomposed each covariance $\mathrm{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}) = \delta_{kk'} + \zeta_{ij}$ into a term $\delta_{kk'}$ depending only on the types of the two agents and a term ζ_{ij}, and similarly for variances. To show the covariance matrix is contained in $\mathcal{W}_{\delta,\zeta'}$, we bound each of these terms suitably.

We begin with ζ_{ij} (and ζ_i). We can write

$$r_{i,t+1} = \sum_j \frac{W_{ij,t+1}}{1 - w^s_{i,t+1}} a_{i,t} = \sum_j \frac{W_{ij,t+1}}{1 - w^s_{i,t+1}} (w^s_{j,t} s_{j,t} + (1 - w^s_{j,t}) r_{j,t}).$$

By the first part of the lemma, the ratio between any two weights (both of the form $W_{ij,t+1}$, $w^s_{i,t+1}$, or $w^s_{j,t}$) corresponding to pairs of agents of the same types is in $[1 - \gamma_1/n, 1 + \gamma_1/n]$ for a constant γ_1. We can use this to bound the variation in covariances of $r_{i,t+1}$ within types by ζ': we take the covariance of $r_{i,t+1}$ and $r_{j,t+1}$ using the expansion above and then bound the resulting summation by bounding all coefficients.

Next we bound $\delta_{kk'}$ (and δ_k). It is sufficient to show that $\mathrm{Var}(r_{i,t+1} - \theta_t)$ is at most δ. To do so, we will give an estimator of θ_t with variance less than β/n, and this will imply $\mathrm{Var}(r_{i,t+1} - \theta_t) < \beta/n = \delta$ (recall $r_{i,t+1}$ is the estimate of θ_t given agent i's social observations in period $t+1$). Since this bounds all the variance terms by δ, the covariance terms will also be bounded by δ in absolute value.

Fix an agent i of network type k and consider some network type k' such that $p_{kk'} > 0$. Then there exists two signal types, which we call A and B, such that i observes $\Omega(n)$ agents of each of these signal types in G_n^k.[26] The basic idea will be that we can approximate θ_t well by taking a linear combination of the average of observed agents of network type k and signal type A and the average of observed agents of network type k and signal type B.

In more detail: Let $N_{i,A}$ be the set of agents of type A in network type k observed by i and $N_{i,B}$ be the set of agents of type B in network type k observed by i. Then fixing some agent j_0 of network type k,

$$\frac{1}{|N_{i,A}|}\sum_{j\in N_{i,A}} a_{j,t-1} = \frac{\sigma_A^{-2}}{1+\sigma_A^{-2}}\theta_t + \frac{1}{1+\sigma_A^{-2}} r_{j_0,t-1} + noise$$

where the noise term has variance of order $1/n$ and depends on signal noise, variation in $r_{j,t}$, and variation in weights. Similarly

$$\frac{1}{|N_{i,B}|}\sum_{j\in N_{i,B}} a_{j,t} = \frac{\sigma_B^{-2}}{1+\sigma_B^{-2}}\theta_t + \frac{1}{1+\sigma_B^{-2}} r_{j_0,t-1} + noise$$

where the noise term has the same properties. Because $\sigma_A^2 \neq \sigma_B^2$, we can write θ_t as a linear combination of these two averages with coefficients independent of n up to a noise term of order $1/n$. We can choose β large enough such that this noise term has variance most β/n for all n sufficiently large. This completes the Proof of Lemma 1.

B.3.4 Step 4:

We now give the two-step version of Lemma 3.

Lemma 5. *Suppose that in period t the matrix $A = A_t$ of covariances of social signals satisfies $A \in \mathcal{W}_{\delta,\zeta}$ and all agents are optimizing in periods $t+1$ and $t+2$. Then there is a γ so that for all n sufficiently large,*

$$\frac{W_{ij,t+2}}{W_{i'j',t+2}} \in \left[1-\frac{\gamma}{n}, 1+\frac{\gamma}{n}\right].$$

whenever i and i' have the same network and signal types and j and j' have the same network and signal types.

[26] We use the notation $\Omega(n)$ to mean greater than Cn for some constant $C > 0$ when n is large.

Given what we established about covariances in Step 3, the lemma follows by the same argument as the proof of Lemma 3. .

Step 5: Now that Lemma 5 is proved, we can apply it to show that

$$\widetilde{\Phi}^2(\mathcal{W}_{\delta,\zeta}) \subset \mathcal{W}_{\delta,\zeta}.$$

We will do this by first writing the time-$(t+2)$ behavior in terms of agents' time-t observations (Step 5(a)), which comes from applying $\widetilde{\Phi}$ twice. This gives a formula that can be used for bounding the covariances[27] of time-$(t+2)$ actions in terms of covariances of time-t actions. Step 5(b) then applies this formula to show we can take ζ_{ij} and ζ_i to be sufficiently small. We split our expression for $r_{i,t+2}$ into several groups of terms and show that the contribution of each group of terms depends only on agents' types up to a small noise term. Step 5(c) notes that we can also take $\delta_{kk'}$ and δ_k to be sufficiently small.

Step 5(a):

$$\begin{aligned}
r_{i,t+2} &= \sum_j \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} \rho a_{j,t+1} \\
&= \rho\Big(\sum_j \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} w^s_{j,t+1} s_{j,t+1} + \sum_{j,j'} \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} W_{jj',t+1} \rho a_{j',t}\Big) \\
&= \rho\Big(\sum_j \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} w^s_{j,t+1} s_{j,t+1} + \rho\Big(\sum_{j,j'} \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} W_{jj',t+1} w^s_{j',t} s_{j',t} \\
&\quad + \sum_{j,j'} \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} W_{jj',t+1} (1-w^s_{j',t}) r_{j',t}\Big)\Big).
\end{aligned}$$

Let $c_{ij',t}$ be the coefficient on $r_{j',t}$ in this expansion of $r_{i,t+2}$. Explicitly,

$$c_{ij',t} = \sum_j \frac{W_{ij,t+2}}{1-w^s_{i,t+2}} W_{jj',t+1} (1-w^s_{j',t}).$$

The coefficient $c_{ij',t}$ adds up the influence of $r_{j',t}$ on $r_{i,t+2}$ over all paths of length two.

First, we establish a lemma about how much these weights vary.

Lemma 6. *For n sufficiently large, when i and i' have the same network types and j' and j'' have the same network and signal types, the ratio $c_{ij',t}/c_{i'j'',t}$ is in $[1-2\gamma/n, 1+2\gamma/n]$.*

[27] We take this term to include variances.

Proof. Suppose $i \in G_k$ and $j' \in G_{k'}$. For each network type k'', the number of agents j of type k'' who are observed by i and who observe j' varies by at most a factor ζ^2 as we change i in G_k and j' in $G_{k'}$. For each such j, the contribution of that agent's action to $c_{ij',t}$ is

$$\frac{W_{ij,t+2}}{1-w^s_{i,t+2}} W_{jj',t+1}(1-w^s_{j',t}).$$

By Lemma 3 applied to each term, this expression varies by at most a factor of γ/n as we change i in G_k and j' in $G_{k'}$. Combining these facts for each type k'' shows the lemma. $\square$

Step 5(b): We first show that fixing the values of $\delta_{kk'}$ and δ_k in period t, the variation in the covariances $\mathrm{Cov}(r_{i,t+2} - \theta_{t+1}, r_{i',t+2} - \theta_{t+1})$ of these terms as we vary i and i' over network types is not larger than ζ. From the formula above, we observe that we can decompose $r_{i,t+2} - \theta_{t+1}$ as a linear combination of three mutually independent groups of terms:

(i) signal error terms $\eta_{j,t+1}$ and $\eta_{j',t}$;

(ii) the errors $r_{j',t} - \theta_t$ in the social signals from period t; and

(iii) changes in state $\nu_t$and ν_{t+1} between periods t and $t+2$.

Note that the terms $r_{j',t} - \theta_t$ are linear combinations of older signal errors and changes in the state. We bound each of the three groups in turn:

(i) Signal Errors: We first consider the contribution of signal errors. When i and i' are distinct, the number of such terms is close to its expected value because E_2 and E_4 hold. Moreover the weights are close to their expected values by Step 2, so the variation is bounded suitably. When i and i' are equal, we use the facts that the weights are close to their expected values and the variance of an average of $\Omega(n)$ signals is small.

(ii) Social Signals: We now consider terms $r_{j',t} - \theta_t$, which correspond to the third summand in our expression for $r_{i,t+2}$. Since we will analyze the weight on ν_t below, it is sufficient to study the terms $r_{j',t} - \theta_{t-1}$.

By Lemma 6, the coefficients placed on $r_{j',t}$ by i and on $r_{j'',t}$ by i' vary by a factor of at most $2\gamma/n$. Moreover, the absolute value of each of these covariances is bounded above by δ and the variation in these terms is bounded above by ζ. We conclude that the variation from these terms has order $1/n^2$.

(iii) Innovations: Finally, we consider the contribution of the innovations $\nu_t$and ν_{t+1}. We treat ν_{t+1} first. We must show that any two agents of the same types place the same weight on the innovation ν_{t+1} (up to an error of order $\frac{1}{n^2}$). This will imply that the contributions of timing to the covariances $\mathrm{Cov}(r_{i,t+2} - \theta_{t+1}, r_{i',t+2} - \theta_{t+1})$ can be expressed as a term that can be included in the relevant $\delta_{kk'}$ and a lower-order term which can be included in $\zeta_{ii'}$.

The weight an agent places on ν_{t+1} is equal to the weight she places on signals from period $t+1$. So this is equivalent to showing that the total weight

$$\rho \sum_j \frac{W_{ij,t+2}}{1 - w^s_{i,t+2}} w^s_{j,t+1}$$

agent i places on period $t+1$ depends only on the network type k of agent i and $O(1/n^2)$ terms. We will first show the average weight placed on time-$(t+1)$ signals by agents of each signal type depends only on k. We will then show that the total weights on agents of each signal type do not depend on n.

Suppose for simplicity here that there are two signal types A and B; the general case is the same. We can split the sum from the previous paragraph into the subgroups of agents with signal types A and B:

$$\rho \sum_{j:\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{1 - w^s_{i,t+2}} w^s_{j,t+1} + \rho \sum_{j:\sigma_j^2=\sigma_B^2} \frac{W_{ij,t+2}}{1 - w^s_{i,t+2}} w^s_{j,t+1}.$$

Letting $W_i^A = \sum_{\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{1-w^s_{i,t+2}}$ be the total weight placed on agents with signal type A and similarly for signal type B, we can rewrite this as:

$$W_i^A \rho \sum_{j:\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{W_i^A(1 - w^s_{i,t+2})} w^s_{j,t+1} + W_i^B \rho \sum_{j:\sigma_j^2=\sigma_B^2} \frac{W_{ij,t+2}}{W_i^B(1 - w^s_{i,t+2})} w^s_{j,t+1}.$$

The coefficients $\frac{W_{ij,t+2}}{W_i^A(1-w^s_{i,t+2})}$ in the first sum now sum to one, and similarly for the second. We want to check that the first sum $\sum_{j:\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{W_i^A(1-w^s_{i,t+2})} w^s_{j,t+1}$ does not depend on k, and the second sum is similar.

For each j in group A,

$$w^s_{j,t+1} = \frac{\sigma_A^{-2}}{\sigma_A^{-2} + (\rho^2 \kappa_{j,t+1} + 1)^{-1}},$$

where we recall that $\kappa^2_{j,t+1} = \mathrm{Var}(r_{j,t+1} - \theta_t)$. Because $\kappa_{j,t+1}$ is close to zero, we can approximate $w^s_{j,t+1}$ locally as a linear function $\mu_1 \kappa_{j,t+1} + \mu_2$ where $\mu_1 < 1$ (up to order $\frac{1}{n^2}$ terms).

So we can write the sum of interest as

$$\sum_{j:\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{W_i^A(1-w_{i,t+2}^s)}(\mu_1 \sum_{j',j''} W_{jj',t+1}W_{jj'',t+1}(\rho^2 \boldsymbol{V}_{j'j'',t}+1)+\mu_2).$$

By Lemma 3, the weights vary by at most a multiplicative factor contained in $[1-\gamma/n, 1+\gamma/n]$. The number of paths from i to j' passing through agents of any network type k'' and any signal type is close to its expected value (which depends only on i's network type), and the weight on each path depends only on the types involved up to a factor in $[1-\gamma/n, 1+\gamma/n]$. The variation in $\boldsymbol{V}_{j'j'',t}$ consists of terms of the form $\delta_{k'k''}$, $\delta_{k'}$, and $\zeta_{j'j''}$, all of which are $O(1/n)$, and terms from signal errors $\eta_{j',t}$. The signal errors only contribute only when $j=j'$, and so only contribute to a fraction of the summands of order $1/n$. So we can conclude the total variation in this sum as we change i within the network type k has order $1/n^2$.

Now that we know each the average weight on private signals of the observed agents of each signal type depends only on k, it remains to check that W_i^A and W_i^B only depend on k. The coefficients W_i^A and W_i^B are the optimal weights on the group averages

$$\sum_{j:\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{W_i^A(1-w_{i,t+2}^s)}\rho a_{j,t+1} \text{ and } \sum_{j:\sigma_j^2=\sigma_B^2} \frac{W_{ij,t+2}}{W_i^B(1-w_{i,t+2}^s)}\rho a_{j,t+1},$$

so we need to show that the variances and covariance of these two terms depend only on k. We check the variance of the first sum: we can expand

$$\sum_{\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{W_i^A(1-w_{i,t+2}^s)}\rho a_{j,t+1} = \sum_{\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{W_i^A(1-w_{i,t+2}^s)}\rho(w_{j,t+1}^s s_{j,t+1}+(1-w_{j,t+1}^s)r_{j,t+1}).$$

We can again bound the signal errors and social signals as in the previous parts of this proof, and show that the variance of this term depends only on k and $O(1/n^2)$ terms. The second variance and covariance are similar, so W_i^A and W_i^B depend only on k and $O(1/n^2)$ terms.

This takes care of the innovation ν_{t+1}. Because we have included any innovations prior to ν_t in the social signals $r_{j',t}$, to complete Step 5(b) we need only show the weight on ν_t depends only on the network type k of an agent.

The analysis is a simpler version of the analysis of the weight on ν_{t+1}. It is sufficient to show the total weight placed on period t social signals depends only on the network type of k of an agent i. This weight is equal to

$$\rho^2 \sum_{j,j'} \frac{W_{ij,t+2}}{1 - w^s_{i,t+2}} \cdot W_{jj',t+1} \cdot (1 - w^s_{j',t}).$$

As in the ν_{t+1} case, we can approximate $(1 - w^s_{j',t})$ as a linear function of $\kappa_{j',t}$ up to $O(1/n^2)$ terms. Because the number of paths to each agent j' though a given type and the weights on each such path cannot vary too much within types, the same argument shows that this sum depends only on k and $O(1/n^2)$ terms.

Step 5(b) is complete.

Step 5(c): The final step is to verify that we can take $\delta_{kk'}$ and δ_k to be smaller than δ. It is sufficient to show that the variance $\mathrm{Var}(r_{i,t+2} - \theta_{t+1})$ of each social signal about θ_{t+1} is at most δ. The proof is the same as in Step 2(b).

C Remaining proofs (for online publication)

C.1 Proof of Proposition 1

Recall from Section 3.3 the map Φ, which gives the next-period covariance matrix $\Phi(\boldsymbol{V}_t)$ for any $\boldsymbol{V}_t$. The expression given there for this map ensures that its entries are continuous functions of the entries of $\boldsymbol{V}_t$. Our strategy is to show that this function maps a compact set to itself, which, by Brouwer's fixed-point theorem, ensures that Φ has a fixed point $\widehat{\boldsymbol{V}}$.

We now define the compact set. Because memory is arbitrary, entries of $\boldsymbol{V}_t$ are covariances between pairs of observations from any periods available in memory. Let k, l be two indices of such observations, corresponding to actions taken by agents i and j respectively, and let $\bar{\sigma}_i = \max\{\sigma_i^2, \rho^{m-1}\sigma_i^2 + \frac{1-\rho^{m-1}}{1-\rho}\}$. Now let $\mathcal{K} \subset \mathcal{V}$ be the subset of symmetric positive semi-definite matrices $\boldsymbol{V}_t$ such that, for any such k, l,

$$V_{kk,t} \in \left[\min\left\{\frac{1}{1+\sigma_i^{-2}}, \frac{\rho^{m-1}}{1+\sigma_i^{-2}} + \frac{1-\rho^{m-1}}{1-\rho}\right\}, \max\left\{\sigma_i^2, \rho^{m-1}\sigma_i^2 + \frac{1-\rho^{m-1}}{1-\rho}\right\}\right]$$
$$\boldsymbol{V}_{kl,t} \in [-\bar{\sigma}_i\bar{\sigma}_j, \bar{\sigma}_i\bar{\sigma}_j].$$

This set is closed and convex, and we claim that $\Phi(\mathcal{K}) \subset \mathcal{K}$.

To show this claim, we will first bound the variance of any agent's action (at any period in memory). Note that a Bayesian agent will not choose an action with a larger variance than her signal, which has variance σ_i^2.

For a lower bound on the variance of the agent's action, note that if she knew the previous period's state and her own signal, then the variance of her action would be $\frac{1}{1+\sigma_i^{-2}}$. Thus an agent observing only noisy estimates of θ_t and her own signal can do no better. By the same reasoning applied to i's self from m periods ago, the variance of the estimate of θ_{t-1} based on i's action from m periods ago cannot exceed $\rho^{m-1}\sigma_i^2 + \frac{1-\rho^{m-1}}{1-\rho}$ or be smaller than $\frac{\rho^{m-1}}{1+\sigma_i^{-2}} + \frac{1-\rho^{m-1}}{1-\rho}$. This establishes bounds on $V_{kk,t}$ for observations k coming from either the most recent or the oldest available period. The corresponding bounds from the periods between $t-m+1$ and t are always weaker than at least one of the two bounds we have described, so we need only take minima and maxima over two terms.

This established the claimed bound on the variances. The bounds on covariances follow from Cauchy-Schwartz.

C.2 Proof of Proposition 2

We first check there is a unique equilibrium and then prove the remainder of Proposition 2.

Lemma 7. *Suppose G has symmetric neighbors. Then there is a unique equilibrium.*

Proof of Lemma. We will show that when the network satisfies the condition in the proposition statement, Φ induces a contraction on a suitable space. For each agent, we can consider the variance of the best estimator for yesterday's state based on observed actions. These variances are tractable because they satisfy the envelope theorem. Moreover, the space of these variances is a sufficient statistic for determining all agent strategies and action variances.

Let $r_{i,t}$ be i's *social signal*—the best estimator of θ_t based on the period $t-1$ actions of agents in N_i—and let $\kappa_{i,t}^2$ be the variance of $r_{i,t} - \theta_t$.

We claim that Φ induces a map $\widetilde{\Phi}$ on the space of variances $\kappa_{i,t}^2$, which we denote $\widetilde{\mathcal{V}}$. We must check the period t variances $(\kappa_{i,t}^2)_i$ uniquely determine all period $t+1$ variances $(\kappa_{i,t+1}^2)_i$: The variance $\boldsymbol{V}_{ii,t}$ of agent i's action, as well as the covariances $\boldsymbol{V}_{ii',t}$ of all pairs of agents i, i' with $N_i = N_{i'}$, are determined by $\kappa_{i,t}^2$. Moreover, by the condition on our network, these variances and covariances determine all agents' strategies in period $t+1$, and this is enough to pin down all period $t+1$ variances $\kappa_{i,t+1}^2$.

The proof proceeds by showing $\widetilde{\Phi}$ is a contraction on $\widetilde{\mathcal{V}}$ in the sup norm.

For each agent j, we have $N_i = N_{i'}$ for all $i, i' \in N_j$. So the period t actions of an agent i' in N_j are

$$a_{i',t} = \frac{(\rho^2\kappa_{i,t}^2 + 1)^{-1}}{\sigma_{i'}^{-2} + (\rho^2\kappa_{i,t}^2 + 1)^{-1}} \cdot r_{i,t} + \frac{\sigma_{i'}^{-2}}{\sigma_{i'}^{-2} + (\rho^2\kappa_{i,t}^2 + 1)^{-1}} \cdot s_{i',t} \tag{11}$$

where $s_{i',t}$ is agent (i')'s signal in period t and $r_{i,t}$ the social signal of i (the same one that i' has). It follows from this formula that each action observed by j is a linear combination of a private signal and a *common* estimator $r_{i,t}$, with positive coefficients which sum to one. For simplicity we write

$$a_{i',t} = b_0 \cdot r_{i,t} + b_{i'} \cdot s_{i',t} \tag{12}$$

(where b_0 and $b_{i'}$ depend on i' and t, but we omit these subscripts). We will use the facts $0 < b_0 < 1$ and $0 < b_{i'} < 1$.

We are interested in how $\kappa_{j,t}^2 = \mathrm{Var}(r_{j,t} - \theta_t)$ depends on $\kappa_{i,t-1}^2 = \mathrm{Var}(r_{i,t-1} - \theta_{t-1})$. The estimator $r_{j,t}$ is a linear combination of observed actions $a_{i',t}$, and therefore can be expanded as a linear combination of signals $s_{i',t}$ and the estimator $r_{i,t-1}$. We can write

$$r_{j,t} = c_0 \cdot (\rho r_{i,t-1}) + \sum_{i'} c_{i'} s_{i',t} \tag{13}$$

and therefore (taking variances of both sides)

$$\begin{aligned}
\kappa_{j,t}^2 = \mathrm{Var}(r_{j,t} - \theta_t) &= c_0 \mathrm{Var}(\rho r_{i,t-1} - \theta_t) + \sum_{i'} c_{i'} \sigma_{i'}^2 \\
&= c_0(\kappa_{i,t-1}^2 + 1) + \sum_{i'} c_{i'} \sigma_{i'}^2
\end{aligned}$$

The desired result, that $\widetilde{\Phi}$ is a contraction, will follow if we can show that the derivative $\frac{d\kappa_{j,t}^2}{d\kappa_{i,t-1}^2} \in [0, \delta]$ for some $\delta < 1$. By the envelope theorem, when calculating this derivative, we can assume that the weights placed on actions $a_{i',t-1}$ by the estimator $r_{j,t}$ do not change as we vary $\kappa_{i,t-1}^2$, and therefore c_0 and the $c_{i'}$ above do not change. So it is enough to show the coefficient c_0 on $\kappa_{i,t-1}^2$ is in $[0, \delta]$. $\square$

The intuition for the lower bound is that *anti-imitation* (agents placing negative weights on observed actions) only occurs if observed actions put too much weight on public information. But if $c_0 < 0$, then the weight on public information is actually negative so there is no reason to anti-imitate. This is formalized in the following lemma.

Lemma 8. *Agent j's social signal places non-negative weight on agent i's social signal from the previous period, i.e., $c_0 \geq 0$.*

Proof. To check this formally, suppose that c_0 is negative. Then the social signal $r_{j,t}$ puts negative weight on some observed action—say the action $a_{k,t-1}$ of agent k. We want to check that the covariance of $r_{j,t} - \theta_t$ and $a_{k,t-1} - \theta_t$ is negative. Using (12) and (13), we compute that

$$\begin{aligned}\mathrm{Cov}(r_{j,t} - \theta_t, a_{k,t-1} - \theta_t) &= \mathrm{Cov}\left(c_0(\rho r_{i,t-1} - \theta_t) + \sum_{i' \in N_j} c_{i'}(s_{i',t} - \theta_t)), b_0(\rho r_{i,t-1} - \theta_t) + b_k(s_{k,t-1} - \theta_t)\right) \\ &= c_0 b_0 \mathrm{Var}(\rho r_{i,t-1} - \theta_t) + c_k b_k \mathrm{Var}(s_{k,t-1} - \theta_t)\end{aligned}$$

because all distinct summands above are mutually independent. We have $b_0, b_k > 0$, while $c_0 < 0$ by assumption and $c_k < 0$ because the estimator $r_{j,t}$ puts negative weight on $a_{k,t-1}$. So the expression above is negative. Therefore, it follows from the usual Gaussian Bayesian updating formula that the best estimator of θ_t given $r_{j,t}$ and $a_{k,t-1}$ puts positive weight on $a_{k,t-1}$. However, this is a contradiction: the best estimator of θ_t given $r_{j,t}$ and $a_{k,t-1}$ is simply $r_{j,t}$, because $r_{j,t}$ was defined as the best estimator of θ_t given observations that included $a_{k,t-1}$.

Now, for the upper bound $c_0 \leq \delta$, the idea is that $r_{j,t}$ puts more weight on agents with better signals while these agents put little weight on public information, which keeps the overall weight on public information from growing too large.

Note that $r_{j,t}$ is a linear combination of actions $\rho a_{i',t-1}$ for $i \in N_j$, with coefficients summing to 1. The only way the coefficient on $\rho r_{i,t-1}$ in $r_{j,t}$ could be at least 1 would be if some of these coefficients on $\rho a_{i',t-1}$ were negative and the estimator $r_{j,t}$ placed greater weight on actions $a_{i',t-1}$ which placed more weight on $r_{j,t}$.

Applying the formula (11) for $a_{i',t-1}$, we see that the coefficient b_0 on $\rho r_{i,t-1}$ is less than 1 and increasing in $\sigma_{i'}$. On the other hand, it is clear that the weight on $a_{i',t-1}$ in the social signal $r_{j,t}$ is decreasing in $\sigma_{i'}$: more weight should be put on more precise individuals. So in fact the estimator $r_{j,t}$ places less weight on actions $a_{i',t-1}$ which placed more weight on $r_{i,t}$.

Moreover, the coefficients placed on private signals are bounded below by a positive constant when we restrict to covariances in the image of $\widetilde{\Phi}$ (because all covariances are bounded as in the proof of Proposition 1). Therefore, each agent $i' \in N_j$ places weight at most δ on the estimator $\rho r_{i,t-1}$ for some $\delta < 1$. Agent j's social signal $r_{j,t}$ is a sum of these agents' actions with coefficients summing to 1 and satisfying the monotonicity property above. We conclude that the coefficient on $\rho r_{i,t-1}$ in the expression for $r_{j,t}$ is at most δ. We conclude that the coefficient on $\rho r_{i,t-1}$ in $r_{j,t}$ is bounded above by some $\delta < 1$. $\square$

This completes the proof of Lemma 7. We now prove Proposition 2.

Proof of Proposition **??**. By Lemma 7 that there is a unique equilibrium on any network G with symmetric neighbors. Let $\varepsilon > 0$.

Consider any agent i. Her neighbors have the same private signal qualities and the same neighborhoods (by the symmetric neighbors assumption). So there exists an equilibrium where for all i, the actions of agent i's neighbors are exchangeable. By uniqueness, this in fact holds at the sole equilibrium.

So agent i's social signal is an average of her neighbors' actions:

$$r_{i,t} = \frac{1}{|N_i|} \sum_{j \in N_i} a_{j,t}.$$

Suppose the ε-perfect aggregation benchmark is achieved. Then all agents must place weight at least $\frac{(1+\varepsilon)^{-1}}{(1+\varepsilon)^{-1}+\sigma^{-2}}$ on their social signals. So at time t, the social signal $r_{i,t}$ places weight at least $\frac{(1+\varepsilon)^{-1}}{(1+\varepsilon)^{-1}+\sigma^{-2}}$ on signals from at least two periods ago. Since the variance of any linear combination of such signals is at least 2, for ε sufficiently small enough the social signal $r_{i,t}$ is bounded away from a perfect estimate of θ_{t-1}. This gives a contradiction. $\square$

Proof of Corollary 1. Consider a complete graph in which all agents have signal variance σ^2 and memory $m = 1$. By Proposition 2, as n grows large the variances of all agents converge to $A > (1+\sigma^{-2})^{-1}$. Choose σ^2 large enough such that $A > 1$.

Now suppose that we increase σ_1^2 to ∞. Then $a_{1,t} = r_{1,t}$ in each period, so all agents can infer all private signals from the previous period. As n grows large, the variance of agent 1 converges to 1 and the variances of all other agents $(1+\sigma^{-2})^{-1}$. By our choice of σ^2, this gives a Pareto improvement. We can see by continuity that the same argument holds for σ_1^2 finite but sufficiently large. $\square$

Proof of Proposition 4. We first check that there is a unique naive equilibrium. As in the Bayesian case, covariances are updated according to equations 4:

$$\boldsymbol{V}_{ii,t} = (w^s_{i,t})^2\sigma_i^2 + \sum W_{ik,t}W_{ik',t}(\rho^2\boldsymbol{V}_{kk',t-1}+1) \text{ and } \boldsymbol{V}_{ij,t} = \sum W_{ik,t}W_{i'k',t}(\rho^2\boldsymbol{V}_{kk',t-1}+1).$$

The weights $W_{ik,t}$ and $w^s_{i,t}$ are now all positive constants that do not depend on $\boldsymbol{V}_{t-1}$. So differentiating this formula, we find that all partial derivatives are bounded above by $1 - w^s_{i,t} < 1$. So the updating map (which we call Φ^{naive}) is a contraction in the sup norm on $\mathcal{V}$. In particular, there is at most one equilibrium.

The remainder of the proof characterizes the variances of agents at this equilibrium. We first construct a candidate equilibrium with variances converging to V_A^∞ and V_B^∞, and then we show that for n sufficiently large, there exists an equilibrium nearby in $\mathcal{V}$.

To construct the candidate equilibrium, suppose that each agent observes the same number of neighbors of each signal type. Then there exists an equilibrium $\widehat{\boldsymbol{V}}^{sym}$ where covariances depend only on signal types, i.e. $\widehat{\boldsymbol{V}}^{sym}$ is invariant under permutations of indices that do not change signal types. We now show variances of the two signal types at this equilibrium converge to V_A^∞ and V_B^∞.

To estimate θ_{t-1}, a naive agent combines observed actions from the previous period with weight proportional to their precisions σ_A^{-2} or σ_B^{-2}. The naive agent incorrectly believes this gives an almost perfect estimate of θ_{t-1}. So the weight on older observations vanishes as $n \to \infty$. The naive agent then combines this estimate of θ_{t-1} with her private signal, with weights converging to the weights she uses if the estimate is perfect.

Agent i observes $\frac{|N_i|}{2}$ neighbors of each signal type, so her estimate $r^{naive}_{i,t}$ of θ_{t-1} is approximately:

$$r^{naive}_{i,t} = \frac{2}{|N_i|(\sigma_A^{-2}+\sigma_B^{-2})}\left[\sigma_A^{-2}\sum_{j\in N_i,\sigma_j^2=\sigma_A^2}\rho a_{j,t-1} + \sigma_B^{-2}\sum_{j\in N_i,\sigma_j^2=\sigma_B^2}\rho a_{j,t-1}\right].$$

The actual variance of this estimate converges to:

$$\mathrm{Var}(r^{naive}_{i,t}-\theta_t) = \frac{\rho^2}{(\sigma_A^{-2}+\sigma_B^{-2})}\left[\sigma_A^{-4}Cov^\infty_{AA} + \sigma_B^{-4}Cov^\infty_{BB} + 2\sigma_A^{-2}\sigma_B^{-2}Cov^\infty_{AB}\right]+1 \qquad (14)$$

where Cov^∞_{AA} is the covariance of two distinct agents of signal type A and Cov^∞_{BB} and Cov^∞_{AB} are defined similarly.

Since agents believe this variance is close to 1, the action of any agent with signal variance σ_A^2 is approximately:

$$a_{i,t} = \frac{r_{i,t}^{naive} + \sigma_A^{-2} s_{i,t}}{1 + \sigma_A^{-2}}.$$

We can then compute the limits of the covariances of two distinct agents of various signal types to be:

$$Cov_{AA}^{\infty} = \frac{\kappa_t^2}{\left(1 + \sigma_A^{-2}\right)^2};\ Cov_{BB}^{\infty} = \frac{\kappa_t^2}{\left(1 + \sigma_B^{-2}\right)^2};\ Cov_{AB}^{\infty} = \frac{\kappa_t^2}{\left(1 + \sigma_A^{-2}\right)\left(1 + \sigma_B^{-2}\right)}.$$

Plugging into 14 we obtain

$$\kappa^{-2} = 1 - \frac{\rho^2}{(\sigma_A^{-2} + \sigma_B^{-2})} \left(\frac{\sigma_A^{-2}}{1 + \sigma_A^{-2}} + \frac{\sigma_B^{-2}}{1 + \sigma_B^{-2}} \right).$$

Using this formula, we can check that the limits of agent variances in $\widehat{\boldsymbol{V}}^{sym}$ match equations 6.

We must check there is an equilibrium near $\widehat{\boldsymbol{V}}^{sym}$ with high probability. Let $\zeta = 1/n$. Let E be the event that for each agent i, the number of agents observed by i with private signal variance σ_A^2 is within a factor of $[1 - \zeta^2, 1 + \zeta^2]$ of its expected value, and similarly the number of agents observed by i with private signal variance σ_B^2 is within a factor of $[1 - \zeta^2, 1 + \zeta^2]$ of its expected value. This event implies that each agent observes a linear number of neighbors and observes approximately the same number of agents with each signal quality. We can show as in the proof of Theorem 1 that for n sufficiently large, the event E occurs with probability at least $1 - \zeta$. We condition on E for the remainder of the proof.

Let $\mathcal{V}_\varepsilon$ be the ε-ball around in $\widehat{\boldsymbol{V}}^{sym}$ the sup norm. We claim that for n sufficiently large, the updating map preserves this ball: $\Phi^{naive}(\mathcal{V}_\varepsilon) \subset \mathcal{V}_\varepsilon$. We have $\Phi^{naive}(\widehat{\boldsymbol{V}}^{sym}) = \widehat{\boldsymbol{V}}^{sym}$ up to terms of $O(1/n)$. As we showed in the first paragraph of this proof, the partial derivatives of Φ^{naive} are bounded above by a constant less than one. For n large enough, these facts imply $\Phi^{naive}(\mathcal{V}_\varepsilon) \subset \mathcal{V}_\varepsilon$. We conclude there is an equilibrium in $\mathcal{V}_\varepsilon$ by the Brouwer fixed point theorem.

Finally, we compare the equilibrium variances to perfect aggregation and to V^∞. It is easy to see these variances are worse than the perfect aggregation benchmark, and therefore by Theorem 1 also asymptotically worse than the Bayesian case when $\sigma_A^2 \neq \sigma_B^2$.

In the case $\sigma_A^2 = \sigma_B^2$, it is sufficient to show that Bayesian agents place more weight on their private signals (since asymptotically action error comes from past changes in the state and not signal errors). Call the private signal variance σ^2. For Bayesian agents, we showed in Theorem 1 that the weight on the private signal is equal to $\frac{\sigma^{-2}}{\sigma^{-2}+(\rho^2 Cov^{\infty}+1)^{-1}}$ where Cov^{∞} solves

$$Cov^{\infty} = \frac{(\rho^2 Cov^{\infty}+1)^{-1}}{[\sigma^{-2}+(\rho^2 Cov^{\infty}+1)^{-1}]^2}.$$

For naive agents, the weight on the private signal is equal to $\frac{\sigma^{-2}}{\sigma^{-2}+1}$, which is smaller since $Cov^{\infty} > 0$. $\square$

Proof of Theorem 2. We provide the proof in the case $m = 1$ to simplify notation. The argument carries through with arbitrary finite memory.

Case (1): Consider an agent l who places positive weight on a rational agent k and positive weight on at least one other agent. Define weights $\overline{W}$ by $\overline{W}_{ij} = W_{ij}$ and $\overline{w}_i^s = w_i^s$ for all $i \neq k$, $\overline{W}_{kj} = (1-\epsilon)W_{kj}$ for all $j \leq n$, and $\overline{w}_k^s = (1-\epsilon)w_k^s + \epsilon$, where W_{ij} and w_i^s are the weights at the initial steady state. In words, agent k places weight $(1-\epsilon)$ on her equilibrium strategy and extra weight ϵ on her private signal. All other players use the same weights as at the steady state.

Suppose we are at the initial steady state until time t, but in period t and all subsequent periods agents instead use weights $\overline{W}$. These weights give an alternate updating function $\overline{\Phi}$ on the space of covariance matrices. Because the weights $\overline{W}$ are positive and fixed, all coordinates of $\overline{\Phi}$ are increasing, linear functions of all previous period variances and covariances. Explicitly, the diagonal terms are

$$[\overline{\Phi}(\boldsymbol{V}_t)]_{ii} = (\overline{w}_i^s)^2 \sigma_i^2 + \sum_{j,j' \leq n} \overline{W}_{ij} \overline{W}_{ij'} V_{jj',t}$$

and the off-diagonal terms are

$$[\overline{\Phi}(\boldsymbol{V}_t)]_{ii'} = \sum_{j,j' \leq n} \overline{W}_{ij} \overline{W}_{i'j'} V_{jj,t'}.$$

So it is sufficient to show the variances $\overline{\Phi}^h(\boldsymbol{V}_t)$ after applying $\overline{\Phi}$ for h periods Pareto dominate the variances in $\boldsymbol{V}_t$ for some h.

In period t, the change in weights decreases the covariance $V_{jk,t}$ of k and some other agent j, who l also observes, by $f(\epsilon)$ of order $\Theta(\epsilon)$. By the envelope theorem, the change in weights only increases the variance V_{kk} by $O(\epsilon^2)$. Taking ϵ sufficiently small, we can ignore

$O(\epsilon^2)$ terms.

There exists a constant $\delta > 0$ such that all initial weights on observed neighbors are at least δ. Then each coordinate $[\Phi(\boldsymbol{V})]_{ii}$ is linear with coefficient at least δ^2 on each variance or covariance of agents observed by i.

Because agent l observes k and another agent, agent l's variance will decrease below its equilibrium level by at least $\delta^2 f(\epsilon)$ in period $t+1$. Because $\overline{\Phi}$ is increasing in all entries and we are only decreasing covariances, agent l's variance will also decrease below its initial level by at least $\delta^2 f(\epsilon)$ in all periods $t' > t+1$.

Because the network is strongly connected and finite, the network has a diameter. After $d+1$ periods, the variances of all agents have decreased by at least $\delta^{2d+2} f(\epsilon)$ from their initial levels. This gives a Pareto improvement.

Case (2): Consider a naive agent k who observes at least two neighbors. We can write agent k's period t action as

$$a_{k,t} = w_k^s s_{i,t} + \sum_{j \in N_i} W_{kj} a_{j,t-1}.$$

Define new weights $\overline{W}$ as in the proof of case (1). Because agent k is naive and the summation $\sum_{j \in N_i} W_{kj} a_{j,t-1}$ has at least two terms, she believes the variance of this summation is smaller than its true value. So marginally increasing the weight on $s_{i,t}$ and decreasing the weight on this summation decreases her action variance. This deviation also decreases her covariance with any other agent. The remainder of the proof proceeds as in case (1). $\square$

Proof of Proposition 5. Suppose agent 1 learns asymptotically. We can assume that agent 1 has at least one neighbor in each G_n. We will discuss the case of rational agents using positive weights.

If agent 1 learns asymptotically, then $\widehat{V}_{ii}(n) < 1$ for n sufficiently large. Fix any n so that $V_{11}(n) < 1$ (to simplify notation we drop references to n for the remainder of the proof).

Then, at equilibrium, any agent i connected to 1 has a best estimator $r_{i,t}$ of θ_t based on observed actions with variance $\kappa_{i,t}^2 = \mathrm{Var}(r_{i,t} - \theta_t)$ less than $1 + \rho^2 \leq 2$. So agent i's action

$$a_{i,t} = \frac{\kappa_{i,t}^{-2} r_{i,t} + \sigma_i^{-2} s_{i,t}}{\kappa_{i,t}^{-2} + \sigma_i^{-2}},$$

puts weight at least $\frac{2}{2+\underline{\sigma}^2}$ on observed actions.

Therefore, in period t, agent 1's best estimator $r_{1,t}$ of θ_{t-1} (indirectly) puts weight at least $\frac{2\rho}{2+\underline{\sigma}^2}$ on actions from period $t-2$. Because

$$\mathrm{Var}(\rho \sum b_j a_{j,t-2} - \theta_{t-1}) = \mathrm{Var}(\rho \sum b_j a_{j,t-2} - \theta_{t-2}) + \mathrm{Var}(\theta_{t-2} - \theta_{t-1}) \geq 1$$

for any *positive* coefficients b_j summing to 1, the variance of $r_{i,t} - \theta_{t-1}$ is at least $\frac{4\rho^2}{(2+\underline{\sigma}^2)^2}$. But then agent 1's action variance is bounded away from $(\sigma_i^{-2}+1)$ and this bound holds for all large enough n, which contradicts our assumption that agent 1 learns asymptotically.

The case where some or all agents are naive agents is similar. $\square$