

# LEARNING FROM NEIGHBORS ABOUT A CHANGING STATE

KRISHNA DASARATHA, BENJAMIN GOLUB, AND NIR HAK

ABSTRACT. Agents learn about a changing state using private signals and past actions of neighbors in a network. When can they learn efficiently about recent changes? We find two conditions are sufficient: (i) each individual’s neighbors have sufficiently diverse types of private information; (ii) agents are sophisticated enough to use this diversity to filter out outdated information and identify recent developments. If either condition fails, learning can be bounded far from efficient levels—even in networks where, with a fixed state, learning is guaranteed to be efficient without (i) or (ii). We thus identify learning externalities that are distinctive to a dynamic environment, and argue that they can be quite severe. The model we develop to make our argument provides a Bayesian foundation for DeGroot learning in networks, permitting new counterfactual and welfare analyses for that commonly-used behavioral model.

---

*Date:* June 30, 2019.

Department of Economics, Harvard University, Cambridge, U.S.A., [dasarath@g.harvard.edu](mailto:dasarath@g.harvard.edu), [bgolub@fas.harvard.edu](mailto:bgolub@fas.harvard.edu), [nirhak@g.harvard.edu](mailto:nirhak@g.harvard.edu). We thank the Foundations for Human Behavior Initiative at Harvard for financial support. We are grateful to (in random order) Alireza Tahbaz-Salehi, Drew Fudenberg, Matt Elliott, Nageeb Ali, Tomasz Strzalecki, Jeroen Swinkels, Margaret Meyer, Philipp Strack, Erik Eyster, Michael Powell, Michael Ostrovsky, Leeat Yariv, Andrea Galeotti, Eric Maskin, Elliot Lipnowski, Jeff Ely, Kevin He, Eddie Dekel, Annie Liang, Iosif Pinelis, Ozan Candogan, Ariel Rubinstein, Bob Wilson, Omer Tamuz, Jeff Zwiebel, Matthew O. Jackson, Matthew Rabin, and Andrzej Skrzypacz for valuable comments, and especially to Evan Sadler for important conversations early in the project, David Hirshleifer for detailed comments on a draft, and Rohan Dutta for a helpful discussion. We also thank numerous seminar participants for helpful questions and comments.

## 1. INTRODUCTION

Consider workers who are trying to assess conditions in a labor market, which naturally vary over time (e.g., due to demand shocks). The workers learn from their own experiences, as well as from friends' and coworkers' actions (e.g., search efforts), which reflect those people's estimates of market conditions. Consider another example: analysts who estimate changing economic variables, such as prices and interest rates varying over the business cycle, have their own data and calculations, but they also use the reports of some others. Because of the costs of acquiring and processing data, in general analysts may not access all relevant estimates generated by others.<sup>1</sup> This paper is a study of how individuals situated in a network, each seeking to learn a changing state of interest, incorporate the information of their neighbors, and whether that leads to timely aggregation of information.

We emphasize three features of the type of setting we are interested in. First, the state of interest is dynamic. Second, individuals learn both from private signals and from others' observed actions or reported estimates. Third, because different agents learn from different sources (i.e., without a central information aggregation device), estimates may differ across a population. A central 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? How does this depend on the way information is distributed throughout society? How much better off is a society of rational updaters compared to more naive ones? These questions are important for understanding the determinants of information aggregation and the welfare implications. They are also relevant in design decisions—e.g., for a planner who can influence group composition or information endowments with the aim of facilitating better learning.

Our first contribution, which is methodological, is to define and study equilibria in a dynamic environment capturing the key features emphasized above. The state evolves according to an AR(1) process, with an exogenous innovation added each period. There is therefore some degree of persistence in the state, but also a constant arrival of innovations to learn about. Agents receive private signals about the current state, with different distributions. Thus, some agents' private signals may exogenously be more precise. There is also endogenous learning: Agents' actions reveal their expectations of the current state; others learn from the actions they have observed before acting. These actions are relevant because they contain information about recent states. A crucial feature of our setup is that the environment in which agents learn is stationary—the distribution of states and signals

---

<sup>1</sup>See [Sethi and Yildiz \(2016\)](#) for a recent study of limits to information flow in networks due to the investments needed to understand sources.

is invariant to time shifts. This removes a time-dependence inherent in many models of network learning, where society accumulates information over time about a fixed state: In those models, the optimal way for agents to update depends on how much time has elapsed since the start of the learning process, and it is often an eventual *rate of learning* about a fixed state that is of interest (Harel et al., 2017). In our model, the state and the population are refreshed over time, and the predictions concern a *steady-state measure of learning quality*. This provides a complementary way to analyze how efficiently a community aggregates information, which is well-suited to learning about a dynamic state.

Importantly, this approach allows us to give a new foundation for a commonly used type of learning rule, the DeGroot model; this is our second contribution. In our model, Bayesians in stationary equilibrium update estimates by taking linear combinations of neighbors' estimates and their own private information, with the same coefficients over time. The weights are endogenously determined because, when each agent extracts information from neighbors' estimates, the information content of those estimates depends on the neighbors' learning rules. Bayesians who understand the environment perfectly will correctly account for all correlations in their observations. But, within the same environment, we can also study agents who ignore or underestimate the extent to which their neighbors' views are correlated by shared earlier observations, or exhibit other behavioral deviations from Bayesian rationality. This gives a new framework for comparative statics, welfare analysis, counterfactuals, and estimation exercises in social learning settings, under various behavioral assumptions.

Turning from the methods to substantive findings, our third contribution is to analyze the steady-state quality of learning, and how it depends on signal endowments, network structure, and agents' rationality or sophistication. First, we focus on a classical benchmark of Bayesian agents in equilibrium, who correctly understand each other's learning rules and respond optimally. The main finding is that, in large populations, an essentially optimal benchmark is achieved in an equilibrium, as long as each individual has access to a set of neighbors that is sufficiently diverse (in the sense of having different signal distributions from each other). If signal endowments are *not* diverse, then this result does not hold. Indeed, equilibrium learning may be necessarily bounded away from efficiency 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. A key mechanism behind the value of diverse signal endowments is that it leads to diversity of neighbors' behavior, which helps agents form better statistical estimators. These estimators

allow them to concentrate on new developments in the dynamic environment—in other words, to successfully *filter* out older, less useful information. Thus, in stark contrast with the analogous model with an unchanging state, Bayesian agents who understand the environment perfectly are not guaranteed to be able to aggregate information well (see Sections 4.4 and 4.5). We first make the points we have been discussing in the simplest way we can—with one-dimensional variation in signals—but then remark in Sections 7.3 and 7.4 that the logic extends to much more general environments.

To achieve good learning when it is possible, agents must have a sophisticated understanding of the correlations in their neighbors’ estimates. Thus, the second contrast we emphasize is between Bayesians who are correctly specified about their learning environment and agents who do not have sufficient sophistication about the correlations among their observations to remove confounds. To make the point that such understanding is essential to good aggregation, we examine some canonical types of learning rules, adapted to our setting, in which the sophistication is absent (Eyster and Rabin, 2010). There, information aggregation is essentially guaranteed to fall short of good aggregation benchmarks for all agents. We then generalize these results to a broader class of learning rules in which agents do not use sophisticated filtering techniques.

We argue that the deficiencies of naive learning rules differ from and more severe than those in similar problems with an unchanging state, and the requirements for good learning more stringent. In analogous fixed-state environments where individuals have sufficiently many observations, if everyone uses certain simple and stationary heuristics (requiring no sophistication about correlations between neighbors’ behavior), they can learn the state quite precisely. This point is made in the most closely analogous environments by Golub and Jackson (2010) and Jadbabaie et al. (2012), where DeGroot-style rules guarantee asymptotically efficient learning. A changing state makes such imitative heuristics quite inefficient. It is worth emphasizing that even Bayesians’ good learning in our environment depends on conditions—namely, signal diversity throughout the network—that differ markedly from the conditions that play a role in the above-mentioned models with an unchanging state.

Many of our results on the quality of learning use large random graphs. We thus also offer evidence that neither the technical conditions imposed these graphs nor a reliance on very large numbers is driving our principal conclusions. To this end, we first prove that even in small networks, learning rules without filtering (in a sense we make precise in Section 5.3) lead to Pareto-inefficient aggregation. To show that the importance of signal diversity manifests in realistic networks, we perform a numerical exercise: we calculate equilibria for

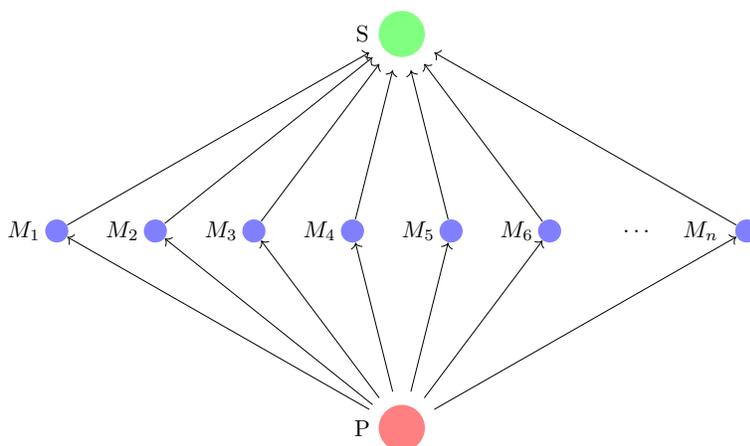
real-world social networks both for diverse and non-diverse signal endowments, and show that signal diversity does enable considerably better learning, even in networks with several hundred agents, each having 10-20 neighbors.

Finally, we discuss some implications of our results for designers who wish to facilitate better learning, and what distributions of expertise they would prefer. In particular, our results provide a distinctive rationale for *informational specialization* in organizations, which we flesh out in Section 7.3.

**An Example.** We now present a simple example that illustrates our dynamic model, highlights obstacles to learning that distinctively arise in a dynamic environment, and gives a sense of some of the main forces that play a role in our results on the quality of learning.

Consider an environment with a single perfectly informed source  $S$ ; many media outlets  $M_1, \dots, M_n$  with access to the source as well as some independent private information; and the general public. The public consists of many individuals who learn only from the media outlets. We are interested in how well each member of the public could learn by following many media outlets. More precisely, we consider the example shown in Figure 1.1 and think of  $P$  as a generic member of the large public.

FIGURE 1.1. The network used in the value of diversity example



The state  $\theta_t$  follows a Gaussian random walk:  $\theta_t = \theta_{t-1} + \nu_t$ , where the innovations  $\nu_t$  are standard normal. Each period, the source learns the state  $\theta_t$  and announces this to the media outlets. The media outlets observe the source's announcement from the previous period, which is  $\theta_{t-1}$ . At each time period, they also receive noisy private signals,  $s_{M_i,t} = \theta_t + \eta_{M_i,t}$  with normally distributed, independent, mean-zero errors  $\eta_{M_i,t}$ . They then

announce their posterior means of  $\theta_t$ , which we denote by  $a_{M_i,t}$ . The member of the public, in a given period  $t$ , makes an estimate based on the observations  $a_{M_1,t-1}, \dots, a_{M_n,t-1}$  of media outlets' actions in the previous period. All agents are short-lived: they see actions in their neighborhoods one period ago, and then they take an action that reveals their posterior belief of the state.

If we instead had a fixed state but the same signals and observation structure, learning would trivially be perfect: the media outlets learn the state from the source and report it to the public.

Now consider the dynamic environment. Given that  $P$  has no signal, an upper bound on what she can hope for is to learn  $\theta_{t-1}$  (and use that as her estimate of  $\theta_t$ ). Can this benchmark be achieved, and if so, when?

A typical estimate of a media outlet at time  $t$  is a linear combination of  $s_{M_i,t}$  and  $\theta_{t-1}$  (the latter being the social information that the media outlets learned from the source). In particular, the estimate can be expressed as

$$a_{M_i,t} = w_i s_{M_i,t} + (1 - w_i) \theta_{t-1},$$

where the weight  $w_i$  on the media outlet's signal is increasing in the precision of that signal. We give the public no private signal, for simplicity only.

Suppose first that the media outlets have identically distributed private signals. Because the member of the public observes many symmetric media outlets, it turns out that her best estimate of the state,  $a_{P,t}$ , is simply the average of the estimates of the media outlets. Since each of these outlets uses the same weight  $w_i = w$  on its private signal, we may write

$$a_{P,t} = w \sum_{i=1}^n \frac{s_{M_i,t-1}}{n} + (1 - w) \theta_{t-2} \approx w \theta_{t-1} + (1 - w) \theta_{t-2}.$$

That is,  $P$ 's estimate is an average of media private signals from last period, combined with what the media learned from the source, which tracks the state in the period before that. In the approximate equality, we have used the fact that an average of many private signals is approximately equal to the state, by our assumption of independent errors. No matter how many media outlets there are, and even though each has independent information about  $\theta_{t-1}$ , the public's beliefs are confounded by older information.

What if, instead, half of the media outlets (say  $M_1, \dots, M_{n/2}$ ) have more precise private signals than the other half, perhaps because these outlets have invested more heavily in covering this topic? The media outlets with more precise signals will then place weight  $w_A$  on their private signals, while the media outlets with less precise signals use a smaller weight  $w^B$ . We will now argue that a member of the public can extract more information

from the media in this setting. In particular, she can first compute the averages of the two groups' actions

$$w_A \sum_{i=1}^{n/2} \frac{s_{M_i,t-1}}{n/2} + (1 - w_A)\theta_{t-2} \approx w_A\theta_{t-1} + (1 - w_A)\theta_{t-2}$$

$$w_B \sum_{i=n/2+1}^n \frac{s_{M_i,t-1}}{n/2} + (1 - w_B)\theta_{t-2} \approx w_B\theta_{t-1} + (1 - w_B)\theta_{t-2}.$$

Then, since  $w_A > w_B$ , the public knows two distinct linear combinations of  $\theta_{t-1}$  and  $\theta_{t-2}$ . The state  $\theta_{t-1}$  is identified from these. So the member of the public can form a very precise estimate of  $\theta_{t-1}$ —which, recall, is as well as she can hope to do. The key force is that the two groups of media outlets give different mixes of the old information and the more recent state, and by understanding this, the public can infer both. Note that if agents are naive, e.g., if they think that all of the signals from the media are not (strongly) correlated conditional on the state, they will put positive weights on their observations and will again be bounded from learning the state.

This illustration relied on a particular network with several special features, including one-directional links and no communication among the media outlets or public. We will show that the same considerations determine learning quality in a large class of random networks in which agents have many neighbors, with complex connections among them. Quite generally, diversity of signal endowments in their neighborhoods allows agents to concentrate on new developments in the state while filtering out old, less relevant information and thus estimate the changing state as accurately as physical constraints allow.

**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 sophisticated responses to correlations. Section 6 relates our model and results to the social learning literature. In Section 7, we discuss our numerical exercise, structural estimation of our model, and multidimensional states. Finally, we comment on how the main finding on the importance of signal diversity can be generalized to more general signal distributions.

## 2. MODEL

**2.1. Description.** We describe the environment and game; complete details are formalized in Appendix A.

**State of the world.** There is a discrete set of instants of time,

$$\mathcal{T} = \mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}.$$

(As an alternative, we will also discuss a specification with a starting time,  $\mathcal{T} = \mathbb{Z}_{\geq 0} = \{0, 1, 2, \dots\}$ , but will focus on the doubly-infinite model.)

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

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

where  $\rho$  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 all actions taken 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  $t - m$  at a certain location (node) 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  $\nu_t$  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.<sup>2</sup> By a standard fact about squared-error loss functions, given the distribution of  $(\mathbf{a}_{N_i,t-\ell})_{\ell=1}^m$ , she sets:

$$(2.1) \quad a_{i,t} = \mathbb{E}[\theta_t \mid s_{i,t}, (\mathbf{a}_{N_i,t-\ell})_{\ell=1}^m].$$

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 (2.1) depends on the prior of agent  $(i, t)$  about  $\theta_t$ , which can be any normal distribution or a uniform improper prior (in which case all of  $i$ ’s beliefs about  $\theta_t$  come from her own signal and her neighbors’ actions).<sup>3</sup> 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.

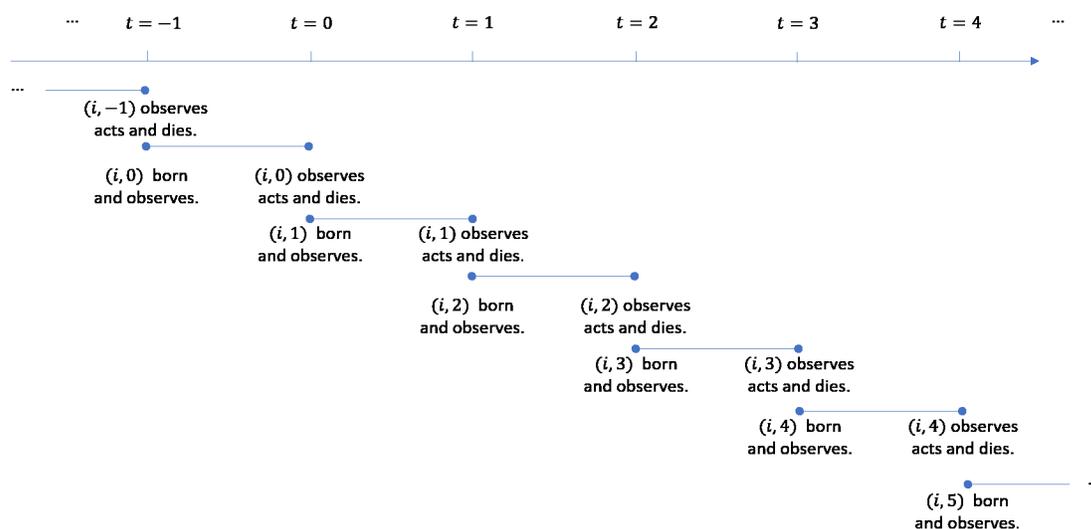
**2.2. Interpretation.** The agents minimize their loss function given the information they have access to and their beliefs’ about others’ play. The canonical case is one in which they are Bayesian and so everyone’s expectations assume the correct distribution of play for others, though the model is readily adapted to other behavioral assumptions, as we will discuss later. 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

<sup>2</sup>In Section 2.2 we discuss this assumption and how it relates to applications.

<sup>3</sup>With  $0 < \rho < 1$ , one natural choice for a prior is the stationary distribution of the state.

FIGURE 2.1. 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  $a_{i,t}$ .

to changing this assumption—but it is worth explaining. Whereas  $a_{i,t}$  is an observable choice, such as a published evaluation of an asset or a mix of inputs actually used by an agent in production, the private signals are not shareable.<sup>4</sup>

Finally, our agents act once and do not consider future payoffs, which shuts down the possibility that they may distort reports to manipulate the future path of social learning for their successors' benefit. Equivalently, we could simply assume that agents sincerely announce their subjective expectations of the state, as in [Geanakoplos and Polemarchakis \(1982\)](#) and the literature following it. For discussions of this type of assumption in social learning models, and ways to relax it, see, for instance [Mueller-Frank \(2013\)](#) and [Mossel et al. \(2015\)](#).

We discuss extensions of the basic model in various directions in [Section 7](#).

### 3. UPDATING AND EQUILIBRIUM

In this section we study agents' learning behavior and present a notion of stationary equilibrium. We begin with the canonical case of Bayesian agents with correct models

<sup>4</sup>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.

of others' behavior; the analysis is adapted to other behavioral assumptions in Section 5 below.

**3.1. Best-response behavior.** The first step is to analyze best-response updating behavior given others' strategies. The following observations apply whether the agents are in equilibrium or not.

A strategy of an agent is *linear* if the action taken is a linear function of the variables in her information set. We first discuss agents' best responses to linear strategies, showing that they are linear and computing them explicitly.

Suppose predecessors have played linear strategies up to time  $t$ .<sup>5</sup> Then we can express each action up until time  $t$  as a weighted summation of past signals. Because all innovations  $\nu_t$  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 (3.1) below for details).

**3.1.1. Covariance matrices.** The optimal weights for an agent to place on her sources of information depend on the variances and covariances of these sources. 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  $\theta_t$  based on actions  $a_{i,t-\ell}$  during the past  $m$  periods and the current 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  $V_{ij,t}$ .

**3.1.2. Best-response weights.** A strategy profile is a best-response if the weights each agent places on the variables in her information set minimize her posterior variance. We now characterize such weights 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

---

<sup>5</sup>We will discuss this below in the context of our equilibrium concept; but one immediate motivation is that, in the model with a starting time, where  $\mathcal{T} = \mathbb{Z}_{\geq 0}$ , Bayesian agents' updating at  $t = 0$  is a single-agent problem where optimal behavior is a linear function of own signals only, and thus the hypothesis holds.

to neighbors of  $i$ <sup>6</sup> 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  $\theta_t$  is normally distributed with mean

$$(3.1) \quad \frac{\mathbf{1}^T \mathbf{C}_{i,t-1}^{-1}}{\mathbf{1}^T \mathbf{C}_{i,t-1}^{-1} \mathbf{1}} \cdot \begin{pmatrix} \rho \mathbf{a}_{N_i,t-1} \\ \vdots \\ \rho^m \mathbf{a}_{N_i,t-m} \\ s_{i,t} \end{pmatrix}.$$

(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 (3.1) is a linear combination of the agent's signal and the observed actions; the weights in this linear combination depend on the matrix  $\mathbf{V}_{t-1}$  (but not on realizations of any random variables). In (3.1) we use our assumption of an improper prior.<sup>7</sup>

We denote by  $(\mathbf{W}_t, \mathbf{w}_t^s)$  a *weight profile* in period  $t$ , with  $\mathbf{w}_t^s \in \mathbb{R}^n$  being the weights agents place on their private signals and  $\mathbf{W}_t$  recording the weights they place on their other information.<sup>8</sup>

3.1.3. *The evolution of covariance matrices under best-response behavior.* In view of the formula (3.1) 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

$$(3.2) \quad \Phi : \mathbf{V}_{t-1} \mapsto \mathbf{V}_t.$$

This map gives the basic dynamics of the model: how a variance-covariance matrix  $\mathbf{V}_{t-1}$  maps to a new one when all agents best-respond to an arbitrary  $\mathbf{V}_{t-1}$ . The variance-covariance matrix  $\mathbf{V}_{t-1}$  (along with parameters of the model) determines (i) the weights agents place on their observations and (ii) the covariances of the random variables that are

<sup>6</sup>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  $\ell \in \{1, \dots, m\}$ .

<sup>7</sup>As we have mentioned, this is for convenience and without loss of generality. Our analysis applies equally to any proper normal prior for  $\theta_t$ : To get an agent's estimate of  $\theta_t$ , the formula in (3.1) 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.

<sup>8</sup>We do not need to describe the indexing of coefficients in  $\mathbf{W}_t$  explicitly in general; this would be a bit cumbersome because there are weights on actions at various lags.

being combined. This yields the deterministic updating dynamic  $\Phi$ . That is, the weights agents place on observations are (commonly) known, and do not depend on any random realizations.

**Example 1.** We compute the map  $\Phi$  explicitly in the case  $m = 1$ . 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  $w_{i,t}^s$ . Note we have, from (3.1) above, explicit expressions for these weights. Then

$$(3.3) \quad [\Phi(\mathbf{V})]_{ij} = (w_i^s)^2 \sigma_i^2 + \sum W_{ik} W_{ik'} (\rho^2 V_{kk'} + 1) \quad \text{and} \quad V_{ij} = \sum W_{ik} W_{jk'} (\rho^2 V_{kk'} + 1).$$

**3.2. Stationary equilibrium in linear strategies.** We will now turn our attention to *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 use the word *equilibrium* to refer to one of these unless otherwise noted.

A reason for focusing on equilibria in linear strategies comes from noting that, in the variant of the model with a starting time,  $\mathcal{T} = \mathbb{Z}_{\geq 0}$ , agents begin by using only private signals, and they do this linearly. After that, inductively applying the reasoning of 3.1, best-responses are linear at all future times. Taking time to extend infinitely backward is an idealization that allows us to focus on exactly stationary behavior.

We now show the existence of stationary equilibria in linear strategies.<sup>9</sup>

**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 such an 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  $\widehat{W}_{ij}$  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.

<sup>9</sup>Because time in this game is doubly infinite, there are some subtleties in definitions, which are dealt with in Appendix A.

To find a convex, 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 above and bounded away from zero below. This is because all agents' actions must be at least as precise in estimating  $\theta_t$  as their private signals, and cannot be more precise than estimates given perfect knowledge of  $\theta_{t-1}$  combined with the private signal. Because the Cauchy-Schwartz inequality bounds covariances in terms of the corresponding variances, it follows that there is a compact, convex set containing the image of  $\Phi$ . This and the continuity of  $\Phi$  allow us to apply the Brouwer fixed-point theorem.

In the case of  $m = 1$ , we can use the formula of Example 1, equation (3.3), to write the fixed-point condition  $\Phi(\widehat{\mathbf{V}}) = \widehat{\mathbf{V}}$  explicitly. More generally, for any  $m$ , we can obtain 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 describe the equilibrium  $\widehat{V}_{ij}$  as solving 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.<sup>10</sup>

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.<sup>11</sup>

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.<sup>12</sup> In Section 4.5 and Appendix D we study a variant of the model in which there is a unique equilibrium outcome, without stationarity.

We now briefly touch on 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

<sup>10</sup>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  $\Phi$ , as written in (3.3), to an initial covariance matrix. In all our experiments, the same fixed point has been found, independent of starting conditions.

<sup>11</sup>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.

<sup>12</sup>We have checked numerically that  $\Phi$  is not, in general, a contraction in any of the usual norms (entrywise sup, Euclidean operator norm, etc.). In computing equilibria numerically for many examples, we have not been able to find a case of equilibrium multiplicity.

variances and covariances are available (i.e., the entries of the matrix  $\mathbf{V}_{N_i,t}$  discussed in Section 3.1), then the agent needs only to use these to compute a best estimate of  $\theta_{t-1}$ , which is essentially a linear regression problem.

#### 4. WHEN IS THERE FAST INFORMATION AGGREGATION IN LARGE NETWORKS?

In this section, we consider the quality of learning outcomes in stationary equilibrium, and when good aggregation can be achieved. To make this exercise precise, we first define a benchmark of good information aggregation (Section 4.1). Our main results show that, under certain conditions on the distribution of signals in the population, this benchmark can actually be achieved in equilibrium, in a class of large networks. The conditions require that signals are distributed in a suitably heterogeneous way within each neighborhood. These conditions and results are described in Sections 4.2 and 4.3. We then show, in Section 4.4, that this result is tight: without such signal diversity, aggregation can fail in the unique equilibrium.

We then turn to the robustness of these results and the contrast with the case of a fixed state. As an alternative to the fully stationary model, Section 4.5 discusses a variant where there is an initial time period. Analogous results hold in that setting; indeed, in this variant, which features a unique equilibrium, we are able to strengthen the prediction that good learning obtains with signal diversity. We also use that variant to note that with an unchanging state, good aggregation does not depend on signal diversity.

This section relies on asymptotics of large random graphs. Section 7.1 in the discussion presents numerical results on some real-world networks, and verifies that the asymptotic and random-graph assumptions made to facilitate the main analysis are not crucial to the central point that signal diversity is helpful for learning.

**4.1. The aggregation benchmark.** 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.** Agent  $i$  achieves the  $\varepsilon$ -perfect aggregation benchmark if

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

An equilibrium achieves the  $\varepsilon$ -perfect aggregation benchmark if all agents do.

This says that all agents do nearly as well as if each knew her private signal and yesterday's state. Note agents can never infer yesterday's state perfectly from observed actions in any finite network (even if they have access to all signals previously observed by anyone), and so we must have  $\frac{\widehat{V}_{ii}}{V_{ii}^{\text{benchmark}}} > 1$  for all  $i$  on any fixed network.

We give conditions under which  $\varepsilon$ -perfect aggregation is achieved for any  $\varepsilon > 0$  on large enough networks. To make this formal, we fix  $\rho$  and consider a sequence of networks  $(G_n)_{n=1}^\infty$ , where  $G_n$  has  $n$  nodes.

**Example 2.** We use a very simple example to demonstrate that the  $\varepsilon$ -perfect aggregation benchmark can be achieved for any  $\varepsilon > 0$ . 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  $V_{11}^{\text{benchmark}} = (\sigma_1^{-2} + 1)^{-1} = \frac{1}{2}$  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 the next section, we address this and show that perfect aggregation can be achieved by all agents simultaneously even without anyone having very precise signals.

**4.2. Distributions of networks and signals.** To study learning in large populations, we specify two aspects of the environment: *network distributions* and *signal distributions*. In terms of network distributions, we define a stochastic model that makes the analysis of large networks tractable, but is flexible in that it allows us to encode rich heterogeneity in network positions. We also specify *signal distributions*: how signal precisions are allocated to agents, in a way that may depend on network position. We now describe these two primitives of the model and make some assumptions that are maintained in Section 4.

Fix a set of *network types*  $k \in \mathcal{K} = \{1, 2, \dots, K\}$ . There is a probability  $p_{kk'}$  for each pair of network types, which is the probability that an agent of network type  $k$  has a link to a given agent of network type  $k'$ . 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. There is also a vector  $(\alpha_1, \dots, \alpha_K)$  of *population shares* of each type, which we assume are all positive. Jointly,  $(p_{kk'})_{k,k' \in \mathcal{K}}$  and  $\alpha$  specify the network distribution.<sup>13</sup> These

<sup>13</sup>This type of network is known as a *stochastic block model* (Holland et al., 1983).

parameters can encode differences in expected degree and also features such as homophily (where some groups of types are linked to each other more densely than to others).

We next define signal distributions, which describe the allocation of signal variances to network types. Fix a finite set  $\mathbb{S}$  of private signal variances, which we call signal types.<sup>14</sup> We let  $q_{k\tau}$  be the share of agents of network type  $k$  with signal type  $\tau$ ;  $(q_{k\tau})_{k \in \mathcal{K}, \tau \in \mathbb{S}}$  defines the signal distribution.

**Generating the networks.** Let the nodes in network  $n$  be a disjoint union of sets  $N_n^1, N_n^2, \dots, N_n^K$ , with the cardinality  $|N_n^k|$  equal to  $\lfloor \alpha_k n \rfloor$  or  $\lceil \alpha_k n \rceil$  (rounding so that there are  $n$  agents in the network). We (deterministically) set the signal variances  $\sigma_i^2$  equal to elements of  $\mathbb{S}$  in accordance with the signal shares (again rounding as needed). Let  $(G_n)_{n=1}^\infty$  be a sequence of directed or undirected random networks with these nodes, so that  $i \in N_n^k$  and  $j \in N_n^{k'}$  are linked with probability  $p_{kk'}$ ; these realizations are all independent.

**Diversity of signals.** The environment is described by the linking probabilities  $(p_{kk'})_{k, k' \in \mathcal{K}}$ , the type shares  $\alpha$ , and the signal distribution  $(q_{k\tau})_{k \in \mathcal{K}, \tau \in \mathbb{S}}$ . We say that a signal type  $\tau$  is represented in a network type  $k$  if  $q_{k\tau} > 0$ .

**Definition 2.** We say that the environment satisfies *signal diversity* if at least two distinct signal types are represented in each network type.

We will discuss environments that satisfy this condition as well as ones that do not, and show that it is pivotal for information-aggregation.

**4.3. Diverse signals.** Our first main result is that signal diversity is sufficient for good aggregation. Under this condition, the benchmark is achieved independently of the structural properties of the network.

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  $\varepsilon > 0$ . If an environment satisfies signal diversity, asymptotically almost surely  $G_n$  has a equilibrium where the  $\varepsilon$ -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., an agent has no neighbors). We give an outline of the argument next, and the proof appears in Appendix C.

<sup>14</sup>The assumptions of finitely many signal types and network types are purely technical, and could be relaxed.

*Outline of the argument.* To give intuition for the result, we first describe why the theorem holds on the complete network<sup>15</sup> with two signal types  $A$  and  $B$  in the  $m = 1$  case. This echoes the intuition of the example in the introduction. 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  $\theta_{t-1}$  based on the actions she has observed in her neighborhood. On the complete network, all players have the same social signal, which we call  $r_t$ .<sup>16</sup>

At any equilibrium, each agent's action is a weighted average of her private signal and this social signal:<sup>17</sup>

$$(4.1) \quad a_{i,t} = \widehat{w}_i^s s_{i,t} + (1 - \widehat{w}_i^s) r_t.$$

The weight  $\widehat{w}_i^s$  depends only on the precision of agent  $i$ 's signal. We call the weights of two distinct signal types  $\widehat{w}_A^s$  and  $\widehat{w}_B^s$ .

Now observe that each time- $(t + 1)$  agent can average the time- $t$  actions of each type, which can be written as follows using (4.1) and  $s_{i,t} = \theta_t + \eta_{i,t}$ :

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

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 convex combinations of  $\theta_t$  and  $r_t$ . Because the two weights,  $\widehat{w}_A^s$  and  $\widehat{w}_B^s$ , are distinct, she can approximately solve for  $\theta_t$  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  $\theta_t$  that this strategy gives will place negative weight on  $\frac{1}{n_A} \sum_{i:\sigma_i^2=\sigma_A^2} a_{i,t-1}$ , 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.

<sup>15</sup>Note this is a special case of the stochastic block model.

<sup>16</sup>In particular, agent  $(i, t)$  sees everyone's past action, including  $a_{i,t-1}$ .

<sup>17</sup>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  $\theta_t$ .

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. Showing this is more subtle, because the social signals at an equilibrium are endogenous, and—if the network is arbitrary—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. (“Length two” is not essential: in sparser random graphs, this statement holds with a different length, and the same arguments go through, as we discuss below.) 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  $\Phi^2$  maps the neighborhood of covariance matrices where all social signals are close to perfect to itself, and then we apply a fixed point theorem.  $\square$

Concerning the rate of learning as  $n$  grows, the proof implies that, under the assumptions of the theorem, the error in agents’ estimates of  $\theta_{t-1}$  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 details; in particular, if the difference between signal precisions is small, the constants may be very bad. Section 7.1 offers numerical evidence on the quality of aggregation in networks of practically relevant sizes.

**Sparser random graphs.** The random graphs we defined have held  $(p_{kk'})$  fixed for simplicity. This yields expected degrees that grow linearly in the population size, which may not be the desired asymptotic model. While it is important to have neighborhoods “large enough” (i.e., growing in  $n$ ) to permit the application of laws of large numbers, their rate of growth can be considerably slower than linear: for example, our proof extends directly to degrees that scale as  $n^\alpha$  for any  $\alpha > 0$ . Instead of studying  $\Phi^2$  and second-order neighborhoods, we apply the same analysis to  $\Phi^k$  and  $k^{\text{th}}$ -order neighborhoods for  $k$  larger than  $1/\alpha$ .

**4.4. Non-diverse signals.** So far we have seen that assuming signal diversity, good aggregation is obtained independently of the structure of the large random graph. We now

show this result is tight: in environments without signal diversity, equilibrium aggregation can (necessarily) be much worse.

To gain an intuition for this, note that 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. 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  $\theta_t$ . We now show that without diversity in signal quality, information aggregation may be much worse.

We first study a class of networks with certain symmetries and show that for this class there is a unique equilibrium, and at this equilibrium good aggregation is not achieved (Section 4.4.1). We then present a 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.

A natural question is whether asymmetry in network positions can substitute for asymmetry in signal endowments. We thus move beyond completely symmetric networks and examine random graphs. We show there exists a similar equilibrium without good aggregation on random graphs. This shows that some variation in network positions need not give individuals enough power to identify the state, as they can under signal diversity.

#### 4.4.1. *Graphs with symmetric neighbors.*

**Definition 3.** A network  $G$  has *symmetric neighbors* if  $N_j = N_{j'}$  for any  $i$  and any  $j, j' \in N_i$ .

In the undirected case, the graphs with symmetric neighbors are the complete network and complete bipartite networks.<sup>18</sup> 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  $\sigma^2$ , 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  $\varepsilon$ -perfect aggregation benchmark is not achieved by any agent  $i$  at this equilibrium for any  $n$ .*

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 E.

<sup>18</sup>These are both special cases of our stochastic block model from Section 4.3.

This immediately implies that in environments not satisfying signal diversity, there are network distributions for which aggregation fails in a strong sense.

This 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  $\varepsilon$ -perfect aggregation for any  $\varepsilon > 0$  when  $n$  is large. See Appendix G for a formal statement, proof and numerical results.<sup>19</sup> In this sense, the social learning externalities are quite severe: a fairly small change in weights for each individual could yield a very large benefit in a world of homogeneous signal types.

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.<sup>20</sup> So actions must be unweighted averages of observations. This prevents the sort of inference of  $\theta_t$  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  $\frac{1}{|N_i|} \sum_{j \in N_i} a_{j,t}$ :

$$(4.2) \quad a_{i,t+1} = \widehat{w}_i^s s_{i,t+1} + (1 - \widehat{w}_i^s) \frac{1}{|N_i|} \sum_{j \in N_i} a_{j,t}.$$

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* 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  $V^\infty$  denote the limit, as  $n$  grows large, of the variance of any agent’s error ( $a_{i,t} - \theta_t$ ). Let  $Cov^\infty$  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:

$$(4.3) \quad V^\infty = \frac{1}{\sigma^{-2} + (\rho^2 Cov^\infty + 1)^{-1}}, \quad Cov^\infty = \frac{(\rho^2 Cov^\infty + 1)^{-1}}{[\sigma^{-2} + (\rho^2 Cov^\infty + 1)^{-1}]^2}.$$

This variance and covariance describe behavior not only in the complete graph, but in any graph with symmetric neighbors where degrees tend uniformly to  $\infty$ . In such graphs, too, the variances of all agents converge to  $V^\infty$  and the covariances of all pairs of agents converge to  $Cov^\infty$ , as  $n \rightarrow \infty$ .<sup>21</sup> This implies that, in large graphs, the equilibrium action

<sup>19</sup>We thank Alireza Tahbaz-Salehi for suggesting this analysis.

<sup>20</sup>The proof of the proposition establishes uniqueness by showing that  $\Phi$  is a contraction in a suitable sense.

<sup>21</sup>This is established by the same argument as in the proof of Proposition 3.

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.

4.4.2. *A corollary: Perverse consequences of improving signals.* As a consequence of Theorem 1 and Proposition 2, we can give an example where making one agent's private information less precise helps all agents.

**Corollary 1.** *There exists a network  $G$  and an agent  $i \in G$  such that increasing  $\sigma_i^2$  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.

4.4.3. *Non-diverse signals in large random graphs.* Our results on non-diverse signals have used graphs with symmetric neighbors. In those graphs, 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 Erdos-Renyi random networks, there is an equilibrium with essentially the same learning outcomes when signal precisions are homogeneous.

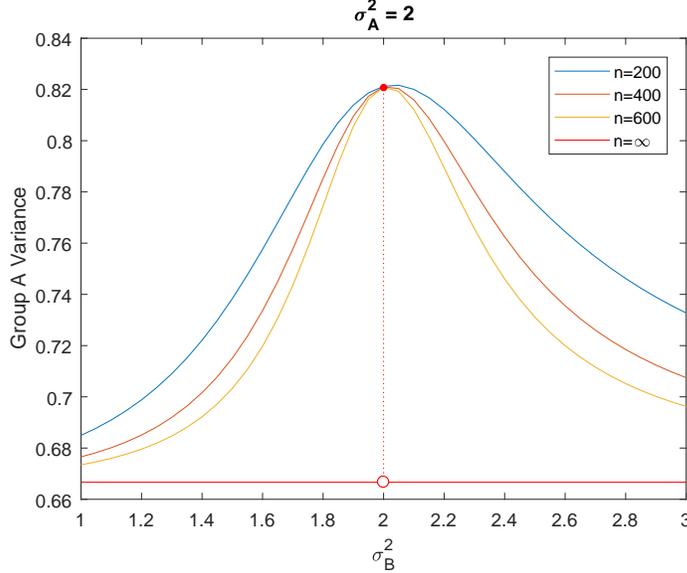
Let  $(G_n)_{n=1}^\infty$  be a sequence of directed or undirected random networks, with  $G_n$  having  $n$  nodes, with any pair of distinct nodes linked (i.i.d.) with positive probability  $p$ . We continue to assume all signal variances are equal to  $\sigma^2$  and  $m = 1$ .

**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  $\varepsilon$ -perfect aggregation benchmark is not achieved.*

The equilibrium covariances in this equilibrium again converge to  $V^\infty$  and  $Cov^\infty$  (for any value of  $p$ ). Thus, when there is only one signal type, we obtain the same learning outcomes asymptotically on a variety of networks.

4.4.4. *How much diversity is needed?* The results in Section 4 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

FIGURE 4.1. Dependence of Learning on Signal Variances



$n \rightarrow \infty$ , 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 to the second is actually gradual: Action error remains well above the perfect aggregation benchmark when signal qualities differ slightly.

In Figure 4.1, 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  $a_{i,t} - \theta_t$  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  $\sigma_B^2$ . 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  $\sigma_B^2 = 3$ .

**4.5. Model with an initial time: Aggregation quality and contrast with a fixed state.** The same questions about the quality of information aggregation can be asked in the alternative model with  $\mathcal{T} = \mathbb{Z}_{\geq 0}$ . This analysis complements our main points above. We summarize the analysis here and defer the details to Appendix D.

In the model with a starting time, there is an unambiguous prediction of behavior. Time-zero agents have only their own signals to use, so they play simple linear strategies, and

subsequent agents inductively set their weights based on the prior ones, as described in Section 3.1.

We study the same environments (signal types and network distributions) discussed in the previous subsections. The counterpart of Theorem 1, says that under the same conditions on the environment as in that theorem, learning as good as the benchmark obtains for all times  $t \geq 1$ . The counterpart of Proposition 2 says that in graphs with symmetric neighbors, without signal diversity, learning short of the benchmark is the unique long-run outcome. And the counterpart of Proposition 3 says that, under the same conditions on the environment, learning short of the benchmark is again the unique long-run outcome.

The first and third results sharpen the analysis relative to the stationary model, in that there is no question of the possible existence of other equilibria. On the other hand, in the  $\mathcal{T} = \mathbb{Z}_{\geq 0}$  model, updating weights are generally dynamic and thus more complex to describe. We view the  $\mathcal{T} = \mathbb{Z}_{\geq 0}$  analysis as an important robustness check: it verifies that exact stationarity is not crucial to the findings, and shows that the good learning outcomes are the ones “selected” by the model with a starting time.

Finally, the model with a starting time allows us to contrast the findings of this section with the case of an unchanging state. (That is, the evolution of  $\theta_t$  is replaced by  $\theta_t = \theta$  for all  $t$ , where  $\theta$  is standard normal, and the model is otherwise the same.) Consider any fixed, finite network where  $i \in N_i$  for each  $i$ , so that each agent observes an own-node predecessor, and each agent has a signal with finite, positive precision. It is straightforward to verify inductively that the precision of information held by  $(i, t)$  about the unchanging state  $\theta$  increases by at least some fixed  $\delta > 0$  every period, because that node can use the  $(i, t - 1)$  estimate as well as a new independent signal. Thus, estimates will eventually be arbitrarily precise. This prediction will not be sensitive to any other details of signal or network structure, echoing the findings of DeMarzo et al. (2003) and Lobel and Sadler (2015) on sufficient conditions good learning for Bayesians. This highlights the fundamental difference between our changing-state model and a static analogue.

## 5. THE IMPORTANCE OF UNDERSTANDING CORRELATIONS

In the positive result on achieving the perfect aggregation benchmark (Theorem 1), a key aspect of the argument involved agents filtering out confounding information from their neighbors’ estimates—i.e., responding in a sophisticated way to the correlation structure of those estimates. In this section, we demonstrate that this sort of behavior is essential for nearly perfect aggregation, and that more naively imitative heuristics yield outcomes far from the benchmark. Empirical studies have found evidence (depending on the setting and

the subjects) consistent with both equilibrium behavior and naive inference in the presence of correlated observations (e.g., [Eyster, Rabin, and Weizsacker \(2015\)](#); [Dasaratha and He \(2017\)](#); [Enke and Zimmermann \(2019\)](#)).

We begin with a canonical model of agents who do not account for correlations among their neighbors’ estimates conditional on the state, and show by example that naive agents achieve much worse learning than Bayesian agents, and thus fail to reach the perfect aggregation benchmark. We then formalize the idea that accounting for correlations in neighbors’ actions is crucial to reaching the benchmark. This is done by demonstrating a general lack of asymptotic learning by agents who use imitative strategies, rather than filtering in a sophisticated way. Finally, we show that even in fixed, finite networks, any positive weights chosen by optimizing agents will be Pareto-dominated.

**5.1. Naive agents.** In this part we introduce agents who misunderstand the distribution of the signals they are facing and who therefore do not update as Bayesians with a correct understanding of their environment. We consider a particular form of misspecification that simplifies solving for equilibria analytically:<sup>22</sup>

**Definition 4.** 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, in our model, of “best-response trailing naive inference” ([Eyster and Rabin, 2010](#)). So naive agents understand that their neighbors’ actions from the previous period are estimates of  $\theta_{t-1}$ . 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. They then play their expectation of the state given this misspecified theory of others’ play.

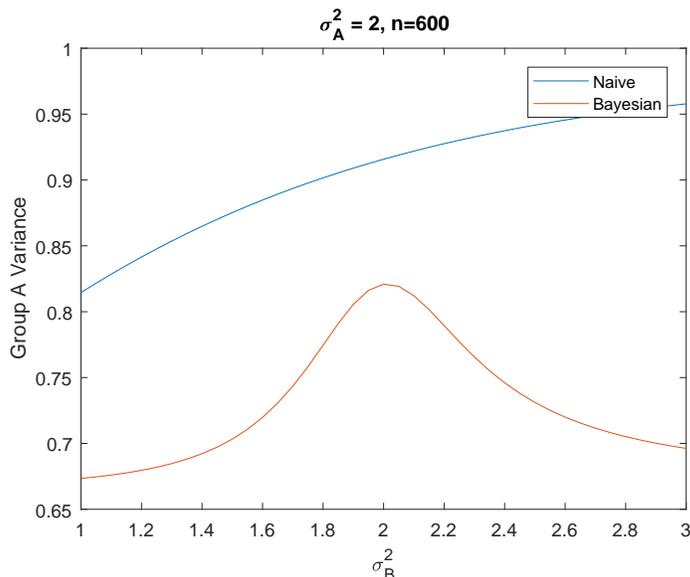
In [Figure 5.1](#), we compare Bayesian and naive learning outcomes. As in [Figure 4.1](#), we consider a complete network where half of agents have signal variance  $\sigma_A^2 = 3$  and we vary the signal variance  $\sigma_B^2$  of the remaining agents. We observe that naive agents learn substantially worse than rational agents, whether signals are diverse or not. Formal analysis and formulas for variances under naive learning can be found in [Appendix F](#).

**5.2. More general learning rules: Understanding correlation is essential for reaching the benchmark.** We now show more generally that a sophisticated response to

---

<sup>22</sup>A seminal paper studying boundedly rational learning rules in networks is [Bala and Goyal \(1998\)](#). 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).

FIGURE 5.1. Bayesian and Naive Learning



correlation is needed to achieve perfect aggregation on any sequence of growing networks. To this end, we make the following definition:

**Definition 5.** 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 best-responding to others' actual distributions of play, agents use exogenous weights  $\mathbf{W}$  in all periods.

By a straightforward application of the contraction mapping theorem, if agents use any non-negative weights under which covariances remain bounded at all times, there is a unique steady state.

Consider a sequence of networks  $(G_n)_{n=1}^{\infty}$  with  $n$  agents in  $G_n$ .

**Proposition 4.** Fix any sequence of steady states under non-negative weights on  $G_n$ . Suppose that all private signal variances are bounded below by  $\underline{\sigma}^2 > 0$  and that all agents place weight at most  $\bar{w} < 1$  on their private signals. Then there is an  $\varepsilon > 0$  such that, for all  $n$ , the  $\varepsilon$ -aggregation benchmark is not achieved by any agent  $i$  at steady state.

The essential idea is that at time  $t + 1$  observed time- $t$  actions all put weight on actions from period  $t - 1$ , which causes  $\theta_{t-1}$  to have a (positive weight) contribution to all observed

actions. Agents do not know  $\theta_{t-1}$  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 see why the weights on private signals must be bounded away from one, note that an individual agent could learn well without adjusting for correlations by observing many autarkic agents who simply report their private signals. But in this case, all of these autarkic agents would be far from the benchmark. Without the bound on private signal weights, *some* agent must still fail to achieve perfect aggregation.

On undirected networks, the proposition implies that perfect aggregation is not achieved by naive agents or agents exhibiting various other specifications of non-Bayesian inference. Moreover, the same argument shows that in any sequence of *Bayesian* equilibria on undirected networks where all agents use positive weights, no agent can learn well.

**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.

**Proposition 5.** *Suppose the network  $G$  is strongly connected and some agent has more than one neighbor. Given any naive equilibrium or any Bayesian equilibrium where all weights are positive, the variances at that equilibrium are Pareto-dominated by variances at another steady state.*

The basic argument behind Proposition 5 is that if agents place marginally more weight on their private signals, this introduces more independent information that eventually benefits everyone. In the proof in Appendix E, we state and prove a more general result with weaker hypotheses on behavior.

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). Proposition 5 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 proposition.** We next briefly discuss the sufficient conditions in the proposition statement. First, 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.<sup>23</sup> Second, 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 force also works 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 3.** 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

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

<sup>23</sup> In fact, the result of Proposition 5 (with the same proof) applies to a larger class of networks: rather than strongly connectedness it is sufficient that, starting at each agent, there are two paths of some length  $k$  to a rational agent and another distinct agent.

by Hayek (1945)<sup>24</sup> and of primary interest in some applied problems, e.g., in real business cycle models with consumers and firms learning about evolving state. Nevertheless, there is relatively little modeling of Bayesian learning of dynamic states in the large literature on social learning and information aggregation in networks. We now put our contribution in the context of the most closely related papers on social learning in networks.<sup>25</sup>

Play in the stationary linear equilibria of our model closely resembles behavior in the DeGroot (1974) model, where agents update by linearly aggregating network neighbors' past estimates, with constant weights on neighbors over time. DeMarzo, Vayanos, and Zweibel (2003), in an environment with an unchanging state, derive DeGroot learning as a boundedly-rational heuristic. Each agent has one initial signal about the state, and this determines her  $t = 0$  estimate, which is shared with neighbors. Assuming all randomness is Gaussian, the Bayesian rule for forming estimates at  $t = 1$  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. Molavi, Tahbaz-Salehi, and Jadbabaie (2018) have offered new bounded-rationality foundations for the DeGroot rule. 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 stationary environment admits a stationary equilibrium in which fixed updating rules are best responses.<sup>26</sup> In the introduction, we have discussed the contrast between the predictions about outcomes that our model makes and the analysis of fixed-state benchmarks.

Several recent papers in computer science and engineering study dynamic environments similar to ours. Shahrampour, Rakhlin, and Jadbabaie (2013) study an exogenous-weights version, interpreted as a set of Kalman filters under the control of a planner; they focus on computing or bounding various measures of welfare in terms of network invariants and the persistence of the state process ( $\rho$ ). Frongillo, Schoenebeck, and Tamuz (2011) study (in our notation) a  $\theta_t$  that follows a random walk ( $\rho = 1$ ). They examine agents who

---

<sup>24</sup>“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... there still remains the problem of communicating to [each individual] such further information as he needs.” Hayek’s main concern was aggregation of information through markets, but the same questions apply more generally.

<sup>25</sup>For more complete surveys of different parts of this literature, see, among others, Acemoglu and Ozdaglar (2011), Golub and Sadler (2016), and Mossel and Tamuz (2017). See Moscarini et al. (1998) for an early model in a binary-action environment, where it is shown that a changing state can break information cascades.

<sup>26</sup>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.

learn using fixed, exogenous weights on arbitrary networks. They characterize the steady-state distribution of behavior with arbitrary (non-equilibrium) fixed weights. They also give a formula for equilibrium weights on a complete network and show these weights are inefficient. Our Proposition 5 on Pareto-inefficiency on an arbitrary network documents a closely related inefficiency. Our main question, the quality of equilibrium learning in large networks, is a topic not considered in these papers.

In economics, the model in [Alatas, Banerjee, Chandrasekhar, Hanna, and Olken \(2016\)](#) most closely resembles ours. There, agents are not Bayesian, ignoring the correlation between social observations, similarly to our naive models. 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 7.2.

Some recent learning models—e.g., [Sethi and Yildiz \(2012\)](#) and [Harel et al. \(2017\)](#)—consider other obstacles to learning in environments with Gaussian signals and a fixed state. In [Sethi and Yildiz \(2012\)](#), learning outcomes depend on whether individuals’ (heterogeneous) priors are independent or correlated. [Harel et al. \(2017\)](#), among others, study the rate of social learning in a Gaussian setting. Bad learning corresponds to this rate being low. In our setting, we can focus on steady-state errors, which provide an alternative measure of the quality of learning.

Finally, a robust aspect of rational learning in sequential models is the phenomenon of anti-imitation ([Eyster and Rabin, 2014](#)). They 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. [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#) and [Lobel and Sadler \(2015\)](#) show that such a heuristic leads to asymptotic learning in a sequential model. Our dynamic learning environment is different, as shown in Proposition 4: 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

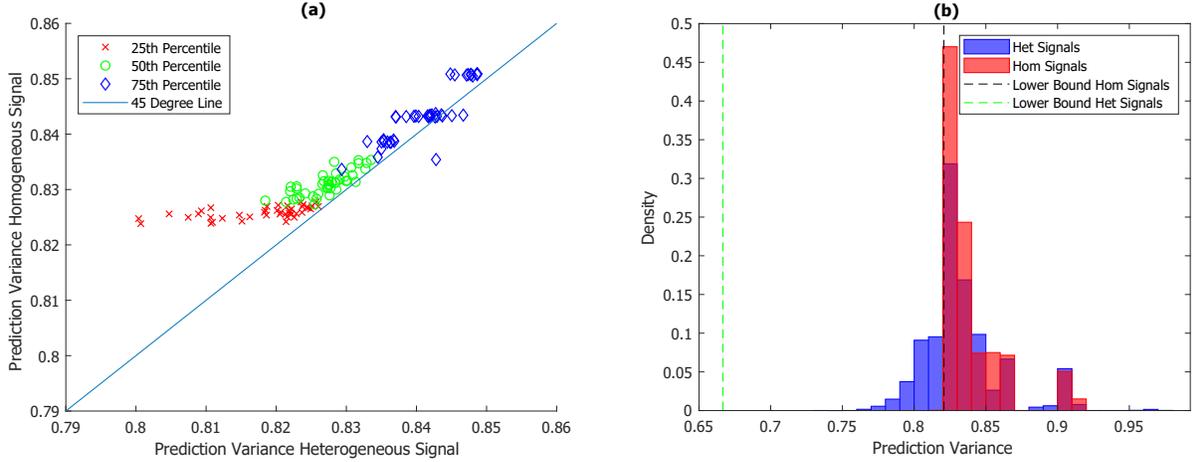
**7.1. Aggregation and its absence without asymptotics: Numerical results.** To examine whether the large network results above work in realistic networks with moderate degrees, we present numerical evidence based on the data in [Banerjee, Chandrasekhar, Duflo, and Jackson \(2013\)](#). This data set contains the social networks of villages in rural India.<sup>27</sup> 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 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.<sup>28</sup>

In Figure 7.1(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 25<sup>th</sup> 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. They see an improvement of 0.35 standard deviations compared to their homogeneous case counterparts. In Figure 7.1(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 results hold both qualitatively and quantitatively: adding heterogeneity

<sup>27</sup>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.

<sup>28</sup>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 7.1. Prediction Variance In Indian Villages



(a) The error variance of the agent in the 25<sup>th</sup>, 50<sup>th</sup> and 75<sup>th</sup> 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.

helps learning in the population, while even with a small group of agents with the new, inferior signal type the influence on at least some nodes in the network can be sizable.

**7.2. 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  $\rho$  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  $\theta_t$  ex post (perhaps with a long delay).<sup>29</sup>

<sup>29</sup>We can instead assume that the analyst observes (a proxy for) the private signal  $s_{i,t}$  of agent  $i$ ; we mention how below.

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 of our model when it applies. The strategy does not rely on uniqueness of equilibrium. We can identify the weights agents are using through standard vector autoregression methods. In steady state,

$$(7.1) \quad \bar{a}_{i,t} = \sum_j W_{ij} \rho \bar{a}_{j,t-1} + w_i^s \theta_t + \zeta_{i,t},$$

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.<sup>30</sup> 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.<sup>31</sup>

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 (3.3) for  $\Phi$ . In view of this formula, the signal variances  $\sigma_i^2$

<sup>30</sup>This system defines a VAR(1) process (or generally VAR( $m$ ) for memory length  $m$ ).

<sup>31</sup>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.

are uniquely determined by the other variables:

$$(7.2) \quad \widehat{V}_{ii} = \sum_j \sum_k \widehat{W}_{ij} \widehat{W}_{ik} (\rho^2 \widehat{V}_{jk} + 1) + (\widehat{w}_i^s)^2 \sigma_i^2.$$

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

Note that our basic VAR for recovering the weights relies only on constant linear strategies and does not assume that agents are playing any particular strategy within this class. Thus, if agents are using some other behavioral rule (e.g., optimizing in a misspecified model) we can replace (7.2) 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.

**7.3. Multidimensional states and informational specialization.** So far, we have been working with a one-dimensional state and one-dimensional signals, which varied only in their precisions. Our message about the value of diversity is, however, better interpreted in a mathematically equivalent multidimensional model.

Consider Bayesian agents who learn and communicate about two independent dimensions simultaneously (each one working as in our model). If all agents have equally precise signals about both dimensions, then society may not learn well about either of them. In contrast, if half the agents have superior signals about one dimension and inferior signals about the other (and the other half has the reverse), then society can learn well about both dimensions. Thus, the designer has a strong preference for an organization with informational specialization where some, but not all, agents are expert in a particular dimension.<sup>32</sup>

Of course, there are many familiar reasons for specialization, in information or any other activity. For instance, it may be that more total information can be collected in this case, or that incentives are easier to provide. Crucially, specialization is valuable in our setting for a reason distinct from all these: it helps agents with their inference problems.

---

<sup>32</sup>This raises important questions about what information agents would acquire, and whom they would choose to observe, which are the focus of a growing literature. For recent papers on this in the context of networks, see Sethi and Yildiz (2016) and Myatt and Wallace (2017), among others.

**7.4. General distributions.** The example of the previous subsection involved trivially extending our model to several independent dimensions. We now briefly discuss a more substantive extension, which applies to more realistic signal structures.

Our analysis of stationary linear learning rules relied crucially on the assumptions that the innovations  $\nu_t$  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  $\nu_t$  and signal errors  $\eta_t$ . For simplicity, consider the complete graph and  $m = 1$ .<sup>33</sup> Because  $\theta_{t-1}$  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  $\theta_{t-1}$  and her private signal. An agent with type  $\tau$  (which is observable) who believes  $\theta_{t-1}$  is distributed according to  $\mathcal{D}$  takes an action equal to  $f(\tau, \mathcal{D}, s_{i,t})$ . Here,  $\tau$  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  $\theta_{t-1}$ , the state  $\theta_t$  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.<sup>34</sup>

Now consider a time- $t$  agent  $i$ . Suppose now that any possible distribution that time- $(t-1)$  agents might have over  $\theta_{t-2}$  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  $\tau$  of  $t-1$  agent that  $i$  observes, the distribution of  $f(\tau, d, s_{i,t})$  gives an agent a different measurement of  $d$ , which is a summary of beliefs about  $\theta_{t-2}$ , and  $\theta_{t-1}$ . Assuming there is not too much "collinearity," these measurements of the finitely many parameters of interest should provide linearly

<sup>33</sup>These states and signals may now be multidimensional.

<sup>34</sup>The easiest way to ensure this is to consider a complete network and a large number of agents. However, notice that even if the neighborhood of an agent changes from period to period, if some of the individuals in that neighborhood are randomly sampled, this can provide information about the empirical distribution of actions of various types in the same way. That, in turn, can facilitate the identification of recent states, as we are about to explain.

independent information about  $\theta_{t-1}$ . Thus, as long as the set of signal types  $\tau$  is sufficiently rich, we would expect the identification condition to hold.

The simplest example of this is one in which the state is still AR(1) Gaussian, but now  $d$ -dimensional. Private signals, now also multidimensional, are arbitrary linear functions of  $\theta_t$ . If these linear functions are generic and there are sufficiently many distinct signal types, then observing actions will allow an observer to back out both  $\theta_{t-2}$ , and  $\theta_{t-1}$ . The essential observation about what is needed for good learning by Bayesians is that there are enough linearly independent action rules to identify the underlying  $d$  dimensions of fundamental uncertainty. In the Gaussian case such linear independence is guaranteed by having at least  $d + 1$  generic signal types.

## 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.
- 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.
- BABUS, A. AND P. KONDOR (2018): “Trading and Information Diffusion in Over-the-Counter Markets,” *Econometrica*, 86, 1727–1769.
- BALA, V. AND S. GOYAL (1998): “Learning from neighbours,” *The Review of Economic Studies*, 65, 595–621.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): “The diffusion of microfinance,” *Science*, 341, 1236–1248.
- DASARATHA, K. AND K. HE (2017): “Network Structure and Naive Sequential Learning,” *arXiv preprint arXiv:1703.02105*.
- 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.
- ENKE, B. AND F. ZIMMERMANN (2019): “Correlation neglect in belief formation,” *The Review of Economic Studies*, 86, 313–332.
- 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.
- EYSTER, E., M. RABIN, AND G. WEIZSACKER (2015): “An Experiment on Social Mislarning,” *Working Paper*.

- FRONGILLO, R., G. SCHOENEBECK, AND O. TAMUZ (2011): “Social Learning in a Changing World,” *Internet and Network Economics*, 146–157.
- 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.
- HOLLAND, P. W., K. B. LASKEY, AND S. LEINHARDT (1983): “Stochastic blockmodels: First steps,” *Social Networks*, 5, 109–137.
- JADBABAIE, A., P. MOLAVI, A. SANDRONI, AND A. TAHBAZ-SALEHI (2012): “Non-Bayesian social learning,” *Games and Economic Behavior*, 76, 210–225.
- 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 (2015): “Information diffusion in networks through social learning,” *Theoretical Economics*, 10, 807–851.
- 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.
- MYATT, D. AND C. WALLACE (2017): “Information Acquisition and Use by Networked Players,” University of Warwick. Department of Economics. CRETA Discussion Paper Series (32), available at <http://wrap.warwick.ac.uk/90449/>.
- 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*.
- 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 normal, mutually independent random variables, with  $\nu_t$  having variance 1 and  $\eta_{i,t}$  having variance  $\sigma_i^2$ . Also take a stochastic process  $(\theta_t)_{t \in \mathbb{Z}}$ , such that for each  $t \in \mathbb{Z}$ , 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.<sup>35</sup> Note that if  $\sigma \in \Sigma$ , then it can be checked that  $\tilde{\sigma} = (\sigma'_{(i,t)}, \sigma_{-(i,t)}) \in \Sigma$  whenever  $\sigma'_{(i,t)} \in \tilde{\Sigma}_{(i,t)}$ . 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

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

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  $\Sigma$ .<sup>36</sup> 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  $\sigma_i$  is a linear function. Our analysis focuses on *stationary, linear equilibria*.

## APPENDIX B. EXISTENCE OF EQUILIBRIUM: PROOF OF PROPOSITION 1

Recall from Section 3.1 the map  $\Phi$ , 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 convex, compact set,  $\mathcal{K}$ , to itself, which, by Brouwer's fixed-point theorem, ensures that  $\Phi$  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  $j$

<sup>35</sup>Condition 1 is necessary to rule out response profiles such as the one given by  $\sigma_{i,t}(s_{i,t}, a_{i,t-1}) = |a_{i,t-1}| + 1$ . 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.

<sup>36</sup>I.e.,  $\sigma' = \mathfrak{T}_p \sigma$  is defined by  $\sigma_{(i,t)} = \sigma_{(i,t-p)}$ .

respectively (at potentially different times), and let  $\bar{\sigma}_i^2 = \max\{\sigma_i^2, \rho^{m-1}\sigma_i^2 + \frac{1-\rho^{m-1}}{1-\rho}\}$ . Now let  $\mathcal{K} \subset \mathcal{V}$  be the subset of symmetric positive semi-definite matrices  $\mathbf{V}_t$  such that, for any such  $k, l$ ,

$$V_{kk,t} \in \left[ \min \left\{ \frac{1}{1 + \sigma_i^{-2}}, \frac{\rho^{m-1}}{1 + \sigma_i^{-2}} + \frac{1 - \rho^{m-1}}{1 - \rho} \right\}, \max \left\{ \sigma_i^2, \rho^{m-1}\sigma_i^2 + \frac{1 - \rho^{m-1}}{1 - \rho} \right\} \right]$$

$$V_{kl,t} \in [-\bar{\sigma}_i\bar{\sigma}_j, \bar{\sigma}_i\bar{\sigma}_j].$$

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  $\sigma_i^2$ . For a lower bound, note that if she knew the previous period's state and her own signal, then the variance of her action would be  $\frac{1}{1+\sigma_i^{-2}}$ . Thus an agent observing only noisy estimates of  $\theta_t$  and her own signal can do no better.

By the same reasoning applied to the node- $i$  agent from  $m$  periods ago, the error variance of  $\rho^m a_{i,t-m} - \theta_t$  is at most  $\rho^m \sigma_i^2 + \frac{1-\rho^m}{1-\rho}$  and at least  $\frac{\rho^m}{1+\sigma_i^{-2}} + \frac{1-\rho^m}{1-\rho}$ . This establishes bounds on  $V_{kk,t}$  for observations  $k$  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  $\Phi$ , this means there exists some weight profile  $(\widehat{\mathbf{W}}, \widehat{\mathbf{w}}^s)$  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  $(\widehat{\mathbf{W}}, \widehat{\mathbf{w}}^s)$  forever.

To show this, first observe that if agents have been using the weights  $(\widehat{\mathbf{W}}, \widehat{\mathbf{w}}^s)$  forever, the variance-covariance matrix  $\mathbf{V}_t$  in any period is uniquely determined and does not depend on  $t$ ; call this  $\check{\mathbf{V}}$ .<sup>37</sup> 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  $(\widehat{\mathbf{W}}, \widehat{\mathbf{w}}^s)$  that there is a distribution

<sup>37</sup>The variance-covariance matrices are well-defined because the  $(W, w^s)$  weights yield unambiguous strategy profiles in the sense of Appendix A.

of actions where the variance-covariance matrix is  $\widehat{\mathbf{V}}$  in every period and agents are using weights  $(\widehat{\mathbf{W}}, \widehat{\mathbf{w}}^s)$  in every period. Combining the two statements shows that in fact  $\check{\mathbf{V}} = \widehat{\mathbf{V}}$ , and this completes the proof. Note that this argument also establishes that the response profile we have constructed is a strategy profile: under the responses used, we can write formally the dependence of actions on all prior signals, and verify using the observations on decay of dependence across time that the formula is summable and hence defines unique actions.

### APPENDIX C. PROOF OF THEOREM 1

**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. Take two arbitrary agents  $i$  and  $j$ . Recall that both  $r_{i,t}$  and  $r_{j,t}$  have mean  $\theta_{t-1}$ , because each is an unbiased estimate<sup>38</sup> of  $\theta_{t-1}$ ; 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  $\Phi$  is well-defined and induces a map  $\tilde{\Phi}$  on  $\mathcal{W}$ .

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

1. for any pair of distinct agents<sup>39</sup>  $i \in G_n^k$  and  $j \in G_n^{k'}$ ,

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

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

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

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

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

This is the space of covariance matrices such that each covariance is split into two parts. Considering (1) first,  $\psi_{kk'}$  is an effect that depends only on  $i$ 's and  $j$ 's network types, while

<sup>38</sup>This is because it is a linear combination, with coefficients summing to 1, of unbiased estimates of  $\theta_{t-1}$ .

<sup>39</sup>Throughout this proof, we abuse terminology by referring to agents and nodes interchangeably when the relevant  $t$  is clear or specified nearby.

$\zeta_{ij}$  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  $\mathcal{W}_{\bar{\psi}, \bar{\zeta}}$  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  $\bar{\psi}$  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 equilibrium outcome has a social error covariance matrix  $A_t$  in the set  $\mathcal{W}_{\bar{\psi}, \bar{\zeta}}$ . In particular,  $\text{Var}(r_{i,t} - \theta_{t-1})$  becomes arbitrarily small in this limit. In our constructions, the  $\zeta_{ij}$  (resp.,  $\zeta_i$ ) terms will be set to much smaller values than the  $\psi_{kk'}$  (resp.,  $\psi_k$ ) terms, because group-level covariances are more predictable and less sensitive to idiosyncratic realizations.

C.2.2. *Approach to showing that  $\mathcal{W}_{\bar{\psi}, \bar{\zeta}}$  contains an equilibrium.* To show that the equilibrium outcome has (a.a.s.) a social error covariance matrix  $A_t$  in the set  $\mathcal{W}_{\bar{\psi}, \bar{\zeta}}$ , the plan is to construct a set so that (a.a.s.)  $\bar{\mathcal{W}} \subset \mathcal{W}_{\bar{\psi}, \bar{\zeta}}$  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  $\mathcal{W}_{\bar{\psi}, \bar{\zeta}}$ .

To construct the set  $\bar{\mathcal{W}}$ , we will fix a positive constant  $\beta$  (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  $\lambda$ ,

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

This will allow us to establish that (a.a.s.)  $\bar{\mathcal{W}} \subset \mathcal{W}_{\bar{\psi}, \bar{\zeta}}$  and  $\tilde{\Phi}(\bar{\mathcal{W}}) \subset \bar{\mathcal{W}}$ , with  $\bar{\psi}$  and  $\bar{\zeta}$  being arbitrarily small numbers.

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

**Lemma 1.** *There is a function  $\underline{\lambda}(\beta) \geq 1$  such that the following holds. For all large enough  $\beta$  and all  $\lambda \geq \underline{\lambda}(\beta)$ , for  $n$  sufficiently large we have  $\tilde{\Phi}(\mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n}, \frac{\lambda}{n}}$  with probability at least  $1 - \frac{1}{n}$ .*

**Lemma 2.** *For all large enough  $\beta$ , for  $n$  sufficiently large,  $\tilde{\Phi}^2(\mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}) \subset \mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}}$ , with probability at least  $1 - \frac{1}{n}$ .<sup>40</sup>*

<sup>40</sup>The notation  $\tilde{\Phi}^2$  means the operator  $\tilde{\Phi}$  applied twice.

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  $\overline{\mathcal{W}} = \mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}} \cup \tilde{\Phi}(\mathcal{W}_{\frac{\beta}{n}, \frac{1}{n}})$  is mapped to a subset of itself by  $\tilde{\Phi}$ , 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  $\mathcal{W}_{\psi, \zeta}$ .* 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  $\mathcal{W}_{\psi, \zeta}$ , to yield new covariance matrices. Thus, we will prove these lemmas by studying two periods of updating. The analysis will come in five steps.

**Step 1: No-large-deviations (NLD) networks and the high-probability event.**

Step 1 concerns the “with high probability” part of the lemmas. In the entire argument, we condition on the event of a *no-large-deviations (NLD)* network realization, which says that certain realized statistics in the network (e.g., number of paths between two nodes) are close to their expectations. The expectations in question depend only on agents’ types. Therefore, on the NLD realization, the realized statistics do not vary much based on which exact agents we focus on, but rather depend 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 describe how the covariance matrix  $A_t$  of social signal errors changes under updating. How this works is determined by the “basic” updating map  $\Phi$ , 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  $\mathcal{W}_{\psi, \zeta}$ , 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  $A_t \in \mathcal{W}_{\psi, \zeta}$  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  $\lambda$  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. What we have established so far about weights 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  $\tau$  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  $\tau$  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  $\tau$  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  $\tau$  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  $\tau$  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 Chebyshev's inequality, 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  $\psi = \beta/n$  for a sufficiently large constant  $\beta$ , and let  $\zeta = 1/n$ .

**Lemma 3.** *Suppose that in period  $t$  the matrix  $A = A_t$  of covariances of social signals satisfies  $A \in \mathcal{W}_{\psi,\zeta}$  and all agents are optimizing in period  $t + 1$ . Then there is a  $\gamma$  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  $\psi_{kk'} = 0$ ,  $\psi_k = 0$  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  $\gamma_1$  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  $\mathbf{C}_i(0)^{-1}$ . To do so, we will evaluate the formula

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

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

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

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.<sup>41</sup>

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

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

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

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

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

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

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

for some real  $c_{ij}$  and  $d_i$ .

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

$$(C.2) \quad \begin{aligned} & \sum_k (a_{ik}1_{r_i \times r_k} + b_i \delta_{ik} I_{r_i})(c_{kj}1_{r_k \times r_j} + d_k \delta_{kj} I_{r_k}) = \\ & (a_{ij}d_j + \sum_k (a_{ik}r_k + \delta_{ik}b_k)c_{kj}) 1_{r_i \times r_j} + b_i d_i \delta_{ij} I_{r_i}. \end{aligned}$$

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 (C.2) 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,

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

---

<sup>41</sup>We are very grateful to Iosif Pinelis for suggesting this argument.

where  $br := (b_1r_1, \dots, b_qr_q)$ . 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.  $\square$

Using Lemma 4 we can analyze the column sums of<sup>42</sup>

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

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 C.1 repeatedly in  $x$ , we obtain an expression for the  $k^{\text{th}}$  derivative in 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  $A \in \mathcal{W}_{\psi, \zeta}$ , 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  $\gamma_1/n$  for some  $\gamma_1$ . 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  $\mathcal{W}_{\psi, \zeta'}$ . 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  $\mathcal{W}_{\psi, \zeta}$ , then the covariance matrix in the

<sup>42</sup>Recall we wrote  $A(x) = A(0) + xM$ , and in (C.1) we expressed the derivative of  $f$  in  $x$  in terms of the matrix we exhibit here.

next period is in a somewhat larger neighborhood  $\mathcal{W}_{\psi, \zeta'}$ . 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  $\text{Cov}(r_{i,t} - \theta_{t-1}, r_{j,t} - \theta_{t-1}) = \psi_{kk'} + \zeta_{ij}$  into a term  $\psi_{kk'}$  depending only on the types of the two agents and a term  $\zeta_{ij}$ , and similarly for variances. To show the covariance matrix is contained in  $\mathcal{W}_{\psi, \zeta'}$ , we bound each of these terms suitably.

We begin with  $\zeta_{ij}$  (and  $\zeta_i$ ). 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  $\gamma_1$ . We can use this to bound the variation in covariances of  $r_{i,t+1}$  within types by  $\zeta'$ : 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  $\psi_{kk'}$  (and  $\psi_k$ ). It is sufficient to show that  $\text{Var}(r_{i,t+1} - \theta_t)$  is at most  $\psi$ . To do so, we will give an estimator of  $\theta_t$  with variance less than  $\beta/n$ , and this will imply  $\text{Var}(r_{i,t+1} - \theta_t) < \beta/n = \psi$  (recall  $r_{i,t+1}$  is the estimate of  $\theta_t$  given agent  $i$ 's social observations in period  $t+1$ ). Since this bounds all the variance terms by  $\psi$ , the covariance terms will also be bounded by  $\psi$  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  $G_n^k$ .<sup>43</sup> The basic idea will be that we can approximate  $\theta_t$  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. These bounds on the noise term follow from the assumption

<sup>43</sup>We use the notation  $\Omega(n)$  to mean greater than  $Cn$  for some constant  $C > 0$  when  $n$  is large.

that the covariance matrix of the  $r_{i,t}$  is in a neighborhood  $\mathcal{W}_{\psi,\zeta}$  and our analysis of variation in weights. Similarly

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

where the noise term has the same properties. Because  $\sigma_A^2 \neq \sigma_B^2$ , we can write  $\theta_t$  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  $\beta$  large enough such that this noise term has variance most  $\beta/n$  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  $A \in \mathcal{W}_{\psi,\zeta}$  and all agents are optimizing in periods  $t+1$  and  $t+2$ . Then there is a  $\gamma$  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.*

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

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

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 covariances<sup>44</sup> 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  $\zeta_{ij}$  and  $\zeta_i$  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  $\psi_{kk'}$  and  $\psi_k$  to be sufficiently small.

<sup>44</sup>We take this term to refer to variances, as well.

**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,

$$(C.3) \quad c_{ij',t} = \sum_j \frac{W_{ij,t+2}}{1 - w_{i,t+2}^s} W_{jj',t+1} (1 - w_{j',t}^s).$$

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.** *There exists  $\gamma$  such that 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  $[1 - \gamma/n, 1 + \gamma/n]$ .*

*Proof.* Fix  $i$  and  $j'$ . For each network type  $k''$  and signal type  $s$ , consider the number of agents  $j$  of network type  $k''$  and signal type  $s$  who are observed by  $i$  and who observe  $j'$ . This number varies by at most a factor  $\zeta^2$  as we change  $i$  and  $j'$ , preserving signal and network types. For each such  $j$ , the contribution of that agent's action to  $c_{ij',t}$  is (recalling (C.3))

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

By applying Lemma 3 repeatedly, we can choose  $\gamma_1$  such that each of these contributions varies by at most a factor of  $\gamma_1/n$  as we change  $i$  in  $G_k$  and  $j'$  in  $G_{k'}$ . Thus,  $c_{ij',t}$  is a sum of terms which vary by at most a multiplicative factor of  $\gamma_1/n$  as we change  $i$  and  $j'$  preserving signal and network types. If we can show that the sum of the absolute values of these terms is bounded, then it will follow that  $c_{ij',t}$  varies by at most a multiplicative factor of  $\gamma/n$  for some  $n$ . This bound on the sum of absolute values follows from the calculation of weights in the proof of Lemma 3.  $\square$

**Step 5(b):** We first show that fixing the values of  $\psi_{kk'}$  and  $\psi_k$  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  $\zeta$ . From the formula above, we observe that we can decompose  $r_{i,t+2} - \theta_{t+1}$  as a linear combination of three mutually independent groups of terms:

- (i) signal error terms  $\eta_{j,t+1}$  and  $\eta_{j',t}$ ;
- (ii) the errors  $r_{j',t} - \theta_t$  in the social signals from period  $t$ ; and
- (iii) changes in state  $\nu_t$  and  $\nu_{t+1}$  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 we are conditioning on the events  $E_2$  and  $E_4$  defined in Section C.1. 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  $\nu_t$  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  $\psi$  and the variation in these terms is bounded above by  $\zeta$ . We conclude that the variation from these terms has order  $1/n^2$ .

**(iii) Innovations:** Finally, we consider the contribution of the innovations  $\nu_t$  and  $\nu_{t+1}$ . We treat  $\nu_{t+1}$  first. We must show that any two agents of the same types place the same weight on the innovation  $\nu_{t+1}$  (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  $\psi_{kk'}$  and a lower-order term which can be included in  $\zeta_{ii'}$ .

The weight an agent places on  $\nu_{t+1}$  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  $W_i^A = \sum_{j:\sigma_j^2=\sigma_A^2} \frac{W_{ij,t+2}}{1-w_{i,t+2}^s}$  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 define  $\kappa_{j,t+1}^2 = \text{Var}(r_{j,t+1} - \theta_t)$  to be the error variance of the social signal. 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  $\psi_{k'k''}$ ,  $\psi_{k'}$ , 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  $\nu_{t+1}$ . Because we have included any innovations prior to  $\nu_t$  in the social signals  $r_{j',t}$ , to complete Step 5(b) we need only show the weight on  $\nu_t$  depends only on the network type  $k$  of an agent.

The analysis is a simpler version of the analysis of the weight on  $\nu_{t+1}$ . 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  $\nu_{t+1}$  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  $j'$  though a given type and the weights on each such path cannot vary too much within types, the same argument shows that this sum depends only on  $k$  and  $O(1/n^2)$  terms.

Step 5(b) is complete.

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

#### APPENDIX D. MODEL WITH A STARTING TIME

In introducing the model (Section 2), we made the set of time indices  $\mathcal{T}$  equal to  $\mathbb{Z}$ , the set of all integers. Here we study the variant with an initial time period,  $t = 0$ : thus, we take  $\mathcal{T}$  to be  $\mathbb{Z}_{\geq 0}$ , the nonnegative integers.

We let  $\theta_0$  be drawn according to the stationary distribution of the state process:  $\theta_0 \sim \mathcal{N}\left(0, \frac{1}{1-\rho}\right)$ . After this, the state random variables  $\theta_t$  satisfy the AR(1) evolution

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

where  $\rho$  is a constant with  $0 < |\rho| < 1$  and  $\nu_{t+1} \sim \mathcal{N}(0, \sigma_\nu^2)$  are independent innovations. Actions, payoffs, signals, and observations are the same as in the main model, with the obvious modification that in the initial periods,  $t < m$ , information sets are smaller as there are not yet prior actions to observe.<sup>45</sup> To save on notation, we write actions as if agents had an improper prior, understanding that the adjustment for actions taken under the natural prior  $\theta_t \sim \mathcal{N}\left(0, \frac{1}{1-\rho}\right)$  is immediate.

In this model, there is a straightforward prediction of behavior:

**Fact 1.** *In the model with  $\mathcal{T} = \mathbb{Z}_{\geq 0}$ , there is a unique Nash equilibrium, and it is in linear strategies. The initial generation ( $t = 0$ ) plays a linear strategy based on private signals only. In any period  $t > 0$ , given linear strategies from prior periods, players' best responses are linear. For time periods  $t > m$ , we have*

$$\mathbf{V}_t = \Phi(\mathbf{V}_{t-1}).$$

This fact follows from the observation that the initial ( $t = 0$ ) generation faces a problem of forming a conditional expectation of a Gaussian state based on Gaussian signals, so their optimal strategies are linear. From then on, the analysis of Section 3.1 characterizes best-response behavior inductively.

Our main purpose in this section is to give analogues of the main results on learning in large networks. We use the same definition of an environment—in terms of the distribution of networks and signals—as in Section 4.2. For simplicity, we work with  $m = 1$ , though the arguments extend straightforwardly.

The analogue of Theorem 1 is:

**Theorem 2.** *Consider the  $\mathcal{T} = \mathbb{Z}_{\geq 0}$  model and let  $\varepsilon > 0$ . If an environment satisfies signal diversity, asymptotically almost surely  $G_n$  is such that the  $\varepsilon$ -perfect aggregation benchmark is achieved in the unique Nash equilibrium at all times  $t \geq 1$ .*

We sketch the proof, which uses the material we developed in Appendix C. We define  $A_t$  as in that proof (Section C.1). Take a  $\beta > 0$ , to be specified later, and consider

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

<sup>45</sup>The actions for  $t < 0$  can be set to arbitrary (commonly known) constants.

First, for large enough  $\beta$ , we have that  $A_1 \in \overline{\mathcal{W}}$ : In the unique Nash equilibrium, at  $t = 1$ , agents simply take weighted averages of their neighbors' signals, weighted by their precisions. Formally, this is implied by Lemma 2 for  $\beta$  sufficiently large. Second, we use the previously established fact (recall Section C.2.2) that  $\tilde{\Phi}(\overline{\mathcal{W}}) \subset \overline{\mathcal{W}}$  to deduce that  $A_t \in \overline{\mathcal{W}}$  at all future times. Finally, we use Lemma 1 to deduce that for large enough  $n$  we have  $\overline{\mathcal{W}} \subseteq \mathcal{W}_{\bar{\psi}, \bar{\zeta}}$  with  $\bar{\psi}$  and  $\bar{\zeta}$  being arbitrarily small numbers.

Without signal diversity, bad learning can occur forever, in the unique equilibrium. The analogue of Proposition 2 is immediate. In graphs with symmetric neighbors,  $\Phi$  is a contraction, and so iteration of it arrives at the unique fixed point, and thus a learning outcome far from the benchmark.

Even when a random graph creates some structural diversity, the same result obtains. Let  $(G_n)_{n=1}^\infty$  be a sequence of directed or undirected random networks, with  $G_n$  having  $n$  nodes, with any pair of distinct nodes linked (i.i.d.) with positive probability  $p$ . We assume all signal variances are equal to  $\sigma^2$  and that  $m = 1$ . The analogue of Proposition 3 is now:

**Proposition 6.** *Consider the  $\mathcal{T} = \mathbb{Z}_{\geq 0}$  model and maintain the assumptions in the previous paragraph. There exists  $\varepsilon > 0$  such that asymptotically almost surely in the unique Nash equilibrium, at all times  $t > 1$  the  $\varepsilon$ -perfect aggregation benchmark is not achieved by any agent  $i$ .*

The randomness in the statement is, as usual, in the realization of  $G_n$ .

The proof of this result is analogous to that of Proposition 3, in Section E.1 of the Online Appendix. Step 1 is to observe that, in period  $t = 1$ , for any  $\varepsilon' > 0$ , with high probability all agents achieve  $\varepsilon'$ -perfect aggregation for  $n$  large. Step 2 is to argue that, conditioning on this event, after a finite number of periods  $t = 2, \dots, T$  (in which  $\varepsilon$ -perfect aggregation fails for  $\varepsilon$  sufficiently small), the matrix  $A_T$  lies in a small neighborhood of  $\widehat{\mathbf{V}}^{sym}(n)$ . Step 3, which is identical to Step 5 of the Proposition 3 argument, shows that  $\Phi$  maps this neighborhood to itself, and thus  $A_t$  remains in that neighborhood in future times.

It remains only to fill in the details for Step 2. This is done as follows. Consider first an artificial process, where all  $t = 1$  agents know  $\theta_0$  exactly, and in addition observe a private signal about  $\theta_1$ . Their updating is given by

$$a_{i,1} = w^{(1)}\theta_0 + (1 - w^{(1)})s_{i,1}.$$

Now,  $t = 2$  agents each observe a neighborhood of approximately  $(n - 1)p$  others taking such actions, and a private signal. Thus, they essentially observe

$$R_1 = w^{(1)}\theta_0 + (1 - w^{(1)})\theta_1$$

as well as a private signal. Their updating is given by

$$a_{i,2} = w^{(2)}R_1 + (1 - w^{(2)})s_2.$$

And so on. The sequence of weights can be computed explicitly, and it can be seen that they converge to  $\widehat{\mathbf{V}}^{sym}(n)$ . Then, by arguments analogous to those in the proof of Proposition 3, it is shown that the noise in exact neighborhood realizations does not change the sequence of weights much. This concludes the argument.

## APPENDIX E. REMAINING PROOFS (ONLINE APPENDIX)

**E.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 7.* We will show that when the network satisfies the condition in the proposition statement,  $\Phi$  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  $\theta_{t-1}$  based on the period  $t - 1$  actions of agents in  $N_i$ —and let  $\kappa_{i,t}^2$  be the variance of  $r_{i,t} - \theta_{t-1}$ .

We claim that  $\Phi$  induces a map  $\widetilde{\Phi}$  on the space of variances  $\kappa_{i,t}^2$ , which we denote  $\widetilde{\mathcal{V}}$ . We must check the period  $t$  variances  $(\kappa_{i,t}^2)_i$  uniquely determine all period  $t + 1$  variances  $(\kappa_{i,t+1}^2)_i$ : The variance  $\mathbf{V}_{ii,t}$  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  $\widetilde{\Phi}$  is a contraction on  $\widetilde{\mathcal{V}}$  in the sup norm.

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

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

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

$$(E.2) \quad a_{i',t} = b_0 \cdot r_{i,t} + b_{i'} \cdot s_{i',t}$$

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

$$(E.3) \quad r_{j,t+1} = c_0 \cdot (\rho r_{i,t}) + \sum_{i'} c_{i'} s_{i',t}$$

and therefore (taking variances of both sides)

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

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

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.** *Suppose  $j$  has symmetric neighbors. Then agent  $j$ 's social signal places non-negative weight on a neighbor'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  $r_{j,t+1}$  puts negative weight on some observed action—say the action  $a_{k,t}$  of agent  $k$ . We want to check that the covariance of  $r_{j,t+1} - \theta_t$  and  $a_{k,t} - \theta_t$  is negative. Using (E.2) and (E.3), we compute that

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

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  $r_{j,t+1}$  puts negative weight on  $a_{k,t}$ . So the expression above is negative. Therefore, it follows from the usual Gaussian Bayesian updating formula that the best estimator of  $\theta_t$  given  $r_{j,t+1}$  and  $a_{k,t}$  puts positive weight on  $a_{k,t}$ . However, this is a contradiction: the best estimator of  $\theta_t$  given  $r_{j,t+1}$  and  $a_{k,t}$  is simply  $r_{j,t+1}$ , because  $r_{j,t+1}$  was defined as the best estimator of  $\theta_t$  given observations that included  $a_{k,t}$ .  $\square$

We now complete the proof of Lemma 7.

*Proof.* Now, for the upper bound  $c_0 \leq 1$ , the idea is that  $r_{j,t+1}$  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  $r_{j,t+1}$  is a linear combination of actions  $\rho a_{i',t}$  for  $i' \in N_j$ , with coefficients summing to 1. The only way the coefficient on  $\rho r_{i,t}$  in  $r_{j,t+1}$  could be at least 1 would be if some of these coefficients on  $\rho a_{i',t}$  were negative and the estimator  $r_{j,t+1}$  placed greater weight on actions  $a_{i',t}$  which placed more weight on  $r_{i,t}$ .

Applying the formula (E.1) for  $a_{i',t}$ , we see that the coefficient  $b_0$  on  $\rho r_{i,t}$  is less than 1 and increasing in  $\sigma_{i'}$ . On the other hand, it is clear that the weight on  $a_{i',t}$  in the social signal  $r_{j,t+1}$  is decreasing in  $\sigma_{i'}$ : more weight should be put on more precise individuals. So in fact the estimator  $r_{j,t+1}$  places less weight on actions  $a_{i',t}$  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 one on the estimator  $\rho r_{i,t-1}$ . Agent  $j$ 's social signal  $r_{j,t+1}$  is a sum of these agents' actions with coefficients summing to 1 and satisfying the monotonicity property above. We conclude that the coefficient on  $\rho r_{i,t}$  in the expression for  $r_{j,t+1}$  is at most one.  $\square$

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:

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

Suppose the  $\varepsilon$ -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  $\varepsilon$  sufficiently small the social signal  $r_{i,t}$  is bounded away from a perfect estimate of  $\theta_{t-1}$ . This gives a contradiction.  $\square$

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

Choose  $\sigma^2$  large enough such that  $A > 1$ . To see that we can do this, note that as  $\sigma^2$  grows large, the weight each agent places on their private signal vanishes. So the weight on signals from at least  $k$  periods ago approaches one for any  $k$ . Taking  $\sigma^2$  such that this holds for  $k$  sufficiently large, we have  $A > 1$ .

Now suppose that we increase  $\sigma_1^2$  to  $\infty$ . 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 converge to  $(1 + \sigma^{-2})^{-1}$ . By our choice of  $\sigma^2$ , this gives a Pareto improvement. We can see by continuity that the same argument holds for  $\sigma_1^2$  finite but sufficiently large.

**E.3. Proof of Proposition 3.** We begin by outlining the argument, and then carry out the steps. In Step 1, we construct a symmetric version of the Erdos-Renyi network and show there exists a symmetric equilibrium  $\widehat{V}^{sym}(n)$  on this symmetric network. In Step 2, we show variances and covariances (the entries of  $A_t$ , in the notation of Section C.1) at the equilibrium  $\widehat{V}^{sym}(n)$  converge to  $V^\infty$  and  $Cov^\infty$ . The remainder of the proof shows there is an equilibrium on  $G_n$  near  $\widehat{V}^{sym}(n)$ . 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  $\Phi$  maps a small neighborhood of  $\widehat{V}^{sym}(n)$  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  $V(n)_{ij}$  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  $\Phi^{sym}$  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  $V^\infty$  and  $Cov^\infty$ . 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:

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

So the variance of the social signal about  $\theta_{t-1}$  is

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

Thus the covariance of any two distinct agents solves

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

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  $Cov^\infty$ . 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  $V^\infty$ .

**Step 3:** We will show that when  $\zeta = \frac{1}{n}$ , the updating map  $\Phi$  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

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

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  $[(1 - \frac{1}{n^2})pn, (1 + \frac{1}{n^2})pn]$ .

$E_2$  : 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  $[(1 - \frac{1}{n^2})p^2n, (1 + \frac{1}{n^2})p^2n]$ .

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  $\gamma$  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

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

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

Similarly

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

The term

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

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  $\mathbf{V}_{ii',t+1}$  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  $\varepsilon$ -perfect aggregation benchmark is not achieved at this equilibrium for any  $n$ .

**E.4. Proof of Proposition 4.** For each agent  $i$ , we can write

$$a_{i,t} = w_i^s s_{i,t} + \sum_j W_{ij} \rho a_{j,t-1} = w_i^s s_{i,t} + \sum_j W_{ij} (\rho w_j^s s_{j,t} + \sum_{j'} W_{jj'} \rho a_{j',t-2}).$$

Because we assume  $w_i^s < \bar{w} < 1$  and  $w_j^s < \bar{w} < 1$  for all  $j$ , the total weight  $\sum_{j,j'} W_{ij} W_{jj'} \rho$  on terms  $a_{j',t-2}$  is bounded away from zero. Because the error variance of each of these terms is greater than 1, this implies agent  $i$  fails to achieve perfect aggregation.

**E.5. Proof of Proposition 5.** We prove the following statement, which includes the proposition as special cases.

**Proposition 7.** *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.*

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

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

Suppose we are at the initial steady state until time  $t$ , but in period  $t$  and all subsequent periods agents instead use weights  $\overline{\mathbf{W}}$ . These weights give an alternate updating function  $\overline{\Phi}$  on the space of covariance matrices. Because the weights  $\overline{\mathbf{W}}$  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  $\epsilon$  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  $\delta$ . Then each coordinate  $[\Phi(\mathbf{V})]_{ii}$  is linear with coefficient at least  $\delta^2$  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

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

Define new weights  $\bar{W}$  as in the proof of case (1). Because agent  $k$  is naive and the summation  $\sum_{j \in N_i} W_{kj} \rho a_{j,t-1}$  has at least two terms, she believes the variance of this summation is smaller than its true value. So marginally increasing the weight on  $s_{k,t}$  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).

## APPENDIX F. NAIVE AGENTS (ONLINE APPENDIX)

In this section we provide rigorous detail for the analysis given in 5.1. We will describe outcomes with two signal types,  $\sigma_A^2$  and  $\sigma_B^2$ .<sup>46</sup> We use the same random network model as in Section 4.3 and assume each network type contains equal shares of agents with each signal type.

We can define variances

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

where

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

Naive agents' equilibrium variances converge to these values.

**Proposition 8.** *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  $\delta$  of  $V_A^\infty$  and  $V_B^\infty$ .*
- (3) *There exists  $\varepsilon > 0$  such that asymptotically almost surely the  $\varepsilon$ -perfect aggregation benchmark is not achieved, and when  $\sigma_A^2 = \sigma_B^2$  asymptotically almost surely all agents' variances are larger than  $V^\infty$ .*

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  $\infty$  and any sequence of equilibria, the equilibrium action variances of all agents converge to  $V_A^\infty$ .

**F.1. Proof of Proposition 8.** We first check that there is a unique naive equilibrium. As in the Bayesian case, covariances are updated according to equations ??:

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

<sup>46</sup>The general case, with many signal types, is similar.

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  $\Phi^{naive}$ ) 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  $V_A^\infty$  and  $V_B^\infty$ , 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  $V_A^\infty$  and  $V_B^\infty$ .

To estimate  $\theta_{t-1}$ , a naive agent combines observed actions from the previous period with weight proportional to their precisions  $\sigma_A^{-2}$  or  $\sigma_B^{-2}$ . The naive agent incorrectly believes this gives an almost perfect estimate of  $\theta_{t-1}$ . So the weight on older observations vanishes as  $n \rightarrow \infty$ . The naive agent then combines this estimate of  $\theta_{t-1}$  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  $\theta_{t-1}$  is approximately:

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

The actual variance of this estimate converges to:

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

where  $Cov_{AA}^\infty$  is the covariance of two distinct agents of signal type  $A$  and  $Cov_{BB}^\infty$  and  $Cov_{AB}^\infty$  are defined similarly.

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

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

Plugging into F.2 we obtain

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

Using this formula, we can check that the limits of agent variances in  $\widehat{\mathbf{V}}^{sym}$  match equations F.1.

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

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

Finally, we compare the equilibrium variances to perfect aggregation and to  $V^\infty$ . 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  $\sigma^2$ . 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  $Cov^\infty$  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  $Cov^\infty > 0$ .

APPENDIX G. SOCIALLY OPTIMAL LEARNING OUTCOMES WITH NON-DIVERSE SIGNALS  
(ONLINE APPENDIX)

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 9.** *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  $\varepsilon$ -perfect aggregation benchmark is achieved.*

*Proof.* An agent with a social signal equal to  $\theta_{t-1}$  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,

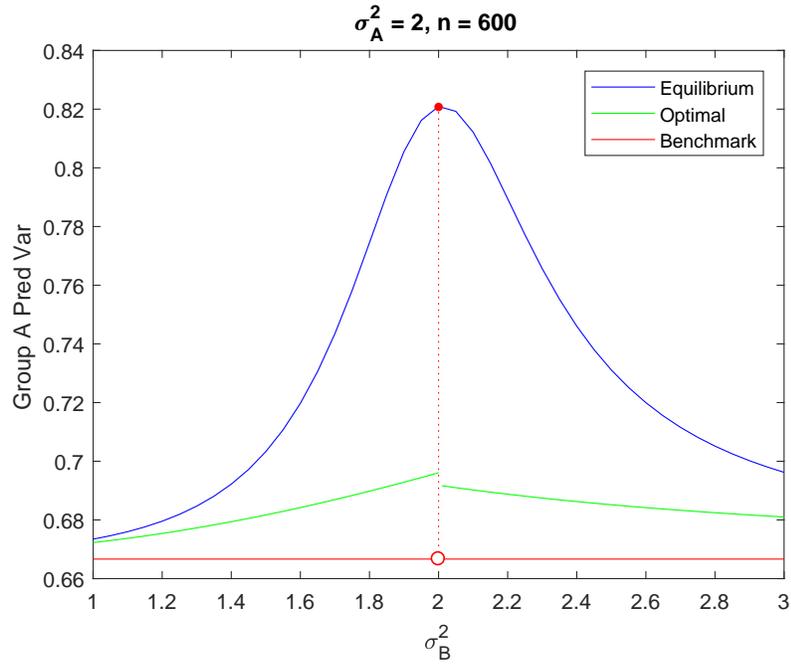
$$\frac{1}{\lfloor n/2 \rfloor} \sum_{j=1}^{\lfloor n/2 \rfloor} a_{j,t-1} = w_A^s \theta_{t-1} + (1 - w_A^s) r_{t-1} + O(n^{-1/2}),$$

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

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  $\delta$  sufficiently small and then  $n$  sufficiently large, we find that  $\varepsilon$ -perfect aggregation is achieved.  $\square$

In Figure G.1, we conduct the same exercise as in Figure 4.1 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

FIGURE G.1. Social Planner and Bayesian Learning



socially optimal weights agents learn very well, and heterogeneity in signal variances only has a small impact.