

SOCIAL LEARNING IN A DYNAMIC ENVIRONMENT

KRISHNA DASARATHA, BENJAMIN GOLUB, AND NIR HAK

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 state? First, if each agent has access to neighbors with sufficiently diverse kinds of signals, then Bayesian learning achieves good information aggregation. Second, without such diversity, there are cases in which Bayesian information aggregation necessarily falls far short of efficient benchmarks. Third, good aggregation can be achieved only if agents “anti-imitate” some neighbors: otherwise, equilibrium estimates are inefficiently confounded by “echoes.” Agents’ 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. We discuss how the resulting tractability is useful for structural estimation of equilibrium learning models and testing against behavioral alternatives.

Date: August 22, 2018.

Department of Economics, Harvard University, Cambridge, U.S.A., dasarath@g.harvard.edu, bgolub@fas.harvard.edu, nirhak@g.harvard.edu. We are grateful to (in random order) Alireza Tahbaz-Salehi, Evan Sadler, Muhamet Yildiz, Philipp Strack, Drew Fudenberg, Nageeb Ali, Tomasz Strzalecki, Jeroen Swinkels, Margaret Meyer, Michael Powell, Michael Ostrovsky, Leeat Yariv, Eric Maskin, Elliot Lipnowski, Kevin He, Eddie Dekel, Annie Liang, Iosif Pinelis, Ozan Candogan, Bob Wilson, Omer Tamuz, Jeff Zwiebel, Matthew O. Jackson, Xavier Vives, Matthew Rabin, and Andrzej Skrzypacz for valuable conversations, and especially to David Hirshleifer for detailed comments on a draft. We also thank numerous seminar participants for helpful questions and comments.

1. INTRODUCTION

Consider a group learning over a period of time about an evolving fundamental state, such as future conditions in a market. In addition to making use of public information, individuals learn from their own private information and also from the estimates of others. For instance, farmers who are trying to assess the demand for a crop they produce may learn from neighbors' actions (e.g., how much they are investing in the crop), which reflect those neighbors' estimates of market conditions.¹ In another example, economic analysts or forecasters have their own data and calculations, but they may also have access to the reports of some other analysts.² Importantly, in many such settings, people have access to the estimates of only some others.³ Therefore, without a central information aggregation device, aggregation of information occurs locally and estimates may differ across a population.

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? These questions are important from a positive perspective, to better understand the determinants of information aggregation and the welfare implications. They will also be relevant in design decisions—e.g., for a planner who is considering changing the environment to facilitate better learning.

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).⁴ Nevertheless, there is relatively little modeling of dynamic states in the large literature on social learning and information aggregation in networks,⁵ though we discuss some

¹ A literature in economic development studies situations in which the flow of information crucial to production decisions is constrained by geographic or social distance; see, e.g., Jensen (2007); Srinivasan and Burrell (2013).

² In a class of models of over-the-counter markets, agents learn about one another's valuations when they negotiate or trade (see, e.g., Vives, 1993; Duffie and Manso, 2007; Duffie, Malamud, and Manso, 2009; Babus and Kondor, 2018).

³ Even when information is available, not everyone will acquire all the available information. This may be because information is costly, or—even if the information itself is free—because understanding some information requires investing in familiarity with its source. The latter is an effect emphasized by Sethi and Yildiz (2016).

⁴ "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.

⁵ See, among many others, DeMarzo, Vayanos, and Zweibel (2003), Acemoglu, Dahleh, Lobel, and Ozdaglar (2011), Mueller-Frank (2013), Eyster and Rabin (2014), Mossel, Sly, and Tamuz (2015), Lobel and Sadler

very important antecedents that we build on—including Frongillo, Schoenebeck, and Tamuz (2011), Shahrampour, Rakhlin, and Jadbabaie (2013), and Alatas, Banerjee, Chandrasekhar, Hanna, and Olken (2016)—in Section 6. Our first contribution is to define and study equilibria in a dynamic environment that captures two essential dimensions emphasized above: the state of the world changes over time, and communication occurs in an arbitrary network. The second contribution is to derive conditions under which decentralized information aggregation works well. We find that when states are dynamic, the answers to the aggregation question are considerably different relative to models where the state is static.

Our main substantive finding is that in large populations decentralized learning *can* approach an essentially optimal benchmark, as long as (i) each individual has access to a set of neighbors that is sufficiently diverse, in the sense of having different signal distributions from each other; and (ii) updating rules are Bayesian, responding to correlations in a sophisticated way. If signal endowments are *not* diverse, then social learning can be inefficiently confounded and far from optimal, even though each agent has access to an unbounded number of observations, each containing independent information. Diversity is, in a sense, more important than precision: giving everyone better signals can hurt aggregation severely if it makes those signals homogeneous. Rationality is also very important: If, instead of using Bayesian learning rules, agents use certain intuitive heuristics that are known to work well in static-state models, learning outcomes are, again, far from efficient.

In addition to specific predictions about social learning, the results also have methodological implications. An organizing question for the literature on learning in networks is: “How does the interplay between network structure and information structure affect the quality of learning outcomes?” The answers in our setting depend on aspects of the environment different from those that matter in canonical social learning models, as we discuss below and in Section 6. So our dynamic-state perspective reveals a new set of forces and offers tools for studying them, expanding the theory of social learning in networks.

We now describe the model and some key intuitions in more detail. The state, θ_t , drifts around according to a stationary, discrete-time AR(1) process given by $\theta_{t+1} = \rho\theta_t + \nu_{t+1}$, with $0 < \rho < 1$, 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. 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 . In each period, the information the agent observes

(2015a), Akbarpour, Saberi, and Shamsi (2017), and Molavi, Tahbaz-Salehi, and Jadbabaie (2018). There is also a literature in finance on information aggregation in complex environments, where the models are mainly static: see, for instance, Malamud and Rostek (2017), Lambert, Ostrovsky, and Panov (2018), and the papers cited there.

consists of an independent, normal signal $s_{i,t} \sim \mathcal{N}(\theta_t, \sigma_i^2)$ of the current state and some past estimates of her neighbors in an arbitrary network; these estimates are relevant for estimating θ_t because they depend on recent states.⁶ 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 their beliefs: 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.

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 standard one: estimating a linear statistical model. However, the linear combinations of the underlying fundamentals that everyone observes depend on others' strategies, i.e., updating rules.⁷ Agents' (linear) strategies in a given period determine the covariance matrix of their estimates, which determines each agent's weights on her information (neighbors' estimates and private signals about the state of the world). Those weights, in turn, determine the covariances of next period's estimates. Each stationary equilibrium corresponds to a covariance matrix that is invariant under this process. On the one hand, the logic of the fixed point, enabled by the stationarity of the model, is simple. Indeed, it extends to more general learning environments, such as ones where agents observe others'

⁶ The network can be directed or undirected; that is, observation opportunities need not be reciprocal.

⁷ In engineering terminology, each agent uses a Kalman filter, but the distribution of her observations is determined by other agents' behavior; thus, we analyze a set of Kalman filters in a network in Nash equilibrium; cf. Olfati-Saber (2007) and Shahrampour, Rakhlin, and Jadbabaie (2013), which model a set of distributed Kalman filters controlled by a planner.

estimates with noise. On the other hand, determining the precisions of agents' estimates is subtle, and these precisions depend on the whole equilibrium.⁸

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. We show—within a random graph model that can capture essentially arbitrary network heterogeneity—these properties are enough to guarantee that information aggregation is as good as it can possibly be. More precisely, 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—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, who get conditionally independent signals 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 many time- t agents (j, t) . Imagine that each of these observed actions, $a_{j,t} = \theta_t + \epsilon_{j,t} + c_i \xi_t$, includes an idiosyncratic shock $(\epsilon_{j,t})$ but is also confounded by a common piece of noise (ξ_t) . As we will see, this is exactly the situation that arises endogenously in our model. Then even if i has many neighbors and their $\epsilon_{j,t}$ are uncorrelated, the presence of the common confound terms may prevent her from learning θ_t precisely. In particular, if all the c_i are the same, there is nothing i can do to filter the confound out of her estimates.

One of our key observations, which we now explain, is that the structure of private signals governs whether this obstruction can arise. The role of the common confound ξ_t is played by the common dependence of everyone's action on the past before period t . When the neighbors of i have diverse private signal distributions, the common confound ξ_t will be *weighted differently* by neighbors with different signal precisions.⁹ As a result, i will be able to filter out the confound and identify θ_t precisely. But when the neighbors all have identically distributed private signals, at least in some networks, their actions will incorporate data from the past before time t in a symmetric way.¹⁰ That, in turn, will force any estimate i makes to

⁸ Equilibrium weights satisfy a system of polynomial equations, but the system usually has high degree and is not amenable to explicit solutions, so additional work is required to characterize properties of these solutions.

⁹ We assume that agents know enough about the environment—either the primitives or at least equilibrium distributions of neighbors' estimates—to set these weights optimally. We discuss these assumptions more precisely in the formal presentation of results.

¹⁰ This is related to a force identified by [Sethi and Yildiz \(2012\)](#) in a fixed-state world, where what matters is a specific kind of uncertainty about privately-known prior beliefs. In their setting, agents cannot learn well

depend on echoes of the distant past. That the possibility of effective filtering *depends on the diversity of signal distributions* distinguishes our dynamic learning environment from a model with an unchanging state ($\theta_t = \theta$ for all t). Indeed, with an unchanging state, efficient aggregation does not depend on such details of signal structure when it can be obtained (DeMarzo, Vayanos, and Zweibel, 2003; Banerjee and Fudenberg, 2004; Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi, 2012). And, on the other hand, obstacles that confound learning also do not depend on whether signals are diverse or homogeneous in their distributions (Smith and Sørensen, 2000; Sethi and Yildiz, 2012; Harel, Mossel, Strack, and Tamuz, 2017). Section 4 fleshes out the ideas we have sketched here.

To argue that neither our technical conditions nor a reliance on very large numbers is driving the conclusions, we calculate equilibria numerically for real social networks both for diverse and non-diverse signal endowments, and show that signal diversity does enable much better learning, even with degrees that are not very large (about 15 on average).

Efficient aggregation depends on another factor, beyond the sufficient diversity of signal precisions we have been discussing: 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 error. 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 a feature of some well-known non-Bayesian heuristic learning rules in network models, most prominently the DeGroot model and its elaborations (Golub and Jackson, 2010; Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi, 2012). Under this condition, information aggregation in our model is essentially guaranteed to fall short of good aggregation benchmarks for all agents. This distinguishes our dynamic setting from models with an unchanging state, in which, as the work just mentioned shows, even without much sophistication agents are able to learn well.

Many of our theoretical 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.

when this uncertainty takes an uncorrelated form. See also Ostrovsky (2012) and the literature cited there for related obstacles to inference in trading models.

A design implication of our results is that if agents are sophisticated (i.e., Bayesian), then an authority (e.g., a manager) who values good aggregation would like to endow them with a variety of signal distributions, even at the cost of some signals now having lower precision. Practically, this can be thought of as the planner valuing the group having a diversity of perspectives in their decentralized discussion. This diversity eases identification problems and enables the group to make the best use of local information.

Modeling contributions. Several features of the model permit progress in places where social learning models often present technical challenges, as we now highlight. Because of its Gaussian, stationary structure, we can study equilibria with linear learning rules. Linearity facilitates our numerical analysis of the model as well as several other kinds of exercises:

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

Much of the prior work of linear social learning rules treats updating rules as heuristics involving exogenous weights. Our approach, in which agents endogenously choose weights to maximize the precisions of their estimates, enables several further exercises:

3. Standard welfare analysis: 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 interpretation. 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 the social learning literature. 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. We describe the environment and game; complete details are formalized in Appendix A.

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 ρ is a constant with $0 < |\rho| \leq 1$ 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 $N = \{1, 2, \dots, n\}$. Each node i 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" $\ell \in \{1, 2, \dots, m\}$. (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, as is the informational environment (i.e., all precisions, etc.). We will sometimes take 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

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

The agent makes the optimal choice for the current period given her information—i.e., does not seek to affect future actions.¹¹ By a standard fact about squared-error loss functions, given the distribution of $(\mathbf{a}_{N_i,t-\ell})_{\ell=1}^m$, she sets:

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

Here the notation $\mathbf{a}_{N_i,t}$ refers to the vector $(a_{j,t})_{j \in N_i}$. An action can be interpreted as an agent’s estimate of the state, and we will sometimes use this terminology.

The conditional expectation (1) depends on the prior of agent (i, t) about θ_t , which can be any normal distribution or a uniform improper prior (in which case all of i ’s beliefs about θ_t come from her own signal and her neighbors’ actions).¹² We take priors, like the information structure and network, to be common knowledge. In the rest of the paper, we formally analyze the case where all agents have improper priors. Because actions under a normal prior are related to actions under the improper prior by a simple bijection—and thus have the same information content for other agents—all results extend to the general case.

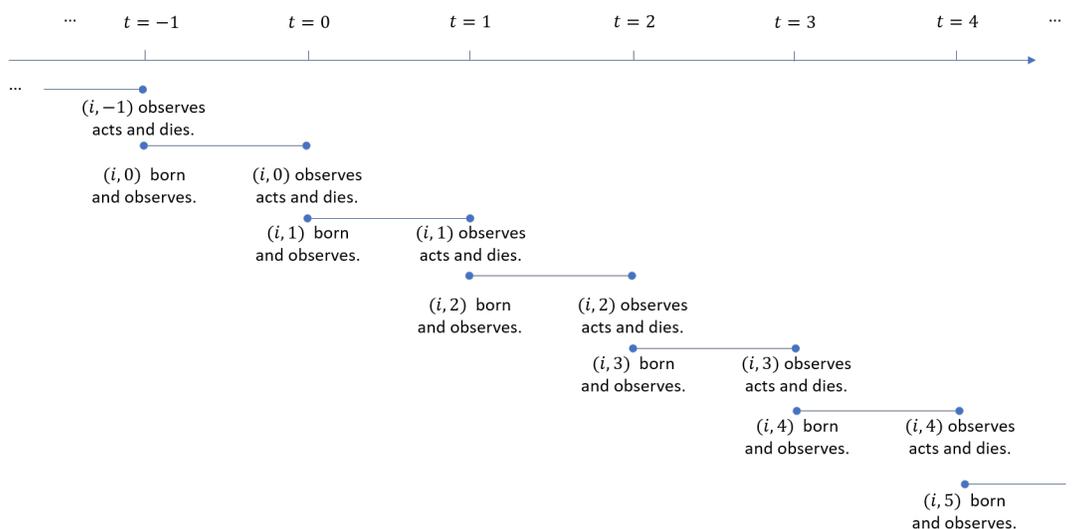
2.2. Interpretation. The agents are fully Bayesian given the information they have access to. Much of our analysis is done for an arbitrary finite m ; we view the restriction to finite memory as an assumption that avoids technical complications, but because m can be arbitrarily large, this restriction has little substantive content. The model generalizes “Bayesian without Recall” agents from the engineering and computer science literature (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 they 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,

¹¹ In Section 2.2 we discuss this assumption and how it relates to applications.

¹² With $0 < \rho < 1$, one natural choice for a prior is the stationary distribution of the state.

FIGURE 1. An illustration of the overlapping generations structure of the model for $m = 2$.



At time $t - 1$, agent (i, t) is born and observes estimates from time $t - 2$. At time t agent (i, t) observes estimates from $t - 1$, her private signal $s_{i,t}$ and submits her estimate $a_{i,t}$.

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.¹³

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 they shut down 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),¹⁴ 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.¹⁵ Equivalently, we could simply assume that agents sincerely announce their subjective expectations of the state, as in Geanakoplos and Polemarchakis (1982) and the extensive literature following it.

¹³ 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.

¹⁴ See also Manea (2011) and Talamàs (2017) in the bargaining literature.

¹⁵ But see Mueller-Frank (2013) for some discussion of why incentives for deception are not a first-order concern in this environment.

3. EQUILIBRIUM

In this section we present the substance of our notion of equilibrium and the basic existence result.¹⁶

3.1. Equilibrium in linear strategies. A strategy of an agent is *linear* if the action taken is a linear function of the variables in her information set. 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, including non-linear ones. 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 weighted summation of past signals.¹⁷ Because all innovations ν_t and signal errors $\eta_{i,t}$ are independent and Gaussian, it follows that 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, $\mathbb{E}[\theta_t \mid s_{i,t}, (\mathbf{a}_{N_i,t-\ell})_{\ell=1}^m]$ is a linear function of $s_{i,t}$ and $(\mathbf{a}_{N_i,t-\ell})_{\ell=1}^m$ (see (2) below for details). It follows that solving for equilibrium can be reduced to searching for the weights agents place on the variables in their information set.

A reason for focusing on equilibria in linear strategies comes from considering how agents would behave in a variant of the model where time began at $t = 0$. In that first period, agents would see only their own signals, and therefore play linear strategies; after that, inductively applying the argument in the previous paragraph shows that strategies would be linear at all future times. This is the thought experiment that motivates our focus on linear strategies; taking time to extend infinitely backward is an idealization that allows us to focus on exactly stationary behavior.

3.2. Covariance matrices. The optimal weights for an agent to place on her sources of information depend on the precisions and covariances of these sources, and so we now study these.

Given a linear strategy profile played up until time t , let \mathbf{V}_t be the $nm \times nm$ covariance matrix of the vector $(\rho^\ell a_{i,t-\ell} - \theta_t)_{i \in N, 0 \leq \ell \leq m-1}$. The entries of this vector are the differences between the best predictors of θ_t based on actions $a_{i,t-\ell}$ during the past m periods and the current

¹⁶ Because time in this game is doubly infinite, there are some subtleties in definitions, which are dealt with in Appendix A.

¹⁷ 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.

state of the world. In the case $m = 1$, this is simply the covariance matrix $\mathbf{V}_t = \text{Cov}(a_{i,t} - \theta_t)$. The matrix \mathbf{V}_t records covariances of action errors: diagonal entries measure the accuracy of each action, while off-diagonal entries indicate how correlated the two agents' action errors are. The entries of \mathbf{V}_t are denoted by $V_{ij,t}$.

3.3. Best-response weights. A strategy profile is an equilibrium if the weights each agent places on the variables in her information set minimize her posterior variance.

We now characterize these in terms of the covariance matrices we have defined. Consider an agent at time t , and suppose some linear strategy profile has been played up until time t . Let $\mathbf{V}_{N_i,t-1}$ be a sub-matrix of \mathbf{V}_{t-1} that contains only the rows and columns corresponding to neighbors of i ¹⁸ and let

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

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

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

(see Example 4.4 of Kay (1993)). This gives $\mathbb{E}[\theta_t \mid s_{i,t}, (\mathbf{a}_{N_i,t-\ell})_{\ell=1}^m]$ (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 \mathbf{V}_{t-1} (but not on realizations of any random variables). In (2) we have taken, for convenience, agents' prior beliefs about the state to be an improper distribution giving all states equal weight.¹⁹

¹⁸ Explicitly, $\mathbf{V}_{N_i,t-1}$ are the covariances of $(\rho^\ell a_{j,t-\ell} - \theta_t)$ for all $j \in N_i$ and $\ell \in \{1, \dots, m\}$.

¹⁹ Our analysis applies equally to any proper normal prior for θ_t : To get an agent's estimate of θ_t , the formula in (2) would simply be averaged with a constant term accounting for the prior, and everyone could invert this deterministic operation to recover the same information from others' actions. Our approach simply saves on notation.

We denote by (W_t, w_t^s) a *weight profile* 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.²⁰ When $m = 1$, we refer to the weight agent i places on $a_{j,t-1}$ (agent 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 \mathbf{V}_t from the previous covariance matrix. This defines a map $\Phi : \mathcal{V} \rightarrow \mathcal{V}$, given by

$$(3) \quad \Phi : \mathbf{V}_{t-1} \mapsto \mathbf{V}_t$$

which we study in characterizing equilibria.

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

$$(4) \quad \mathbf{V}_{ii,t} = (w_{i,t}^s)^2 \sigma_i^2 + \sum W_{ik,t} W_{ik',t} (\rho^2 \mathbf{V}_{kk',t-1} + 1) \text{ and } \mathbf{V}_{ij,t} = \sum W_{ik,t} W_{ik',t} (\rho^2 \mathbf{V}_{kk',t-1} + 1).$$

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

3.4. Equilibrium existence. Consider the map Φ defined in (3). Stationary equilibria in linear strategies correspond to fixed points of the map Φ .²¹

Our first result concerns the existence of equilibrium:

Proposition 1. *A stationary equilibrium in linear strategies exists, and is associated with a covariance matrix $\widehat{\mathbf{V}}$ such that $\Phi(\widehat{\mathbf{V}}) = \widehat{\mathbf{V}}$.*

The proof appears in Appendix B.

At the stationary equilibrium, the covariance matrix and all agent strategies are time-invariant. Actions are linear combinations of observations with stationary weights (which we refer to as $\widehat{W}_{ij,t}$ and \widehat{w}_i^s). The form of these rules has some resemblance to static equilibrium notions studied in the rational expectations literature (e.g., Vives (1993); Babus and Kondor (2018); Lambert, Ostrovsky, and Panov (2018); Mossel, Mueller-Frank, Sly, and Tamuz (2018)), but here we explicitly examine the dynamic environment in which these emerge as steady states. We discuss the relationship between our model and DeGroot learning, which has a related form, in Section 6.

The idea of the argument is as follows. The goal is to apply the Brouwer fixed-point theorem to show there is a covariance matrix $\widehat{\mathbf{V}}$ that remains unchanged under updating. To find a

²⁰ We do not need to describe the indexing of coefficients in W_t explicitly in general; this would be a bit cumbersome because there are weights on actions at various lags.

²¹ More generally, one can use this map to study how covariances of actions evolve given any initial distribution of play. Note that the map Φ is deterministic, so we can study this evolution without considering the particular realizations of signals.

compact set to which we can apply the fixed-point theorem, we use the fact that when agents best respond to any beliefs about prior actions, all variances are bounded and bounded away from 0. This is 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, it follows that every element of any \mathbf{V} in the image of Φ is bounded.²² These bounds along with the continuity of Φ allow us to apply the Brouwer fixed-point theorem.

We have given, in discussing equation (4), an explicit description of Φ . We can now use this to give an explicit version of the fixed-point condition: The equilibrium variances and covariances $\widehat{\mathbf{V}}$ satisfy

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

(or corresponding equations with $m > 1$). We can use (2), which gives a formula in terms of $\widehat{\mathbf{V}}$ for the weights \widehat{W}_{ij} and \widehat{w}_i^s in the best response to $\widehat{\mathbf{V}}$, in order to write the equilibrium \widehat{V}_{ij} as the solutions to a system of polynomial equations. These equations have large degree and cannot be solved analytically except in very simple cases, but they can readily be used to solve for equilibria numerically.²³

The main insight is that we can find equilibria by studying action covariances; this idea applies equally to many extensions of our model. We give two examples: (1) We assume that agents observe neighbors perfectly, but one could define other observation structures. For instance, agents could observe actions with noise, or they could observe some set of linear combinations of neighbors' actions with noise. (2) We assume agents are Bayesian and best-respond rationally to the distribution of actions, but the same proof would also show that equilibria exist under other behavioral rules.²⁴

We show later, as part of Proposition 2, that there is a unique stationary linear equilibrium in networks having a particular structure. In general, uniqueness of the equilibrium is an open question that we leave for future work; our efforts to use standard approaches for proving

²² 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.

²³ Indeed, we have used numerical solutions to study the system and to conjecture many of our results. In practice, a fixed point $\widehat{\mathbf{V}}$ is found by repeatedly applying Φ , as written in (4), to an initial covariance matrix. In all our experiments, the same fixed point has been found, independent of starting conditions.

²⁴ What is important in the proof is that actions depend continuously on the covariance structure of an agent's observations; the action variances are uniformly bounded under the rule agents play; and there is a decaying dependence of behavior on the very distant past.

uniqueness have run into obstacles.²⁵ Nevertheless, in computing equilibria numerically for many examples, we have not been able to find a case of equilibrium multiplicity.

We now discuss how agents could come to play the strategies posited above. If other agents are using stationary equilibrium strategies, then best-responding is easy to do under some conditions. For instance, if historical empirical data on neighbors' error variances and covariances are available (i.e., the entries of the matrix $\mathbf{V}_{N_i,t}$ discussed in Section 3.3), then the agent needs only to use these to compute a best estimate of θ_{t-1} , which is essentially a linear regression problem.²⁶

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 expected payoff that 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.) Let $V_{ii}^{\text{benchmark}}$ be the error variance that player i achieves at this benchmark: namely, $V_{ii}^{\text{benchmark}} = (\sigma_i^{-2} + 1)^{-1}$.

Definition 1. An equilibrium achieves the ε -perfect aggregation benchmark if, for all i ,

$$\frac{\widehat{V}_{ii}}{V_{ii}^{\text{benchmark}}} \leq 1 + \varepsilon.$$

This says that all agents do nearly 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 the error variance V_{ii} at a steady state, which need not come from rational agents optimizing.²⁷ Note agents can never infer yesterday's state perfectly from observed actions in any finite network, and so we must have $\frac{\widehat{V}_{ii}}{V_{ii}^{\text{benchmark}}} > 1$ for all i on any fixed network.

²⁵ We have checked numerically that Φ is not, in general, a contraction in any of the usual norms (entrywise sup, Euclidean operator norm, etc.), nor does it seem clear how to prove uniqueness by defining a Lyapunov function.

²⁶ Even if other agents are not yet using equilibrium strategies, the procedure we have described will result in best-responding to the historical average of \mathbf{V}_t . Thus, if the dynamics of repeatedly applying Φ repeatedly converge in an appropriate sense, this will lead to an equilibrium. (The proof of Proposition 2 shows this is a class of graphs.) But we leave a full theory of learning to play equilibria in such an environment for future work.

²⁷ One can also extend the definition to cover cases where there is not a steady state, for example by considering the lim sup of error variances: $\limsup_t V_{ii,t}$.

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 ε -perfect aggregation benchmark can be achieved for small ε . Suppose each G_n for $n \geq 2$ has a connected component with exactly two agents, 1 and 2, with $\sigma_1^2 = 1$ and $\sigma_2^2 = 1/n$. Then agent 2's weight on her own signal converges to 1 as $n \rightarrow \infty$. So \widehat{V}_{11} converges to $(\sigma_1^{-2} + 1)^{-1} = \frac{1}{2}$ as $n \rightarrow \infty$. Thus, the learning benchmark is achieved.

The environment we have devised in this example is quite special: agent 1 can essentially infer last period's state because someone else has arbitrarily precise information. A much more interesting question is whether anything similar can occur without anyone having extremely precise signals. In Section 4.1 we address this and show that perfect aggregation can be achieved by all agents simultaneously even without anyone having very precise signals.

4.1. Diverse signals. One of our main results is that the perfect aggregation benchmark is achieved if private signals are diverse in their quality and agents' neighborhoods are large enough. The reasons for this are fairly intuitive and we give them below. At the same time, we seek to make this point reasonably generally, and the most general version of this result we have been able to find is set in a certain kind of random graph environment. There are many agents, each having one of several possible types of signals, and links among agents are realized randomly, but with a very flexible structure (formally, a stochastic block model).

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 sets $N_n^1, N_n^2, \dots, N_n^K$. We say the agents in N_n^k have network type k . An 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. An assumption we maintain on these probabilities is that each network type k observes at least one network type (possibly k itself) with positive probability.

Suppose there are finitely many possible private signal variances.²⁸ We say a signal variance is represented in a network type if, as $n \rightarrow \infty$, at least a positive share of agents of that type have that signal variance.²⁹ Finally, we assume that, for each network type, there are at least two distinct signal variances that are represented in that type.

²⁸ The assumptions of finitely many signal and network types are purely technical, and could likely be relaxed.

²⁹ Formally, the \liminf of the share of agents with that signal variance is positive.

We say an event occurs *asymptotically almost surely* if for any $\varepsilon > 0$, the event occurs with probability at least $1 - \varepsilon$ for n sufficiently large.

Theorem 1. *Let $\varepsilon > 0$. Under the assumptions in this subsection, asymptotically almost surely G_n has a equilibrium where the ε -perfect aggregation benchmark is achieved.*

So on large networks, society is very likely to aggregate information as well as possible. The uncertainty in this statement 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). We give an outline of the argument next, and the proof appears in Appendix C.

Outline of the argument. To give intuition for the result, we first describe why the theorem holds on the complete network³⁰ in the $m = 1$ case. We then discuss the challenges involved in generalizing the result to our general stochastic block model networks, and the techniques we use to overcome those challenges.

Consider a time- t agent, (i, t) . We define her *social signal* $r_{i,t}$ to be the optimal estimate of θ_{t-1} based on the actions she has observed in her neighborhood. On the complete network, all players have the same social signal, which we call r_t .³¹

At any equilibrium, each agent's action is a weighted average of her private signal and this social signal:³²

$$(5) \quad a_{i,t} = \widehat{w}_s^i s_{i,t} + (1 - \widehat{w}_s^i) r_t.$$

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

Now observe that each time- $(t+1)$ agent can average the time- t actions of each type, which can be written as follows using (5) and $s_{i,t} = \theta_t + \eta_{i,t}$:

$$\begin{aligned} \frac{1}{n_A} \sum_{i:\sigma_i^2=\sigma_A^2} a_{i,t} &= \widehat{w}_s^A \theta_t + (1 - \widehat{w}_s^A) r_t + O(n^{-1/2}), \\ \frac{1}{n_B} \sum_{i:\sigma_i^2=\sigma_B^2} a_{i,t} &= \widehat{w}_s^B \theta_t + (1 - \widehat{w}_s^B) r_t + O(n^{-1/2}). \end{aligned}$$

Here n_A and n_B denote the numbers of agents of each type, and the $O(n^{-1/2})$ error terms come from averaging the signal noises $\eta_{i,t}$ of agents in each group. In other words, by the law of large numbers, each time- $(t+1)$ agent can obtain precise estimates of two different

³⁰ Note this is a special case of the stochastic block model.

³¹ In particular, agent (i, t) sees everyone's past action, including $a_{i,t-1}$.

³² Agent i 's weights on her observations $s_{i,t}$ and $\rho a_{j,t-1}$ sum to 1, because the optimal action is an unbiased estimate of θ_t .

convex combinations of θ_t and r_t . Because the two weights, \widehat{w}_s^A and \widehat{w}_s^B , are distinct, she can approximately solve for θ_t as a linear combination of the average actions from each type (up to signal error). It follows that in the equilibrium we are considering, the agent must have an estimate at least as precise as what she can obtain by the strategy we have described, and will thus be very close the benchmark. The estimator of θ_t that this strategy gives will place negative weight on $\frac{1}{n_A} \sum_{i:\sigma_i^2=\sigma_B^2} a_{i,t-1}$, thus *anti-imitating* the agents of signal type A. It can be shown that the equilibrium we construct in which agents learn will also have agents anti-imitating others.

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. This is not easy, because the social signals at an equilibrium are endogenous, and in a general network will depend to some extent on many details of the network.

A key insight allowing us to overcome this difficulty and get a handle on social signals is that the *number of paths of length two* between any two agents is nearly deterministic in our random graph model. While any two agents of the same network type may have very different neighborhoods, their connections at distance two will typically look very similar.³³ This gives us a nice expression for the social signal as a combination of private signals and the social signals from two periods earlier. 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. We use this to show that Φ^2 maps the neighborhood of covariance matrices where all social signals are close to perfect to itself, and then we apply a fixed point theorem. \square

We have assumed 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 neighbors' signal types, she could solve her estimation problem 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 older social signals to form a precise estimate of the previous state, thus achieving the benchmark. 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

³³ As we note below, “length two” here is not essential: in other classes of random graphs, this statement holds with a different length, and the same arguments go through.

of laws of large numbers, their rate of growth can be considerably slower than linear: for example, our proof extends to degrees that scale as n^α for any $\alpha > 0$.³⁴

Concerning the rate of learning as n grows, the proof implies that, under the assumptions of the theorem, the error in agents' estimates of θ_{t-1} is $O(n^{-1/2})$; thus they learn at the same rate as in the central limit theorem, though the constants will depend considerably on the network. Section 4.3 offers numerical evidence on the quality of aggregation in networks of practically relevant sizes.

4.2. Non-diverse signals. It is essential to the argument from the previous subsection that different agents have different signal precisions. Recall the complete graph case we examined in our outline of the argument in the previous subsection. From the perspective of an agent $(i, t+1)$, the fact that type A and type B neighbors place different weights on the social signal r_t allows i to prevent the social signal used by her neighbors from confounding her estimate of θ_t . We now show that without diversity in signal quality, information aggregation may be much slower.

We first study a class of very structured networks and show that for this class there is a unique equilibrium and at this equilibrium good aggregation is not achieved. We then present an immediate corollary of this result showing that improving some agents' signals can hurt learning, which distinguishes this regime not only in terms of its outcomes but also in its comparative statics. Finally, we show that our conclusions go beyond the very structured networks: in particular, there is a similar equilibrium without good aggregation on random graphs.

4.2.1. Graphs with symmetric neighbors.

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.³⁵ 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 .*

³⁴ 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 $1/\alpha$.

³⁵ These are both special cases of our stochastic block model from Section 4.1.

All agents are bounded away from our learning benchmark at the unique equilibrium. So all agents learn poorly compared to the diverse signals case. The proof of this proposition, and the proofs of all subsequent results, appear in Appendix D.

The failure of good aggregation is not due simply to a lack of sufficient information in the environment: On the complete graph with exchangeable (i.e., non-diverse) signals, a social planner who exogenously set weights for all agents could achieve ε -perfect aggregation for any $\varepsilon > 0$ when n is large. See Appendix E for a formal statement, proof and numerical results.³⁶

We now give intuition for Proposition 2. In a graph with symmetric neighbors, in the unique equilibrium, the actions of any agent's neighbors are exchangeable.³⁷ So actions must be unweighted averages of observations. This prevents the sort of inference of θ_t that occurred with diverse signals. This is easiest to see on the complete graph, where *all* observations are exchangeable. So, in any equilibrium, each agent's action at time $t + 1$ is equal to a weighted average of his own signal and $\frac{1}{n} \sum_{j \in N_i} a_{j,t}$:

$$(6) \quad a_{i,t+1} = \widehat{w}_s^i s_{i,t+1} + (1 - \widehat{w}_s^i) \frac{1}{n} \sum_{j \in N_i} a_{j,t}.$$

By iteratively using this equation, we can see that actions must place substantial weight on the average of signals from, e.g., two periods ago. 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 never “washes out” of estimates, and this is what prevents perfect aggregation.

We can also explicitly characterize the limit action variances and covariances. Consider again the complete graph and the (unique) symmetric equilibrium. Let V^∞ denote the limit, as n grows large, of the variance of any agent's error ($a_{i,t} - \theta_t$). Let Cov^∞ denote the limit covariance of any two agent's errors. By direct computations, these can be seen to be related by the following equations, which have a unique solution:

$$(7) \quad V^\infty = \frac{1}{\sigma^{-2} + (\rho^2 Cov^\infty + 1)^{-1}}, \quad Cov^\infty = \frac{(\rho^2 Cov^\infty + 1)^{-1}}{[\sigma^{-2} + (\rho^2 Cov^\infty + 1)^{-1}]^2}.$$

This variance and covariance describe behavior not only in the complete graph, but to any graph with symmetric neighbors where degrees tend uniformly to ∞ . In such graphs, too, the variances of all agents converge to V^∞ and the covariances of all pairs of agents converge to Cov^∞ , as $n \rightarrow \infty$.³⁸ This implies that, in large graphs, the equilibrium action distributions are close to symmetric. Indeed, it can be deduced that these actions are equal to an appropriately discounted sum of past $\theta_{t-\ell}$, up to error terms (arising from $\eta_{i,t-\ell}$) that vanish asymptotically.

³⁶ We thank Alireza Tahbaz-Salehi for suggesting this analysis.

³⁷ The proof of the proposition establishes uniqueness by showing that Φ is a contraction in a suitable sense.

³⁸ This is established by the same argument as in the proof of Proposition 3.

4.2.2. *A corollary: Perverse consequences of improving signals.* As a consequence of Theorem 2 and Proposition 2, 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 Proposition 2, all agents do substantially worse than perfect aggregation. If we instead give agent 1 a very uninformative signal, all players can anti-imitate agent 1 and achieve nearly perfect aggregation. When the signals at the initial configuration are sufficiently imprecise, this gives a Pareto improvement.

4.2.3. *Non-diverse signals in large random graphs.* Our results so far have in this subsection have relied on graphs with symmetric neighbors. There, the unique prediction is that learning outcomes fall far short of the perfect aggregation benchmark. We would like to show that exact symmetry is not essential, and that the lack of good aggregation is robust to adding noise. To this end, we now show that in Erdős–Rényi random networks, there is an equilibrium with essentially the same learning outcomes when signal precisions are homogeneous.

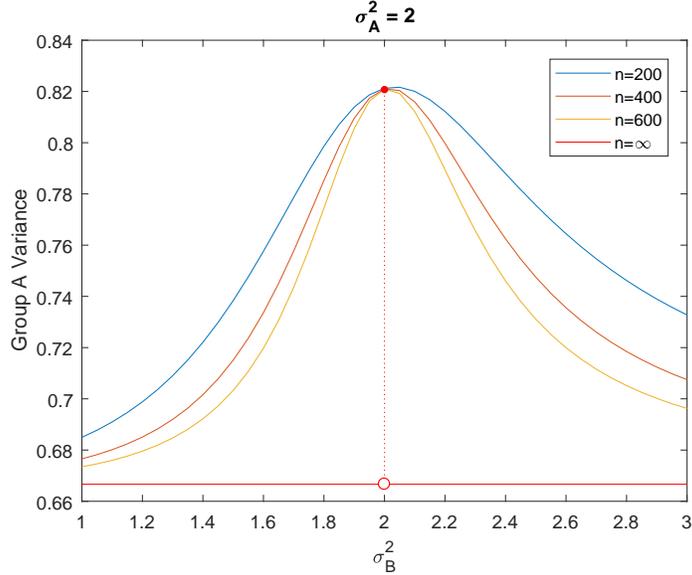
Let $(G_n)_{n=1}^\infty$ be a sequence of undirected random networks, with G_n having n nodes, with any pair of distinct nodes linked (i.i.d.) with positive probability p . We continue to assume all signal variances are equal to σ^2 and $m = 1$.

Proposition 3. *Under the assumptions in the previous paragraph, there exists $\varepsilon > 0$ such that asymptotically almost surely there is an equilibrium on G_n where the ε -perfect aggregation benchmark is not achieved.*

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

4.3. Aggregation and its absence without asymptotics: Numerical results. 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 at all, then as $n \rightarrow \infty$, perfect aggregation is achieved; with exactly homogeneous signal endowments, agents' variances are much higher. In this section, we show numerically that for fixed values of n , the transition from the first regime

FIGURE 2. Distinct Variances Result in Learning



to the second is actually gradual: Action error remains well above the perfect aggregation benchmark when signal qualities differ slightly.

In Figure 2, we study the complete network with $\rho = 0.9$. The private signal variance of agents of signal type A is fixed at $\sigma_A^2 = 2$. We then vary the private signal variance of agents of type B (the horizontal axis), and compute the equilibrium variance of $a_{i,t} - \theta$ for agents of type A (plotted on the vertical axis). The variance of type A agents at the benchmark is $2/3$. We note several features: First, the change in aggregation quality 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 discontinuity 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 5.2% in precision when $\sigma_B^2 = 3$.

To examine whether the large network results above work in realistic networks with moderate degrees, we present numerical evidence based on the data in Banerjee, Chandrasekhar, Duflo, and Jackson (2013). This data set contains the social networks of villages in rural India.³⁹ There are 43 networks in the data, with an average network size of 212 nodes (standard deviation = 53.5), and an average degree of 19 (standard deviation = 7.5). For each network, we calculated the equilibrium for two different situations. The first is the homogeneous case, with all signal variances set to 2. The latter is a heterogeneous case, where a majority has the

³⁹ We take the networks that were used in the estimation in Banerjee, Chandrasekhar, Duflo, and Jackson (2013). As in their work, we take every reported relationship to be reciprocal for the purposes of sharing information. This makes the graphs undirected.

same signal distribution as in the first case, but a minority has a substantially worse signal. More precisely, we kept the signal variances of people that have access to electricity (92% of the nodes) at 2, while setting the signal variances of the rest at 5.⁴⁰

In Figure 3(a), the green points show that in the vast majority of networks, the median agent in terms of learning quality has a lower error variance (i.e., more precise estimates of the state) in the heterogeneous case. Now consider an agent who is at the 25th percentile in terms of error variance (and thus estimates the state better than 75 percent of agents); the red points show that the advantage of the heterogeneous case becomes even more stark for these agents. In Figure 3(b), we pool all the agents together across all networks and depict the empirical distribution of error variance. In the homogeneous case (red histogram), there is bunching around the asymptotic variance for the homogeneous-signal case. When we introduce heterogeneity in signal quality (blue histogram), a substantial share of households have prediction variance below this boundary, thus benefiting from the heterogeneity. Overall, we see that even in networks with relatively small degree our qualitative results hold: adding heterogeneity helps learning in the population, even with a small group of agents with the new signal type.

5. THE IMPORTANCE OF ANTI-IMITATION

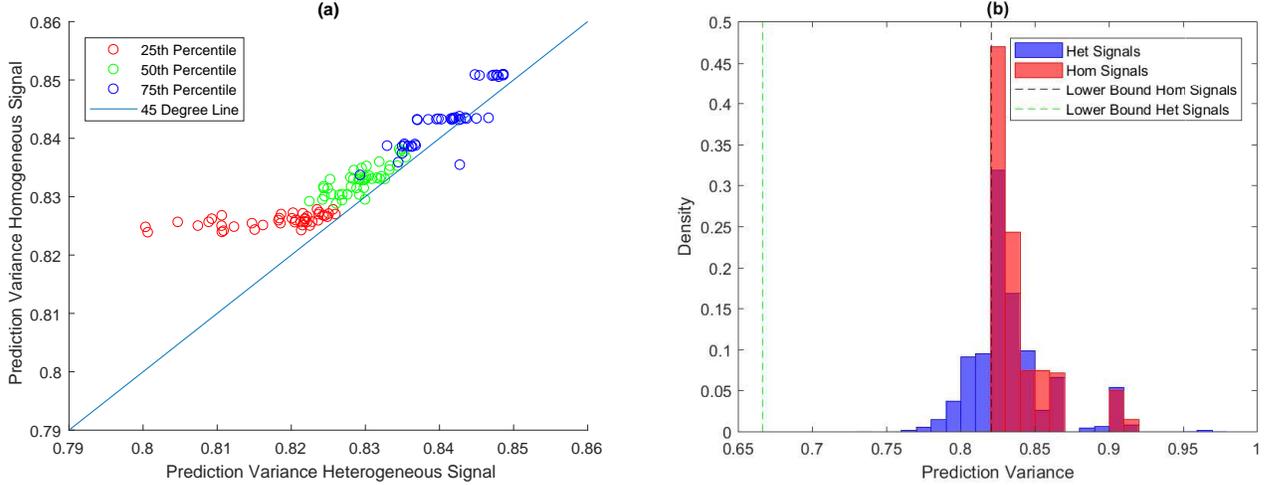
In the proof of our positive result on achieving the perfect aggregation benchmark (Theorem 1), a key aspect of the argument involved agents placing negative weights on some neighbors' estimates to solve their signal extraction problem. In this section, we demonstrate that this behavior, called *anti-imitation*,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

A natural reason for agents not to anti-imitate is that they do not account for correlations among their neighbors' estimates, conditional on the state. We first consider a particular model of naive agents who make this error and show such agents fall far short of perfect aggregation. We then formalize the idea that, more generally, anti-imitation is crucial to reaching the benchmark. This is done by demonstrating a general lack of asymptotic learning when all weights are positive, whether this is due to naiveté or some other reason. Finally, we show that even in fixed, finite networks, any positive weights chosen by optimizing agents will be Pareto-dominated.

5.1. Naive agents. As noted above, one reason agents might not use negative weights is that they misunderstand the distribution of the signals they are facing. In this part we

⁴⁰ We made the heterogeneous signals dependent on electricity status because we believe signal precision would in practice be correlated with, e.g., access to communication technology (or similar attributes). In the figure, we plot outcomes of the nodes with access to electricity—i.e., those whose signal variances did not change in our exercise.

FIGURE 3. Prediction Variance In Indian Villages



(a) The error variance of the agent in the 25th, 50th and 75th percentiles in each village, in the homogeneous and heterogeneous cases. (b) Histograms of error variance (we pool all the agents together across all networks) for the homogeneous (red) and heterogeneous (blue) case. Vertical lines show the asymptotic variance for the complete graph as $n \rightarrow \infty$ for the two cases.

introduce agents who do not process their observations of neighbors’ decisions in a fully Bayesian manner. We consider a particular form of misspecification that simplifies solving for equilibria analytically.⁴¹ 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 model. So naive agents understand that their neighbors’ actions from the previous period are estimates of θ_{t-1} . But they think each such estimate is independent given the state, and that the precision of the estimate is equal to the signal precision of the corresponding agent.

With heterogeneous signal qualities, naive learning outcomes depend on the network more than they do in the equilibrium we have studied with rational agents: 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.

⁴¹ 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, Banerjee, Chandrasekhar, Hanna, and Olken 2016 consider one such variant).

We will describe outcomes with two signal types, σ_A^2 and σ_B^2 .⁴² 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

$$(8) \quad V_A^\infty = \frac{\kappa_t^2 + \sigma_A^{-2}}{(1 + \sigma_A^{-2})^2}, \quad V_B^\infty = \frac{\kappa_t^2 + \sigma_B^{-2}}{(1 + \sigma_B^{-2})^2}$$

where

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

Naive agents' equilibrium variances converge to these values.

Proposition 4. *Let $\delta > 0$. Under the assumptions in this subsection:*

- (1) *There is a unique equilibrium on G_n .*
- (2) *Given any $\delta > 0$, asymptotically almost surely all agents' equilibrium variances are within δ of V_A^∞ and V_B^∞ .*
- (3) *There exists $\varepsilon > 0$ such that asymptotically almost surely the ε -perfect aggregation benchmark is not achieved, and when $\sigma_A^2 = \sigma_B^2$ asymptotically almost surely all agents' variances are larger than V^∞ .*

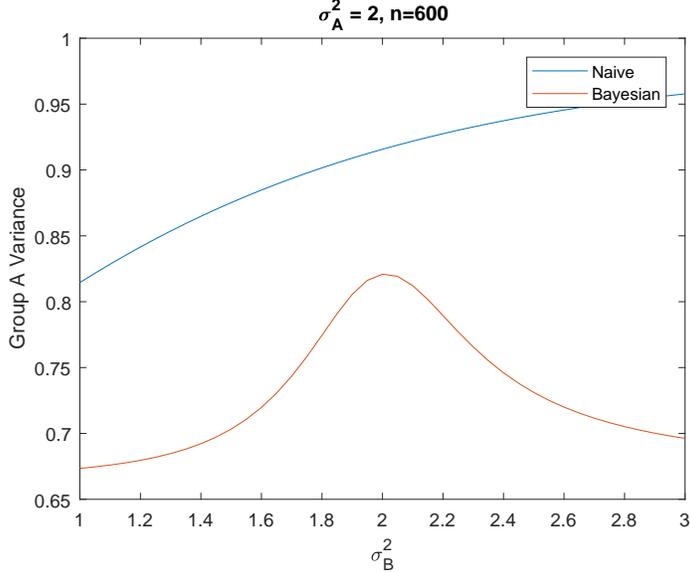
Aggregating information well requires a sophisticated response to the correlations in observed actions. Because naive agents completely ignore these correlations, their learning outcomes are poor. In particular their variances are larger than at the equilibria we discussed in the Bayesian case, even 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 $(G_n)_{n=1}^\infty$ of (deterministic) networks with the minimum degree diverging to ∞ and any sequence of equilibria, the equilibrium action variances of all agents converge to V_A^∞ .

In Figure 4, we compare Bayesian and naive learning outcomes. As in Figure 2, 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 $n \rightarrow \infty$.

⁴² The general case, with many signal types, is similar.

FIGURE 4. Bayesian and Naive Learning



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 weights are positive, we show that these equilibria are bounded away from perfect aggregation.

Proposition 5. *Consider a sequence of undirected networks $(G_n)_{n=1}^\infty$ with n agents in G_n and assume that all private signal variances are bounded below by $\underline{\sigma}^2 > 0$. Either*

(1) suppose all agents are Bayesian. Consider any sequence of equilibria on G_n in which all agents are using positive weights; or

(2) suppose all agents are naive, and consider any sequence of equilibria on G_n .

Then there is an $\varepsilon > 0$ 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 action 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. Related obstructions to learning play an important role in [Harel, Mossel, Strack, and Tamuz \(2017\)](#).

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 \mathbf{W} and \mathbf{w}^s is the (unique) covariance matrix \mathbf{V}^* such that if actions have a variance-covariance matrix given by $\mathbf{V}_t = \mathbf{V}^*$ and next-period actions are set using weights $(\mathbf{W}, \mathbf{w}^s)$, then $\mathbf{V}_{t+1} = \mathbf{V}^*$ 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 \mathbf{W} and \mathbf{w}^s and suppose they are all positive, with an associated steady state \mathbf{V}_t . Suppose either*

- (1) *there is an agent i whose weights are a Bayesian best response to \mathbf{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 \mathbf{V}_t , and who observes multiple neighbors.*

Then the steady state \mathbf{V}_t is Pareto-dominated by 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. We state the result under relatively weak hypotheses on behavior, but special cases include equilibria, where all agents are rational, or naive equilibria.

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.*

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 result of Theorem 2 (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 her 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 (as opposed to best-response) weights is monotonic in terms of the variances of guesses. 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 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 \hat{w}^s on private signals and weight $\frac{1}{2}(1 - \hat{w}^s)$ on each observed action.

When $\sigma^2 = 10$, the equilibrium weight is $\hat{w}^s = 0.192$ 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 $\hat{w}^s = 0.570$ 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 literature on social learning and learning in networks.⁴³

6.1. DeGroot learning and related network models. Play in the stationary linear equilibria of our model closely resembles behavior in the DeGroot (1974) heuristic, where agents update by linearly aggregating network neighbors' past estimates, with constant weights on neighbors over time.⁴⁴ We now discuss how our model compares to existing work on these kinds of models—both in terms of justifications and foundations, as well as in terms of what matters for good learning.

DeMarzo, Vayanos, and Zweibel (2003) justified the DeGroot heuristic by assuming that 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; then agents 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.

⁴³ For surveys of different parts of this literature, see, among others, Acemoglu and Ozdaglar (2011), Golub and Sadler (2016), and Mossel and Tamuz (2017).

⁴⁴ It is also closely related to the model of Friedkin and Johnsen (1997), who allow each individual to place weight on a persistent estimate of her own.

Recently, Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012) and Molavi, Tahbaz-Salehi, and Jadbabaie (2018) have offered powerful new analyses of these types of heuristics, and have introduced flexible forms suited to a state that changes 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.⁴⁵

Concerning learning outcomes under the DeGroot learning rule, DeMarzo, Vayanos, and Zweibel (2003) emphasized that in their model, the stationary rule could in general be far from optimal: 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) showed that DeGroot agents could nevertheless converge to precise estimates in large networks as long as no agent has too prominent a network position, and not otherwise. Different and less demanding sufficient conditions for good learning were given by Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi (2012) in a world with a fixed state but ongoing arrival of information, where agents communicate distributions and update them using an adapted DeGroot rule. An overall message that emerges from these papers is that in these fixed-state environments, certain simple heuristics that require no sophistication about correlations between neighbors' behavior, can perform quite well: they can allow agents to guess the state quite precisely. Of course, the findings in this literature are not uniformly optimistic for naive learning. Beyond the caveats already mentioned, recent work by Akbarpour, Saberi, and Shamel (2017) analyzed a version of the DeGroot model with a fixed state and overlapping generations, and found that a changing society can cause DeGroot learning to be very inefficient, because agents do not know how precise the guesses of various contacts are. Our findings highlight new obstructions to naive learning that arise from a changing state. To learn well in such a world, agents need a sophisticated response to the *correlation* in neighbors' estimates that arises from those neighbors' past learning; knowing the marginal distributions of precision is not enough.

Moreover, in contrast to the environment of DeMarzo, Vayanos, and Zweibel (2003), even Bayesian agents who understand the environment perfectly are not guaranteed to be able to aggregate information well (Section 4.2). Bayesians' good learning in our environment, and its failure, depend on conditions—namely, signal diversity throughout the network—that differ

⁴⁵ Indeed, agents behaving according to the DeGroot heuristic even when it is not appropriate might have to do with their experiences in stationary environments where it is closer to optimal.

markedly from the ones that play a role in the fixed-state environments discussed above. Thus, a changing state presents new forces relevant to the analysis of heuristic learning models, and also changes the relevant Bayesian benchmarks.

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 ($\rho = 1$).⁴⁶ 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 shows 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 weights in our model are determined in a Nash equilibrium, we can consider how they respond endogenously to changes in the environment (e.g., network). We also give conditions for good learning even when agents are optimizing for themselves, as opposed to being programmed to achieve 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 such aggregation processes. Our model provides foundations for structural estimation with Bayesian behavior as well as testing of the Bayesian model against behavioral alternatives such as that of [Alatas, Banerjee, Chandrasekhar, Hanna, and Olken \(2016\)](#); we discuss this below in Section 7.1.

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

⁴⁶ For an early model of repeated learning about a changing state based on social information, see [Ellison and Fudenberg \(1995\)](#); that model differs in that there is no persistence in the state over time.

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, Dahleh, Lobel, and Ozdaglar, 2011; Eyster and Rabin, 2014; Lobel and Sadler, 2015a,b). These models thus feature an incomplete network of observation opportunities.

A major concern of this literature is the potential for information aggregation to stop after some finite time due to inference problems. The discreteness of individuals’ actions often plays an important role. Our focus is different in that 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 for 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, Dahleh, Lobel, and Ozdaglar (2011), and Lobel and Sadler (2015a).⁴⁷ Another difference concerns the modeling of the network: Our agents are at the nodes of a finite, unchanging network, and there is bidirectional observation in the sense that node A learns from node B, which then learns from node A. Thus, new sorts of “feedback” considerations can emerge when we analyze the comparative statics of our model (see, e.g., Section 5.3).

Despite the modeling differences, the inefficiencies that drive information cascades are related in some ways to ones that play a role in our Proposition 2 and Theorem 2. When actions are correlated in certain ways, the information underlying those actions cannot be extracted efficiently. The correlation occurs, in turn, because agents weight their private signals differently from the socially optimal way. Because our learning weights are described by fixed points, we have new ways to examine structural features of the informational environment that affect the magnitude of the inefficiency, as illustrated by Corollary 1.

Another point of contact with the classical social learning literature 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. Nevertheless, this early study has a key conceptual connection to ours in stressing that in a

⁴⁷ Those conditions require either that some signals are very informative, or that some agents have access to a large number of samples of behavior that are not based on any common signals—neither of which we assume in our main results.

dynamic world, social learning outcomes can look very different from those of the canonical static models.

Finally, 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 an important feature in our dynamic model, and in our context is crucial for good learning. Despite this similarity, there is an important contrast between our findings and standard sequential models. In those models, while rational agents *do* prefer to anti-imitate, in many cases individuals 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 \(2011\)](#) and [Lobel and Sadler \(2015a\)](#). Our dynamic learning environment is different, as shown in [Proposition 5](#): to have any hope of approaching good aggregation benchmarks, agents must respond in a sophisticated way, with anti-imitation, to their neighbors’ (correlated) estimates.

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 easily 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 $\bar{a}_{i,t} = a_{i,t} + \xi_{i,t}$ 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).⁴⁸

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.

⁴⁸ We can instead assume that the analyst observes (a proxy for) the private signal $s_{i,t}$ of agent i ; we mention how below.

We can identify the weights agents are using through standard vector autoregression methods. In steady state,

$$(9) \quad \bar{a}_{i,t} = \sum_j W_{ij} \rho \bar{a}_{j,t-1} + w_i^s \theta_t + \zeta_{i,t},$$

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 $\bar{a}_{j,t-1}$ from the previous period.⁴⁹ 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, network links can generically be identified by testing whether the recovered social weights are nonzero. For such tests (and more generally) the standard errors in the estimators can be obtained by standard techniques.⁵⁰

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 determine the signal variances using the fact that equilibrium is characterized by $\Phi(\widehat{\mathbf{V}}) = \widehat{\mathbf{V}}$ 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:

$$(10) \quad \widehat{V}_{ii} = \sum_j \sum_k \widehat{W}_{ij} \widehat{W}_{ik} (\rho^2 \widehat{V}_{jk} + 1) + (\widehat{w}_i^s)^2 \sigma_i^2.$$

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

⁴⁹ This system defines a VAR(1) process (or generally VAR(m) for memory length m).

⁵⁰ 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, the techniques required will generally be more complicated.

agents are using some other behavioral rule (e.g., optimizing in a misspecified model) we can replace (10) by a suitable analogue that reflects the bounded rationality in agents' inference. If such a 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 5.1.

7.2. Signal acquisition. In this section we analyze what would happen if agents were to choose their precision $\sigma_{i,t}^{-2}$ in every period.⁵¹ We assume that there is a convex cost function in precisions $c(\sigma_{i,t}^{-2})$.

Corollary. *Let G be a complete network, and suppose each agent maximizes*

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

There is a unique equilibrium in which $\sigma_{i,t}^2 = \sigma_{j,t'}^2$ for all i, j, t and t' , so there exists $\varepsilon > 0$ such that the ε -perfect aggregation benchmark is not achieved at this equilibrium for any n .

The proof is simple—since the network is complete, all agents have the same social signal r_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 stationary linear equilibrium given signal precisions $\sigma_{i,t}^{-2}$. Because the optimal $\sigma_{i,t}^{-2}$ is increasing in the variance of the social signal $r_{i,t}$, there is a unique value of $\sigma_{i,t}^{-2}$ that is a best response in every period. Hence the equilibrium of the signal acquisition game is unique.

This result shows that while the result of non-diverse signals may seem like an knife-edge 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 statistic 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

⁵¹ For recent work on endogenous signal acquisition in social learning, see Mueller-Frank and Pai (2016) and Ali (2018).

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. Similarly, information need not be one-dimensional, or characterized only by its precision.

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 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.
- ACEMOGLU, D. AND A. OZDAGLAR (2011): "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, 1, 3–49.
- AKBARPOUR, M., A. SABERI, AND A. SHAMELI (2017): "Information Aggregation in Overlapping Generations," Available at SSRN, <https://ssrn.com/abstract=3035178>.
- 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.
- ALI, S. N. (2018): "Herding with costly information," *Journal of Economic Theory*, 175, 713–729.
- BABUS, A. AND P. KONDOR (2018): "Trading and information diffusion in over-the-counter markets," *Econometrica*, forthcoming.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): "The diffusion of microfinance," *Science*, 341, 1236498.
- BANERJEE, A. AND D. FUDENBERG (2004): "Word-of-mouth learning," *Games and economic behavior*, 46, 1–22.
- BANERJEE, A. V. (1992): "A simple model of herd behavior," *Quarterly Journal of Economics*, 107, 797–817.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 100, 992–1026.

- 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.
- ELLISON, G. AND D. FUDENBERG (1995): “Word-of-mouth communication and social learning,” *The Quarterly Journal of Economics*, 110, 93–125.
- 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.
- 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.
- GEANAKOPOLOS, 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): “Groupthink and the Failure of Information Aggregation in Large Groups,” *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.
- LAMBERT, N. S., M. OSTROVSKY, AND M. PANOV (2018): “Strategic trading in informationally complex environments,” *Econometrica*, 86, 1119–1157.
- 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.
- MALAMUD, S. AND M. ROSTEK (2017): “Decentralized exchange,” *American Economic Review*, 107, 3320–62.

- 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.
- MOSSEL, E., M. MUELLER-FRANK, A. SLY, AND O. TAMUZ (2018): “Social learning equilibria,” *ArXiv:1207.5895*.
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): “Strategic learning and the topology of social networks,” *Econometrica*, 83, 1755–1794.
- MOSSEL, E. AND O. TAMUZ (2017): “Opinion exchange dynamics,” *Probability Surveys*, 14, 155–204.
- MUELLER-FRANK, M. (2013): “A general framework for rational learning in social networks,” *Theoretical Economics*, 8, 1–40.
- MUELLER-FRANK, M. AND M. PAI (2016): “Social Learning with Costly Search,” *American Economic Journal: Microeconomics*, 8, 83–109.
- OLFATI-SABER, R. (2007): “Distributed Kalman filtering for sensor networks,” in *Decision and Control, 2007 46th IEEE Conference on*, IEEE, 5492–5498.
- OSTROVSKY, M. (2012): “Information aggregation in dynamic markets with strategic traders,” *Econometrica*, 80, 2595–2647.
- PINELIS, I. (2018): “Inverse of matrix with blocks of ones,” *MathOverflow*, URL:<https://mathoverflow.net/q/296933> (version: 2018-04-04).
- RAHIMIAN, M. A. AND A. JADBABAIE (2017): “Bayesian learning without recall,” *IEEE Transactions on Signal and Information Processing over Networks*.
- SETHI, R. AND M. YILDIZ (2012): “Public disagreement,” *American Economic Journal: Microeconomics*, 4, 57–95.
- (2016): “Communication with unknown perspectives,” *Econometrica*, 84, 2029–2069.
- 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*.
- VIVES, X. (1993): “How fast do rational agents learn?” *The Review of Economic Studies*, 60, 329–347.
- WEIZSÄCKER, G. (2010): “Do we follow others when we should? A simple test of rational expectations,” *The American Economic Review*, 100, 2340–2360.

APPENDIX A. DETAILS OF DEFINITIONS

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 the AR(1) process or, in the case of $\rho = 1$, of 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) 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 $\sigma_{(i,t)} : \mathbb{R} \times (\mathbb{R}^{|N(i)|})^m \rightarrow \mathbb{R}$, mapping her own signal and her neighborhood's actions, $(s_{i,t}, (\mathbf{a}_{N_i, t-\ell})_{\ell=1}^m)$, to a real-valued action $a_{i,t}$. We call the set of these functions $\tilde{\Sigma}_{(i,t)}$. Let $\tilde{\Sigma} = \prod_{(i,t) \in \mathcal{A}} \tilde{\Sigma}_{(i,t)}$ be the set of response profiles. We now define the set of (*unambiguous*) *strategy profiles*, $\Sigma \subset \tilde{\Sigma}$. We say that a response profile $\sigma \in \tilde{\Sigma}$ 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

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

2. Any two tuples 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 essentially 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.⁵² Note that if $\sigma \in \Sigma$, then $\tilde{\sigma} = (\sigma'_{(i,t)}, \sigma_{-(i,t)}) \in \Sigma$ whenever $\sigma'_{(i,t)} \in \tilde{\Sigma}_{(i,t)}$. 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

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

⁵² 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 necessary to rule out response profiles such as the one given by $\sigma_{i,t}(s_{i,t}, a_{i,t-1}) = a_{i,t-1}$, which have many satisfying action paths, leaving payoffs undetermined.

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 $\tilde{\sigma} \in \Sigma$ such that $\tilde{\sigma} = (\sigma'_{(i,t)}, \sigma_{-(i,t)})$ for some $\sigma'_{(i,t)} \in \Sigma_{(i,t)}$, we have

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

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 .⁵³ 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*.

APPENDIX B. EXISTENCE OF EQUILIBRIUM: PROOF OF PROPOSITION 1

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

We begin by defining the compact set \mathcal{K} . Because memory is arbitrary, entries of \mathbf{V}_t are covariances between pairs of neighbor actions from any periods available in memory. Let k, l be two indices of such actions, corresponding to actions taken at nodes 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 \mathbf{V}_t such that, for any such k, l ,

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

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

To show this claim, we will first find upper and lower bounds on the variance of any neighbor's action (at any period in memory). For the upper bound, 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, note that if she knew the previous period's state and her own signal, then

⁵³ I.e., $\sigma' = \mathfrak{T}_p \sigma$ is defined by $\sigma_{(i,t)} = \sigma_{(i,t-p)}$.

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 the node- i agent from m periods ago, the variance of the estimate of θ_{t-1} based on i 's action from m periods ago is at most $\rho^{m-1}\sigma_i^2 + \frac{1-\rho^{m-1}}{1-\rho}$ and at least $\frac{\rho^{m-1}}{1+\sigma_i^{-2}} + \frac{1-\rho^{m-1}}{1-\rho}$. 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.

We have now established that there is a variance-covariance matrix $\widehat{\mathbf{V}}$ such that $\Phi(\widehat{\mathbf{V}}) = \widehat{\mathbf{V}}$. By definition of Φ , this means there exists some weight profile (W, w^s) such that, when applied to prior actions that have variance-covariance matrix $\widehat{\mathbf{V}}$, produce variance-covariance matrix $\widehat{\mathbf{V}}$. However, it still remains to show that this is the variance-covariance matrix reached when agents have been using the weights (W, w^s) forever.

To show this, first observe that if agents have been using the weights (W, w^s) forever, the variance-covariance matrix \mathbf{V}_t in any period is uniquely determined and does not depend on t ; call this $\check{\mathbf{V}}$.⁵⁴ This is because actions can be expressed as linear combinations of private signals with coefficients depending only on the weights. Second, it follows from our construction above of the matrix $\widehat{\mathbf{V}}$ and the weights (W, w^s) that there is a distribution of actions where the variance-covariance matrix is $\widehat{\mathbf{V}}$ in every period and agents are using weights (W, w^s) in every period. Combining the two statements shows that in fact $\check{\mathbf{V}} = \widehat{\mathbf{V}}$, and this completes the proof.

⁵⁴ The variance-covariance matrices are well-defined because the (W, w^s) weights yield unambiguous strategy profiles in the sense of Appendix A.

APPENDIX C. PROOF OF THEOREM 1 (FOR ONLINE PUBLICATION)

C.1. Notation and key notions. Let \mathbb{S} be the (by assumption finite) set of all possible signal variances, and let $\bar{\sigma}^2$ be the largest of them. The proof will focus on the covariances of errors in social signals. Recall that both $r_{i,t}$ and $r_{j,t}$ have mean θ_{t-1} , because each is an unbiased estimate⁵⁵ 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 $(\text{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}))_{i,j}$ and let \mathcal{W} be the subset of such covariance matrices. For all i, j note that $\text{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}) \in [-\bar{\sigma}^2, \bar{\sigma}^2]$ using the Cauchy-Schwarz inequality and the fact that $\text{Var}(r_{i,t} - \theta_{t-1}) \in [0, \bar{\sigma}^2]$ 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 $\tilde{\Phi}$ on \mathcal{W} .

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

1. for any pair of distinct agents⁵⁶ $i \in G_n^k$ and $j \in G_n^{k'}$,

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

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 $i \in G_n^k$,

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

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 i 's and j 's 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.

C.2. Proof strategy.

C.2.1. A set $\mathcal{W}_{\bar{\delta}, \bar{\zeta}}$ of outcomes with good learning. Our goal is to show that as n grows large, $\text{Var}(r_{i,t} - \theta_{t-1})$ becomes very small, which then implies that the agents asymptotically learn. We will take $\bar{\delta}$ and $\bar{\zeta}$ to be arbitrarily small numbers and show that for large enough n , with high probability (which we abbreviate ‘‘asymptotically almost surely’’ or ‘‘a.a.s.’’) the

⁵⁵ This is because it is a linear combination, with coefficients summing to 1, of unbiased estimates of θ_{t-1} .

⁵⁶ Throughout this proof, we abuse terminology by referring to agents and nodes interchangeably when the relevant t is clear or specified nearby.

equilibrium outcome has a social error covariance matrix A_t in the set $\mathcal{W}_{\bar{\delta}, \bar{\zeta}}$. In particular, $\text{Var}(r_{i,t} - \theta_{t-1})$ becomes arbitrarily small in this limit. 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.

C.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.) $\bar{\mathcal{W}} \subset \mathcal{W}_{\bar{\delta}, \bar{\zeta}}$ and $\tilde{\Phi}(\bar{\mathcal{W}}) \subset \bar{\mathcal{W}}$. 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 $\bar{\mathcal{W}}$, we will fix a positive constant β (to be determined later), and define

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

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

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

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 $\bar{\mathcal{W}}$.

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

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 under⁵⁷ $\tilde{\Phi}^2$, i.e., $\tilde{\Phi}^2(\mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}$.*

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

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

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

C.2.3. *Proving the lemmas by analyzing how $\tilde{\Phi}$ and $\tilde{\Phi}^2$ act on sets $\mathcal{W}_{\delta, \zeta}$.* The lemmas are about how $\tilde{\Phi}$ and $\tilde{\Phi}^2$ 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.

⁵⁷ The notation $\tilde{\Phi}^2$ means the operator $\tilde{\Phi}$ applied twice.

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 $\tilde{\Phi}$ and $\tilde{\Phi}^2$, 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 agent (i, t') , 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: $\tilde{\Phi}(\mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n}, \frac{\lambda}{n}}$. 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: $\tilde{\Phi}^2(\mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}$. 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 $A_{t+2} \in \mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}$.

C.3. Carrying out the steps.

C.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 $X_{i,\tau k}^{(1)}$ 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:

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

(E_2) Let $X_{ii',\tau k}^{(2)}$ 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:

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

(E_3) Let $X_{i,\tau k,j}^{(3)}$ 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:

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

(E_4) Let $X_{ii',\tau k,j}^{(4)}$ 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:

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

(E_5) Let $X_{i,\tau k,jj'}^{(5)}$ 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:

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

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 .

C.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:

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

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 $\mathbf{C}_{i,t}^{-1}$. 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 of $f(A) = \mathbf{C}_{i,t}^{-1}$ 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 $\mathbf{C}_i(x)$ 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 $\mathbf{C}_i(x) - \mathbf{C}_i(x')$ 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 $\mathbf{C}_i(x)^{-1}$ are close to the column sums of $\mathbf{C}(0)_i^{-1}$. To do so, we will evaluate the formula

$$(11) \quad \frac{\partial f(A(x))}{\partial x} = \frac{\partial \mathbf{C}_i(x)^{-1}}{\partial x} = \mathbf{C}_i(x)^{-1} \frac{\partial \mathbf{C}_i(x)}{\partial x} \mathbf{C}_i(x)^{-1}$$

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 $\mathbf{C}_i(0)^{-1}$, 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

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

Therefore, $\mathbf{C}_i(0)$ can be written as a block matrix with blocks $a_{ij}\mathbf{1}_{r_i \times r_j} + b_i\delta_{ij}I_{r_i}$ 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.⁵⁸

Lemma 4 (Pinelis (2018)). *Let C be a matrix consisting of $S \times S$ blocks, with its (i, j) block given by*

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

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

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

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

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

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

⁵⁸ We are very grateful to Iosif Pinelis for suggesting the argument in the lemma.

for some real c_{ij} and d_i .

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

$$(12) \quad \sum_k (a_{ik} 1_{r_i \times r_k} + b_i \delta_{ik} I_{r_i}) (c_{kj} 1_{r_k \times r_j} + d_k \delta_{kj} I_{r_k}) = a_{ij} d_j + \sum_k (a_{ik} r_k + \delta_{ik} b_k) c_{kj} 1_{r_i \times r_j} + b_i d_i \delta_{ij} I_{r_i}.$$

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 display (12) 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 $br := (b_1 r_1, \dots, b_q r_q)$. Therefore as $n \rightarrow \infty$ the off-diagonal blocks will be $O(1/n^2)$ while the diagonal blocks are

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

as desired. □

Using Lemma 4 we can analyze the column sums of⁵⁹

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

In more detail, we use the formula of the lemma to estimate both copies of $\mathbf{C}_i(0)^{-1}$, and then expand this to write an expression for any column sum of $\mathbf{C}_i(0)^{-1} M \mathbf{C}_i(0)^{-1}$. It follows straightforwardly from this calculation that all these 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 11 repeatedly in x , we obtain an expression for the k^{th} derivative in

⁵⁹ Recall we wrote $A(x) = A(0) + xM$, and in (11) we expressed the derivative of f in x in terms of the matrix we exhibit here.

terms of $\mathbf{C}_i(0)^{-1}$ and M :

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

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.

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

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

By the first part of the lemma, the ratio between any two weights (both of the form $W_{ij,t+1}$, $w_{i,t+1}^s$, or $w_{j,t}^s$) 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 $\text{Var}(r_{i,t+1} - \theta_t)$ 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 $\text{Var}(r_{i,t+1} - \theta_t) < \beta/n = \delta$ (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 .⁶⁰ 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 ,

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

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

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

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.

C.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.

⁶⁰ 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

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

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 $\tilde{\Phi}$ twice. This gives a formula that can be used for bounding the covariances⁶¹ 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. (Recall the notation introduced in Section C.1 above.) 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): We calculate:

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

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

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

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 $[1 - 2\gamma/n, 1 + 2\gamma/n]$.*

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

⁶¹ We take this term to refer to variances, as well.

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

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

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 $\text{Cov}(r_{i,t+2} - \theta_{t+1}, r_{i',t+2} - \theta_{t+1})$ 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 ν_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 ν_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 $\text{Cov}(r_{i,t+2} - \theta_{t+1}, r_{i',t+2} - \theta_{t+1})$ 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

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

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 :

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

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

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

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

For each j in group A ,

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

where we recall that $\kappa_{j,t+1}^2 = \text{Var}(r_{j,t+1} - \theta_t)$. Because $\kappa_{j,t+1}$ is close to zero, we can approximate $w_{j,t+1}^s$ 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

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

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 $\mathbf{V}_{j',j'',t}$ 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 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

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

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

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

As in the ν_{t+1} case, we can approximate $(1-w_{j',t}^s)$ as a linear function of $\kappa_{j',t}$ up to $O(1/n^2)$ terms. Because the number of paths to each agent j' through 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 $\text{Var}(r_{i,t+2} - \theta_{t+1})$ of each social signal about θ_{t+1} is at most δ . The proof is the same as in Step 2(b).

APPENDIX D. REMAINING PROOFS (FOR ONLINE PUBLICATION)

D.1. 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 $\tilde{\Phi}$ on the space of variances $\kappa_{i,t}^2$, which we denote $\tilde{\mathcal{V}}$. 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 $\mathbf{V}_{ii,t}$ of agent i 's action, as well as the covariances $\mathbf{V}_{ii',t}$ 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 $\tilde{\Phi}$ is a contraction on $\tilde{\mathcal{V}}$ 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

$$(13) \quad a_{i',t} = \frac{(\rho^2 \kappa_{i,t}^2 + 1)^{-1}}{\sigma_{i'}^{-2} + (\rho^2 \kappa_{i,t}^2 + 1)^{-1}} \cdot r_{i,t} + \frac{\sigma_{i'}^{-2}}{\sigma_{i'}^{-2} + (\rho^2 \kappa_{i,t}^2 + 1)^{-1}} \cdot s_{i',t}$$

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

$$(14) \quad a_{i',t} = b_0 \cdot r_{i,t} + b_{i'} \cdot s_{i',t}$$

(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 $\kappa_{j,t}^2 = \text{Var}(r_{j,t} - \theta_t)$ depends on $\kappa_{i,t-1}^2 = \text{Var}(r_{i,t-1} - \theta_{t-1})$. 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

$$(15) \quad r_{j,t} = c_0 \cdot (\rho r_{i,t-1}) + \sum_{i'} c_{i'} s_{i',t}$$

and therefore (taking variances of both sides)

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

The desired result, that $\tilde{\Phi}$ 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 (14) and (15), we compute that

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

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 (13) 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 $\tilde{\Phi}$ (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 2. By Lemma 7 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 $1 + \rho$, for ε sufficiently small the social signal $r_{i,t}$ is bounded away from a perfect estimate of θ_{t-1} . This gives a contradiction. \square

D.2. 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.

D.3. Proof of Proposition 3. We outline the argument. In Step 1, we construct a symmetric version of the Erdos-Renyi network and show there exists a symmetric equilibrium $\widehat{\mathbf{V}}^{sym}(n)$ on this symmetric network. In Step 2, we show variances and covariances at the equilibrium $\widehat{\mathbf{V}}^{sym}(n)$ converge to V^∞ and Cov^∞ . The remainder of the proof shows there is an equilibrium on G_n near $\widehat{\mathbf{V}}^{sym}(n)$. Step 3 defines a *no-large-deviations* event depending on the realized Erdos-Renyi network, and we condition on this event. Step 4 shows that Φ maps a small neighborhood of $\widehat{\mathbf{V}}^{sym}(n)$ to itself. Finally, in Step 5 we apply the Brouwer fixed point theorem to conclude there exists an equilibrium on G_n in this neighborhood.

Step 1: We first consider a symmetric and deterministic version G_n^{sym} of the network G_n on which all agents observe exactly pn other agents and any pair of agents commonly observes exactly p^2n other agents.

Let $\mathcal{V}^{sym} \subset \mathcal{V}$ be the space of covariance matrices for which each entry $\mathbf{V}(n)_{ij}$ depends only on whether i and j are equal and not on the particular agents. Even if such a G_n^{sym} network does not exist (for combinatorial reasons), updating as if on such a network induces a well-defined map $\Phi^{sym} : \mathcal{V}^{sym} \rightarrow \mathcal{V}^{sym}$. This map Φ^{sym} must have a fixed point, which we call $\widehat{\mathbf{V}}^{sym}(n)$. We will next show that the variances and covariances at $\widehat{\mathbf{V}}^{sym}(n)$ converge to V^∞ and Cov^∞ . The remainder of the proof will show that for n large enough, there exists an equilibrium $\widehat{\mathbf{V}}(n)$ on G_n close to the equilibrium $\widehat{\mathbf{V}}^{sym}(n)$ on G_n^{sym} .

Step 2: At $\widehat{\mathbf{V}}^{sym}(n)$, each agent's social signal is:

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

So the variance of the social signal about θ_t is

$$\kappa_{i,t} = \frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{11,t} + 1)}{pn} + \frac{(pn - 1)(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)}{pn}.$$

Thus the covariance of any two distinct agents solves

$$\widehat{\mathbf{V}}^{sym}(n)_{12,t} = \frac{\kappa_{i,t}^{-2}}{(\sigma^{-2} + \kappa_{i,t}^{-1})^2} \left(\frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{11,t} + 1)}{p^2 n} + \frac{(p^2 n - 1)(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)}{p^2 n} \right).$$

As $n \rightarrow \infty$, the right-hand side approaches

$$\frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)^{-1}}{[\sigma^{-2} + (\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)^{-1}]^2},$$

and the unique real solution to this equation is Cov^∞ . Computing $\widehat{\mathbf{V}}^{sym}(n)_{11,t}$ in terms of $\widehat{\mathbf{V}}^{sym}(n)_{12,t}$, we also see the variances converge to V^∞ .

Step 3: We will show that when $\zeta = \frac{1}{n}$, the updating map Φ on the network G_n maps a small neighborhood around $\widehat{\mathbf{V}}^{sym}(n)$ to itself. Let $\mathcal{V}_n \subset \mathcal{V}$ be the subset of covariance matrices such that

$$\mathbf{V}(n)_{ij} \in [\widehat{\mathbf{V}}^{sym}(n) - \zeta, \widehat{\mathbf{V}}^{sym}(n) + \zeta]$$

for all i and j . We will show in Steps 3 and 4 that $\Phi(\mathcal{V}_n) \subset \mathcal{V}_n$ for n large enough.

We first show that the network is close to symmetric with high probability. We will consider the event $E = E_1 \cap E_2$, where the E_i are defined by:

(E_1) : The degree of each agent i is between $1 - \zeta^2$ times its expected value and $1 + \zeta^2$ times its expected value, i.e., in $[(1 - \frac{1}{n^2})pn/2, (1 + \frac{1}{n^2})pn]$.

(E_2) : For any two agents i and i' , the number of agents observed by both i and i' between $1 - \zeta^2$ times its expected value and $1 + \zeta^2$ times its expected value, i.e., in $[(1 - \frac{1}{n^2})p^2 n/2, (1 + \frac{1}{n^2})p^2 n]$.

We can show as in the proof of Theorem 1 that the probability of the complement of event E vanishes exponentially in n . We will condition on the event E , which occurs with probability converging to 1, for the remainder of the proof.

Step 4: Assume all agents observe period t actions with covariances in \mathcal{V}_n and then act optimally in period $t+1$. We can show as in the proof of Lemma 3 that there exists a constant γ such any agent's weight $W_{ij,t+1}$ on an observed neighbor is in $[(1 - \gamma/n)\frac{1}{n}, (1 + \gamma/n)\frac{1}{n}]$. The relevant matrix $\mathbf{C}_i(0)$ now has only one block because we have only signal type, so the calculation is in fact simpler.

We have

$$r_{i,t+1} = \sum_j W_{ij,t+1} a_{j,t},$$

and therefore for i and i' distinct,

$$\text{Cov}(r_{i,t+1} - \theta_{t+1}, r_{i',t+1} - \theta_{t+1}) = \sum_{j,j'} W_{ij,t+1} W_{i'j',t+1} (\rho^2 \mathbf{V}_{jj',t} + 1).$$

The terms $W_{ij,t+1} W_{i'j',t+1}$ sum to 1, and each non-zero term is contained in $[\frac{(1-\gamma/n)^2}{n^2}, \frac{(1+\gamma/n)^2}{n^2}]$. The terms $\mathbf{V}_{jj',t}$ are each contained in $[\widehat{\mathbf{V}}^{sym}(n)_{12} - \frac{1}{n}, \widehat{\mathbf{V}}^{sym}(n)_{12} + \frac{1}{n}]$ (for j and j' distinct) and the terms $\mathbf{V}_{jj,t}$ are each contained in $[\widehat{\mathbf{V}}^{sym}(n)_{11} - \frac{1}{n}, \widehat{\mathbf{V}}^{sym}(n)_{11} + \frac{1}{n}]$. So

$$\left| \text{Cov}(r_{i,t+1} - \theta_{t+1}, r_{i',t+1} - \theta_{t+1}) - \frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{11,t} + 1)}{p^2 n} - \frac{(p^2 n - 1)(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)}{p^2 n} \right| \leq \frac{\rho^2}{n} + O\left(\frac{1}{n^2}\right),$$

where the terms of order $\frac{1}{n^2}$ come from variation in weights and variation in the network. The term

$$\frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{11,t} + 1)}{p^2 n} + \frac{(p^2 n - 1)(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)}{pn}$$

is the covariance of two distinct social signals in G_n^{sym} .

Similarly

$$\left| \text{Var}(r_{i,t+1} - \theta_{t+1}, r_{i',t+1} - \theta_{t+1}) - \frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{11,t} + 1)}{pn} - \frac{(pn - 1)(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)}{pn} \right| \leq \frac{\rho^2}{n} + O\left(\frac{1}{n^2}\right).$$

The term

$$\frac{(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{11,t} + 1)}{pn} + \frac{(pn - 1)(\rho^2 \widehat{\mathbf{V}}^{sym}(n)_{12,t} + 1)}{pn}$$

is the covariance of two distinct social signals in G_n^{sym} .

We compute from these inequalities that the variances $\mathbf{V}_{ii,t+1}$ and covariances $\mathbf{V}_{i'i',t+1}$ of actions are within $\frac{1}{n}$ of $\widehat{\mathbf{V}}^{sym}(n)_{11}$ and $\widehat{\mathbf{V}}^{sym}(n)_{12}$, respectively. This shows that $\Phi(\mathcal{V}_n) \subset \mathcal{V}_n$.

Step 5: By the Brouwer fixed point theorem, there exists an equilibrium $\widehat{\mathbf{V}}(n)$ on G_n with the desired properties. Because $V^\infty > (1 + \sigma^{-2})^{-1}$, there exists $\varepsilon > 0$ such that the ε -perfect aggregation benchmark is not achieved at this equilibrium for any n .

D.4. 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:

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

The weights $W_{ik,t}$ and $w_{i,t}^s$ are now all positive constants that do not depend on \mathbf{V}_{t-1} . So differentiating this formula, we find that all partial derivatives are bounded above by $1 - w_{i,t}^s < 1$. 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 $\widehat{\mathbf{V}}^{sym}$ where covariances depend only on signal types, i.e., $\widehat{\mathbf{V}}^{sym}$ 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 $n \rightarrow \infty$. 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 $r_{i,t}^{naive}$ of θ_{t-1} is approximately:

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

The actual variance of this estimate converges to:

$$(16) \quad \text{Var}(r_{i,t}^{naive} - \theta_t) = \frac{\rho^2}{(\sigma_A^{-2} + \sigma_B^{-2})} [\sigma_A^{-4} Cov_{AA}^\infty + \sigma_B^{-4} Cov_{BB}^\infty + 2\sigma_A^{-2}\sigma_B^{-2} Cov_{AB}^\infty] + 1$$

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:

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

Plugging into 16 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{V}^{sym} match equations 8.

We must check there is an equilibrium near \widehat{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}_ε be the ε -ball around in \widehat{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 $\Phi^{naive}(\widehat{V}^{sym}) = \widehat{V}^{sym}$ 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}_ε 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 $Cov^\infty > 0$.

D.5. 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 \bar{W} by $\bar{W}_{ij} = W_{ij}$ and $\bar{w}_i^s = w_i^s$ for all $i \neq k$, $\bar{W}_{kj} = (1 - \epsilon)W_{kj}$ for all $j \leq n$, and $\bar{w}_k^s = (1 - \epsilon)w_k^s + \epsilon$, 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 \bar{W} . These weights give an alternate updating function $\bar{\Phi}$ on the space of covariance matrices. Because the weights \bar{W} are positive and fixed, all coordinates of $\bar{\Phi}$ are increasing, linear functions of all previous period variances and covariances. Explicitly, the diagonal terms are

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

and the off-diagonal terms are

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

So it is sufficient to show the variances $\bar{\Phi}^h(\mathbf{V}_t)$ after applying $\bar{\Phi}$ for h periods Pareto dominate the variances in \mathbf{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 $[\bar{\Phi}(\mathbf{V})]_{ii}$ 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 $\bar{\Phi}$ 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 \bar{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).

D.6. 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 $\kappa_{i,t}^2 = \text{Var}(r_{i,t} - \theta_t)$ 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 $\frac{2}{2+\sigma^2}$ on observed actions.

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

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

for any *positive* coefficients b_j summing to 1, the variance of $r_{i,t} - \theta_{t-1}$ is at least $\frac{4\rho^2}{(2+\sigma^2)^2}$. 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.

APPENDIX E. SOCIALLY OPTIMAL LEARNING OUTCOMES WITH NON-DIVERSE SIGNALS (FOR ONLINE PUBLICATION)

In this section, we show that a social planner can achieve asymptotically perfect aggregation even when signals are non-diverse. Thus, the failure to achieve perfect aggregation at equilibrium with non-diverse signals is a consequence of individual incentives rather than a necessary feature of the environment.

Let G_n be the complete network with n agents. Suppose that $\sigma_i^2 = \sigma^2$ for all i and $m = 1$.

Proposition 6. *Let $\varepsilon > 0$. Under the assumptions in this section, for n sufficiently large there exist weights \mathbf{W} and \mathbf{w}^s such that at the corresponding steady state on G_n , the ε -perfect aggregation benchmark is achieved.*

Proof. An agent with a social signal equal to θ_{t-1} would place weight $\frac{\sigma^{-2}}{\sigma^{-2}+1}$ on her private signal and weight $\frac{1}{\sigma^{-2}+1}$ on her social signal. Let $w_A^s = \frac{\sigma^{-2}}{\sigma^{-2}+1} + \delta$ and $w_B^s = \frac{\sigma^{-2}}{\sigma^{-2}+1} - \delta$, where we will take $\delta > 0$ to be small.

Assume that the first $\lfloor n/2 \rfloor$ agents place weight w_A^s on their private signals and weight $1 - w_A^s$ on a common social signal r_t we will define, while the remaining agents place weight w_B^s on their private signals and weight $1 - w_B^s$ on the social signal r_t . As in the proof of Theorem 2,

$$\begin{aligned} \frac{1}{\lfloor n/2 \rfloor} \sum_{j=1}^{\lfloor n/2 \rfloor} a_{j,t-1} &= w_A^s \theta_{t-1} + (1 - w_A^s) r_{t-1} + O(n^{-1/2}), \\ \frac{1}{\lfloor n/2 \rfloor} \sum_{j=\lfloor n/2 \rfloor+1}^n a_{j,t-1} &= w_B^s \theta_{t-1} + (1 - w_B^s) r_{t-1} + O(n^{-1/2}). \end{aligned}$$

There is a linear combination of these summations equal to $\theta_{t-1} + O(n^{-1/2})$, and we can take r_t equal to this linear combination. Taking δ sufficiently small and then n sufficiently large, we find that ε -perfect aggregation is achieved. \square

In Figure 5, we conduct the same exercise as in Figure 2 with $n = 600$. The difference is that we now also add the prediction variance of group A when a social planner minimizes the total prediction variance (of both groups). The weights that each agent puts on her own private signal and the other agents are set to depend only on the groups. Under these socially optimal weights agents learn very well, and heterogeneity in signal variances only has a small impact.

FIGURE 5. Social Planner and Bayesian Learning

