Empiricist Learning Rules on Social Networks: Convergence and Quality of Information Aggregation

Thi Mai Anh Nguyen*

October 8, 2022

Abstract

This paper proposes a novel learning model on social networks that captures settings where individuals interact frequently on multiple, relatively short-lived topics. In this model, each period features a new draw of nature and multiple rounds in which information arrives, gets aggregated, and diffuses through network links. The repetitive nature of interactions across periods allows for a separation between learning about the environment and aggregating information about the current state. A class of empiricist learning rules achieve convergence of learning on all networks. On clique trees, these learning rules further achieve strong efficiency in information aggregation. The paper also presents a converse to the positive efficiency result and identifies distinct reasons why efficiency is hard to obtain in general circumstances, even though convergence of learning holds generally.

1 Introduction

People interact in complicated networks while often having little understanding about the structure of these networks. For example, using network data and surveys from 75 villages in India, Breza et al. (2018) find that 46% of respondents are not certain enough to elicit a guess about whether two given individuals have financial, social or informational links.

*Department of Economics, Massachusetts Institute of Technology. Email: anhng@mit.edu. I am particularly indebted to Glenn Ellison and Stephen Morris. I am also grateful to Krishna Dasaratha, Drew Fudenberg, Giacomo Lanzani, Claudio Mezzetti, Alex Wolitzky and participants in the MIT Theory Lunch for helpful comments and suggestions.

Furthermore, when the respondents do make guesses, only 37% of them are correct. The reality that knowledge about the network structure is local poses a challenge to information aggregation, since individuals may fail to correctly assess the reliability of information from their neighbors.

Previous papers in the literature that relax common knowledge of the full network have abstracted from learning.[1] An open question is whether individuals with limited knowledge of their complicated environment, through repeated interactions, could learn to aggregate information from their neighbors optimally. Furthermore, how does such learning affect the quality of information aggregation by the network as a whole?

To answer these questions, I propose a novel model of information diffusion and aggregation in settings like Facebook or Twitter. Three key features of these social platforms motivate my modeling choices. First, individuals interact frequently on multiple, relatively short-lived topics. Second, the diffusion of information depends on an underlying network of individuals, who are quick to respond to the posts of their friends (Facebook) or of those they follow (Twitter). Third, the underlying network is rich, making it impossible for an individual to perfectly know others' links. My model captures these features in a stylized way. It features in each period a new draw of nature and random, asynchronous timing of signal arrivals and decision making. Local knowledge of the network structure and local observability are assumed. Despite individuals' limited knowledge and observability, a class of empiricist learning rules achieve convergence of learning on all networks. The induced long-run play is then shown to achieve strong efficiency on clique trees, tree-like networks where a node can be replaced by a set of fully linked individuals. In general circumstances, however, there are systematic ways in which misalignment of interests arises, challenging efficiency.

My model departs from the literature in two main ways:

First, it has two layers of timing that separate learning about the environment from aggregating information about the current state. The infinite number of periods put a repetitive structure on the interactions between each individual and her environment, facilitating learning about the environment. The multiple rounds within each period allow information about a newly drawn state of the world to spread through network links. Together with the

[1] Most papers in the social learning literature maintain common knowledge of the network topology and focus on Perfect Bayesian equilibria. DeGroot-style learning papers, for example, relax these common knowledge assumptions but use heuristics of updating rules.

network structure, the random arrivals of newly drawn private signals determine the order of play and the set of actions observed by each individual in each period. This means that each period is a game of information aggregation where individuals know little about how information from different sources have diffused throughout the network before reaching them.

Second, each individual in my model faces vast uncertainty about the underlying environment. This includes objective uncertainty, captured by the network structure, the distribution of private signals and their arrivals, and strategic uncertainty, encoded in how other individuals map their private observations to actions. Specifically, each individual only knows the set of her neighbors, her own private signal as well as her own belief-updating rule and decision rule. Moreover, others' belief-updating rules and decision rules are treated as primitives of the learning environment rather than strategic choices. This paper's objective is to find a class of learning rules whose performance is robust to individuals' vast uncertainty about their environment.

One such class is the class of *empiricist learning rules*. These are learning rules that believe in stationary environments and asymptotically believe in the empirical distribution of observations. Empiricist learning rules rest on two ideas. First, to optimally aggregate information in stationary environments, an individual only need to learn the stationary distributions of her observation conditional on each state of the world. In other words, the collection of these stationary conditional distributions, referred to as a *local model*, is the key object of learning for her. Second, since each individual possesses an *independent* and *informative* signal, there is a one-to-one map between an individual's local model and the unconditional distribution of her observation. This further reduces the object of learning to the unconditional distribution of an individual's observation, which, in stationary environments, can be learned asymptotically from its empirical counterpart. In Proposition 1, I show that if an individual's environment is indeed stationary, empiricist learning rules ensure asymptotic learning of her true local model and thus, asymptotic optimality.

What if all individuals on the network actively learn with empiricist learning rules, rendering the underlying environment nonstationary? In which cases will play converge and information aggregation be socially optimal?

Theorem 1 shows that convergence of beliefs and of play holds in all environments where individuals adopt empiricist learning rules and "smooth" optimal decision rules.

In each round, as the play of those neighbors that acted before an individual converges, the local environment of that individual becomes approximately stationary. Empiricist learning rules then ensure convergence of her belief about her local environment, in that round, and smooth decision rules translate convergence of beliefs into convergence of play. Crucially, empiricist learning rules separate learning in different rounds, thus avoiding potential mislearning in later rounds to contaminate learning in earlier rounds.

In fact, the convergence result in my setting extends to richer settings, suggesting that adaptations of my model could provide a framework for studying other games of information aggregation. This framework starts with building a model where the game of interest is played repeatedly. This model is then analyzed in two steps. First is to achieve convergence of learning about local models, or elements of the environment that are relevant to each individual's decision making. Second is to evaluate the per-period game where each individual knows her local model perfectly. Compared to a direct analysis of the game of interest, this framework generates predictions about individuals' play as the long-run outcome of a learning process. In doing so, it makes minimal assumptions on individuals' knowledge of their environment and of others' play.

The remainder of this paper focuses on the long-run quality of information aggregation in my setting when all individuals use empiricist learning rules and smooth decision rules.

Theorem 2 establishes a strong positive result: on *clique trees*, the long-run play induced by empiricist learning rules and smooth decision rules is *strongly efficient*, among all stationary plays that treat the two states symmetrically. Two properties of clique trees guarantee this result. First is *conditional independence of neighbors' actions*. Under symmetric treatment of the two states of the world, this property implies that the overall informativeness of all neighbors' actions is increasing in the informativeness of each neighbor's action. Second, clique trees ensure *local alignment of interests*: an individual maximizes the informativeness of her action to her neighbors exactly by optimally aggregating information from other neighbors of hers. The proof of Theorem 2 uses an inductive argument to show that under conditional independence of neighbors' actions, local alignment of interests leads to global alignment of interests. Since the long-run play induced by empiricist learning rules and smooth decision rules is individually optimal, it follows that this long-run play is strongly efficient; that is, it is optimal for every individual on the network.

What if the underlying network is not a clique tree, or the notion of efficiency is weakened, or the environment is not symmetric? Theorem 3, which is a weak converse to Theorem 2, illustrates the role of clique trees in achieving efficiency: on any network that is not a clique tree, substantive misalignment of interests arises for some diffusion process and signal structure. This theorem generalizes two examples of networks that are not clique trees, where either conditional independence of neighbors' actions or local alignment of interests fails. Example 1 shows that when there is conditional correlation between different information sources, it might be possible to reduce the informativeness of an individual source in a way that breaks the correlation and improves the informativeness of all sources combined. Example 2 illustrates how an individual might play optimally against a distribution that pools diffusions irrelevant to her neighbors, and thus fails to optimize the informativeness of her action to these neighbors. Moreover, forces that challenge strong efficiency are likely to also challenge Pareto efficiency. In two examples that extend Example 1 and Example 2, I show that Pareto efficiency fails because individuals can trade favors across rounds. Lastly, coordinated biases about the two states of the world can improve the informativeness of combined sources. This makes strong efficiency hard to achieve when asymmetric plays are considered for comparison.

In the setting of this paper, that individuals cannot trace the path of their information does not hurt learning but it generates misalignment of interests. In other words, empiricist learning rules solve the problem of local knowledge of the network structure, ensuring asymptotic learning and individual optimality. However, the gap between individual optimality and social optimality remains. This gap depends on all elements of the objective environment: the network, the diffusion process, and the quality of private signals.

My paper does not follow any previous work closely but it shares elements with different branches of literature. The modeling of each period is similar to recent models of social learning by Acemoglu et al. (2011) and Lobel and Sadler (2015) in that individuals play sequentially, after observing a random subset of others' actions. My model additionally generates randomness in the order of play, tying both the realized order of play and observation sets to the underlying network structure. Moreover, their papers focus on late decision makers in large populations, as standard in the social learning literature, while my paper concerns the quality of information aggregation of every individual.

The idea of exploiting the stationarity of the objective environment across periods to learn other individuals' persistent types is shared with Sethi and Yildiz (2016, 2019), who study advice-seeking networks where advice reveals information both about the current state of the world and about the perspective of the advice giver. In their papers, confoundedness of sources is absent and advice-seeking is an active choice. In contrast, individuals in my paper receives information passively on exogenous networks, but their types as perceived by others arise endogenously from how they learn and aggregate information.

My paper draws a connection between the network literature and the literature on learning in games, where learning rules are often motivated by stationary problems. Specifically, the class of empiricist learning rules defined in my paper relates to Fudenberg and Kreps (1993) in the idea that the long-run belief about a stationary distribution should be concentrated on its empirical frequency. Empiricist learning rules take a further step. They use properties that hold in all permissible network environments to back out the conditional distributions of observables from the unconditional distribution.

Lastly, the motivation from the empirical finding that network knowledge is local is shared with Li and Tan (2020). They study a model of misspecified learning, where each individual believes that the underlying network is only a subgraph including her neighbors and herself. In contrast, I take a robustness approach, constructing learning rules that work well when individuals acknowledge their uncertainty about the environment.

The rest of the paper is structured as follows. Section 2 describes the model. Section 3 defines empiricist learning rules and proves their asymptotic individual optimality in stationary environments (Proposition 1). Section 4 shows the convergence result (Theorem 1) and sketches out directions for extensions. Section 5 shows strong efficiency on clique trees (Theorem 2). Section 6 explores the challenges to achieving efficiency of information aggregation in general circumstances, including a formal converse to the positive efficiency result (Theorem 3). Section 7 reviews related literature and Section 8 concludes.

2 Model

My model features asynchronous arrivals of private signals and asynchronous actions within each period, when a new state of the world is drawn. In each period, an individual's observation updates both her belief about the underlying environment and her belief about the current state, the latter of which determines her action. This section focuses on clarifying the elements of an environment.

An environment consists of an objective environment, described in Subsection 2.1, and a profile of belief-updating rules and decision rules, described in Subsection 2.3. That is, belief-updating rules and decision rules are treated as primitives of the learning environment rather than strategic choices. Subsection 2.2 provides details on the timing of actions within each period. Subsection 2.4 defines convergence of beliefs and of play.

2.1 Objective environment

Fix an undirected network G. Without loss of generality, assume that G is connected. Denote by $N = \{1, ..., n\}$ the set of individuals on this network and by N_i the set of neighbors of individual i.

There are infinitely many periods, each of which has R rounds. In period t, a state of the world θ_t is drawn from $\{-1, 1\}$ with $\Pr(\theta_t = -1) = \Pr(\theta_t = 1) = 1/2$. Denote by $s_{1t}, ..., s_{nt}$ the private signals of individuals $1, ..., n$ in period t. Assume that these private signals are independent conditional on the realized state of the world, with the private signal of each individual i having conditional distribution $\Pr(s_{it} = \theta_t | \theta_t) = q_i$ for all $\theta_t \in \{-1, 1\}$. Assume that $q_i > 1/2$ so that private signals are informative. Refer to q_i as the quality of i's private signal. Private signals arrive to individuals at different rounds according to a signal-timing vector $\tau_t = (\tau_{1t}, ..., \tau_{nt})$ drawn from some distribution $H \in \Delta\left(\{1, ..., R, \infty\}^n\right)$, independently over time and independently of the state of the world; $\tau_{it} = r \in \{1, ..., R\}$ means that in period t a private signal arrives to i in round r and $\tau_{it} = \infty$ means that no private signal is generated for i that period. Assume that H has full support and the arrivals of private signals are independent across individuals, that is, $H = H_1 \times ... \times H_n$ where $H_i \in \Delta(\{1, ..., R, \infty\})$ for each i.

In each period t, individual i decides on an action $a_{it} \in \{-1, 1\}$.[2] In the context of social platforms, an action could be a post expressing one's view on the topic of that period and individuals want to express the correct view. Assume that an individual takes action at the end of the round in which she first receives some private signal or observes some actions of her neighbors (or both). Moreover, it takes one round for one's action to be observable to her neighbors. Assume further that once an individual chooses an action in a period, she does not pay attention to any private signal or actions chosen by her neighbors

[2]For $M \subseteq N$, write $a_{M,t}$ for the profile of actions chosen in period t by individuals in M. Write a_t for the profile of actions of all individuals in period t.

in later rounds of that period. As in the social learning literature, this assumption prevents the same piece of information spreading from one individual to her neighbors and then spreading back to her.[3]

A network G, a distribution H of the signal-timing vector and a vector $(q_i)_{i\in N}$ of signal qualities constitute an *objective environment*.

2.2 Timeline in each period

For a given network G, the diffusion of information and timing of actions in each period t are determined solely by the realized signal-timing vector τ_t. This vector induces for each i an *action round* $r_i(\tau_t)$ and an *observation set* $M_i(\tau_t)$. At the beginning of round $r_i(\tau_t)$, i first receives a private signal or observes the actions of some neighbors of hers. The observation set $M_i(\tau_t)$ is the set of individuals whom i hears from in round $r_i(\tau_t)$, including herself if she has received a private signal. At the end of round $r_i(\tau_t)$, individual i chooses an action.

Let $l(i,j)$ denote the distance between individuals i and j on the network. Given signal-timing vector τ, individual i's action round is

$$r_i(\tau) = \min\left\{\min_{j\in N}\{\tau_j + l(i,j)\}, R\right\}.$$

This reflects the assumption that each individual reacts to the first piece of information she receives and that it takes one round for information to travel across one link. If an individual does not receive any information by the last round, she will take action in the last round. The observation set of i is defined by

$$M_i(\tau) = (N_i \cap \{j : r_j(\tau) = r_i(\tau) - 1\}) \cup \{i\}^{\mathbf{1}\{r_i(\tau)=\tau_i\}}.$$

Figure 1 illustrates how a signal-timing vector determines the diffusion of information and timing of actions on one particular network. In the left panel is the undirected network and in the right panel is the diffusion of information on this network for $\tau = (1,3,2,5,2)$. In the latter, an arrow from $j \in N_i$ to i means that i observes j's action

[3] If feedback and updates of actions were instead allowed, then an individual might benefit from distorting her initial decision in ways that improve the quality of information she would receive before making the final decision.

and an arrow from i to herself means that she receives her private signal. The action rounds are $r_1(\tau) = 1, r_2(\tau) = r_3(\tau) = r_5(\tau) = 2, r_4(\tau) = 3$ and the observation sets are $M_1(\tau) = \{1\}, M_2(\tau) = \{1\}, M_3(\tau) = \{1, 3\}, M_4(\tau) = \{2, 5\}, M_5(\tau) = \{5\}$. In round 1, individual 1 observes her private signal and takes action. Her action takes one round to reach individuals 2 and 3, the latter of which receives a private signal in round 2. In round 2, individual 2 chooses an action based only on individual 1's action while individual 3 aggregates both the action of individual 1 and his private signal. Individual 5 observes her private signal in round 2 and chooses an action. In round 3, individual 4 sees the actions of both individuals 2 and 5 and chooses an action before seeing any private signal.

Figure 1: A realized diffusion with diffusion vector $\tau = (1, 3, 2, 5, 2)$

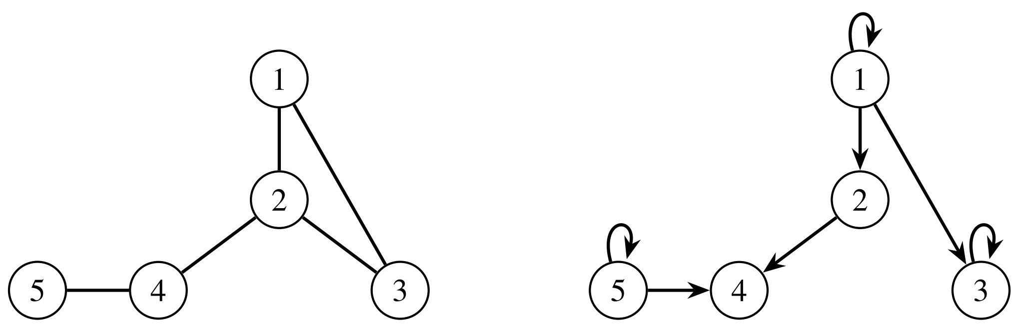

2.3 Belief-updating rules and decision rules

Recall that each individual is assumed to react to the first piece of information she receives in a period and not to pay attention to information arriving after she has chosen an action that period. Thus, an individual's observation at the beginning of her action round in period t includes the actions of some subset of her neighbors and her private signal, and the action round itself. Denote by h_{it} what i observes in period t before choosing an action. That is, $h_{it} = (a_{M_i(\tau_t)\setminus\{i\},t}, s_{it}, r_i(\tau_t))$ if $i \in M_i(\tau_t)$ and $h_{it} = (a_{M_i(\tau_t),t}, r_i(\tau_t))$ if $i \notin M_i(\tau_t)$. Her observation in period t after taking action that period, denoted by $\tilde{h}_{it}$, further includes the action taken. Let $\tilde{h}_i^t = (\tilde{h}_{i1}, ..., \tilde{h}_{it})$.

A *belief-updating rule* of i is a function $\beta_i : (\tilde{h}_i^{t-1}, h_{it}) \mapsto b_i \in \Delta(\{-1, 1\})$ that keeps track of i's belief about the current state of the world at the beginning of her action round each period. A *stationary belief-updating rule* $\bar{\beta}_i : h_{it} \mapsto b_i \in \Delta(\{-1, 1\})$ forms individual i's belief about the current state of the world only from her current observation.

A *decision rule* is a time-dependent function $\sigma_i : (b_i, t) \mapsto d_i \in \Delta(\{-1, 1\})$ that maps i's belief $b_i \in \Delta(\{-1, 1\})$ about the current state of the world to a distribution over feasible actions. A *stationary decision rule* is a time-independent function $\bar{\sigma}_i : b_i \mapsto d_i \in \Delta(\{-1, 1\})$. A natural candidate for consideration is the stationary decision rule that maximizes the subjective probability that i's action matches the current state of the world and breaks ties equally between the two states,

$$\bar{\sigma}_i^*(b_i) = \begin{cases} \delta_1 & \text{if } b_i(1) > 1/2, \\ \delta_{-1} & \text{if } b_i(-1) > 1/2, \\ (1/2, 1/2) & \text{otherwise.} \end{cases}$$

Here δ is the Dirac measure. Notice that this decision rule has a discontinuity at the point where belief splits exactly $(1/2, 1/2)$.

To ensure that convergence of beliefs leads to convergence of play, I consider *smooth decision rules*, which are technical modifications of $\bar{\sigma}_i^*$. A decision rule σ_i^η is a smooth decision rule if

$$\sigma_i^\eta(b_i, t)(\theta) = \frac{\exp(b_i(\theta)/\eta_{it})}{\exp(b_i(\theta)/\eta_{it}) + \exp(b_i(-\theta)/\eta_{it})}.$$

for some sequence $\eta = (\eta_{it})_t$, where $\eta_{it} \to 0$ as $t \to \infty$, and for all $\theta \in \{-1, 1\}$. That is, σ_i^η can be derived as the choice probability of a random utility model where the mean utility is the subjective belief and the logistic error decays with time. Since σ_i^η converges to $\bar{\sigma}_i^*$, smooth decision rules asymptotically optimize the probability of matching an individual's action to the true state of the world if the beliefs are correct. While letting the smoothing parameter decay to zero with time is not important for convergence, this approach allows cleaner statements on asymptotic efficiency than the alternative of keeping a constant smoothing parameter.[4]

An *environment* is a tuple $E = (G, H, (q_i)_{i \in N}, (\beta_i)_{i \in N}, (\sigma_i)_{i \in N})$ that combines an objective environment, with a profile of belief updating rules and decision rules.

[4] Smooth decision rules can also be motivated by the behavior of an individual with smooth ambiguity aversion who faces a decision problem repeatedly and whose uncertainty about the problem gets resolved over time. Such an individual tends to hedge, but her hedging tendency goes away with time. Battigalli et al. (2019) generate decaying hedging tendency from Bayesian learners with smooth ambiguity aversion. For the characterization of smooth ambiguity aversion, see Klibanoff et al. (2005).

2.4 Notions of convergence

To separate the stochasticity of play induced by the objective environment from the implications of different belief-updating rules and decision rules, I define the notion of system states. Fix an objective environment $(G, H, (q_i)_{i\in N})$. For each period t, let $\zeta_t = (\zeta_{it})_{i\in N}$ be a vector of independent standard uniform random variables, which are used instrumentally to capture the randomness in actions induced by mixing decisions. Define the *system state* at time t as the tuple $\omega_t = (\tau_t, s_t, \zeta_t)$, which includes a signal-timing vector, a vector of private signals, and a vector of instrumental draws. Note that system states are drawn independently over time from a stationary distribution, so any nonstationarity in the environment must be induced by the belief-updating rules and decision rules.

For a given objective environment, a profile $(\beta_i)_{i\in N}$ of belief-updating rules and a profile $(\sigma_i)_{i\in N}$ of decision rules define a history-dependent function $\psi : (\omega_t, \tilde{h}^{t-1}, t) \mapsto \{-1, 1\}^n$ that, for each given history $\tilde{h}^{t-1}$, maps the current system state to a profile of actions. Refer to such function as a *function of play* and each component function $\psi_i : (\omega_t, \tilde{h}^{t-1}, t) \mapsto \{-1, 1\}$ as a *function of i's play*. Take a period t and a history $\tilde{h}^{t-1}$, the action profile $a_t = \psi(\omega_t, \tilde{h}^{t-1}, t)$ is constructed as following:

- Let $r_0 = \min\{\tau_{1t}, .., \tau_{nt}\}$. For all i such that $\tau_{it} = r_0$, $h_{it} = (s_{it}, r_0)$. Set $a_{it} = -1$ if $\zeta_{it} \le \sigma_i(\beta_i(\tilde{h}_i^{t-1}, h_{it}))(-1)$, otherwise set $a_{it} = 1$.
- Inductively, for $r \in \{r_0+1, ..., R\}$ and i such that $r_i(\tau_t) = r$, $h_{it} = (a_{M_i(\tau_t)\setminus\{i\},t}, s_{it}, r)$ if $\tau_{it} = r$ and $h_{it} = (a_{M_i(\tau_t),t}, r)$ if $\tau_{it} > r$. Set $a_{it} = -1$ if $\zeta_{it} \le \sigma_i(\beta_i(\tilde{h}_i^{t-1}, h_{it}))(-1)$, otherwise set $a_{it} = 1$.

A *stationary function of play* is a history-independent function $\bar{\psi} : \omega \mapsto a \in \{-1, 1\}^n$, mapping each system state to a profile of actions.

Write $\theta^t = (\theta_1, ..., \theta_t)$ for the sequence of states of the world up to t and write $\omega^t = (\omega_1, ..., \omega_t)$ for the sequence of system states up to t. The objective environment of an environment E defines a probability space on the set of all $(\theta^\infty, \omega^\infty)$. Denote by $\Pr^E$ the corresponding probability measure.

An important question when studying a learning rule is whether, or under which conditions, it leads to convergence of beliefs and of play. This paper focuses on almost sure convergence with respect to the true underlying environment.

Definition 1. *(Convergence) Fix an environment $E = (G, H, (q_i)_{i\in N}, (\beta_i)_{i\in N}, (\sigma_i)_{i\in N})$. A belief-updating rule β_i converges to a stationary belief-updating rule $\bar{\beta}_i$ if*

$$\Pr^E\left(\lim_{t\to\infty}\|\beta_i(\tilde{h}_i^{t-1}, h_{it}) - \bar{\beta}_i(h_{it})\| = 0\right) = 1.$$

The induced function of play ψ converges to a stationary function of play $\bar{\psi}$ if

$$\Pr^E\left(\lim_{t\to\infty}\|\psi(\omega_t, \tilde{h}^{t-1}, t) - \bar{\psi}(\omega_t)\| = 0\right) = 1.$$

3 Empiricist learning rules

This section builds the class of empiricist learning rules and proves their desirable properties in stationary environments. Subsection 3.1 defines learning rules. Subsection 3.2 derives key properties of stationary environments. Subsection 3.3 defines empiricist learning rules and shows that in stationary environments, empiricist learning rules ensure asymptotic learning of relevant elements of the environment and asymptotic individual optimality.

3.1 Learning rules

Over time, individuals update their beliefs about the environment based on what they have observed. To capture the idea that individuals have minimal knowledge of the underlying environment, I assume that each individual is certain only about the set of her neighbors, the quality of her private signal, her own belief-updating rule and decision rule. For each i, denote by $\mathcal{E}_i \subseteq \mathcal{E}$ the set of environments that are consistent with the true tuple $(N_i, q_i, \beta_i, \sigma_i)$ of the underlying environment.

Definition 2. *(Learning rules) A learning rule $\Gamma_i : (\tilde{h}_i^{t-1}, h_{it}) \mapsto \tilde{p} \in \Delta(\mathcal{E}_i)$ of individual i is a map from i's histories to i's beliefs over environments she deems feasible. Belief-updating rule β_i is founded by learning rule Γ_i if for all history $(\tilde{h}_i^{t-1}, h_{it})$ and all θ,*

$$\beta_i(\tilde{h}_i^{t-1}, h_{it})(\theta) = \int_{E\in\mathcal{E}_i}\left(\frac{\Pr^E(h_{it}|\theta_t = \theta, \tilde{h}_i^{t-1})}{\Pr^E(h_{it}|\theta_t = \theta, \tilde{h}_i^{t-1}) + \Pr^E(h_{it}|\theta_t = -\theta, \tilde{h}_i^{t-1})}\right) d\Gamma_i(\tilde{h}_t^{t-1}, h_{it})(E).$$

Note that inside the integral is the likelihood that the current state is θ conditional on

i observing $(\tilde{h}_i^{t-1}, h_{it})$ in environment E. Intuitively, with a learning rule, an individual updates her belief about the underlying environment and then updates her belief about the current state using Bayes' rule.

3.2 Local models for stationary environments

In stationary environments, the key object of learning simplifies substantially. It reduces to a collection of conditional distributions of observations at each pair of an observation set and an action round. This collection is referred to as a local model. I then show that there is a one-to-one map between an individual's local model and the unconditional distribution of her observation, further reducing the object of learning to this unconditional distribution.

Definition 3. *(Stationary environments) An environment* $(G, H, (q_i)_{i\in N}, (\beta_i)_{i\in N}, (\sigma_i)_{i\in N})$ *is stationary for individual* i *if for all* $j \neq i$*, there exists a stationary belief-updating rule* $\bar{\beta}_j$ *such that* $\beta_j(\tilde{h}_j^{t-1}, h_{jt}) = \bar{\beta}_j(h_{jt})$ *for all history* $(\tilde{h}_j^{t-1}, h_{jt})$*, and a stationary decision rule* $\bar{\sigma}_j$ *such that* $\sigma_j(b_j, t) = \bar{\sigma}_j(b_j)$ *for all* t *and* $b_j \in \Delta(\{-1, 1\})$. *That is, the belief-updating rules of other individuals depend only on their current observations and their decision rules do not depend on time.*

Denote by $\bar{\mathcal{E}}_i \subseteq \mathcal{E}_i$ the set of all environments stationary for i. It holds for all $E \in \bar{\mathcal{E}}_i$, histories $(\tilde{h}_i^{t-1}, h_{it})$ and θ that

$$\Pr^E(h_{it}|\theta_t = \theta, \tilde{h}_i^{t-1}) = \Pr^E(h_{it}|\theta_t = \theta).$$

Thus, the likelihood that the current state is θ conditional on histories $(\tilde{h}_i^{t-1}, h_{it})$ is

$$\begin{aligned}&\frac{\Pr^E(h_{it}|\theta_t = \theta)}{\Pr^E(h_{it}|\theta_t = \theta) + \Pr^E(h_{it}|\theta_t = -\theta)}\\ =&\frac{\Pr^E(h_{it}|\theta_t = \theta, M_i(\tau_t) = M_i, r_i(\tau_t) = r)}{\Pr^E(h_{it}|\theta_t = \theta, M_i(\tau_t) = M_i, r_i(\tau_t) = r) + \Pr^E(h_{it}|\theta_t = -\theta, M_i(\tau_t) = M_i, r_i(\tau_t) = r)},\end{aligned}$$

where $M_i \subseteq N_i \cup \{i\}$ is i's observation set and r is i's action round when she observes h_{it}. The equality follows because the signal-timing vector is independent of the state of the world. The above equation means that for i to form her belief about the current state of the world, she only need to specify her belief about the stationary distribution of her observation conditional on the state of the world at each observation set and action round.

For each stationary environment $E \in \bar{\mathcal{E}}_i$, each nonempty set $M_i \subseteq N_i$, and each $r \in \{2, ..., R\}$, denote by $f_i^E(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r)$ the stationary probability that i observes (a_{M_i}, s_i) given observation set $M_i \cup \{i\}$ and action round r, conditional on the state of the world being θ. Similarly, write $f_i^E(a_{M_i}|\theta, M_i, r)$ for the probability that i observes a_{M_i} given observation set M_i and action round r, conditional on the state of the world being θ. Refer to the tuple $f_i^E = (f_i^E(.|\theta, M_i \cup \{i\}, r), f_i^E(.|\theta, M_i, r))_{\theta \in \{-1,1\}, M_i \subseteq N_i, r \in \{2,...,R\}}$ as i's local model in E. Appendix A.1 shows that f_i^E is well defined.[5]

With some abuse of notation, write

$$f_i^E(a_{M_i}|\theta, M_i \cup \{i\}, r) = f_i^E(a_{M_i}, 1|\theta, M_i \cup \{i\}, r) + f_i^E(a_{M_i}, -1|\theta, M_i \cup \{i\}, r)$$

for the conditional probability that i observes a_{M_i} from her neighbors when she receives some private signal. The following lemma proves two properties that hold in all stationary learning environments.

Lemma 1. *Take any $i \in N$ and any $E \in \bar{\mathcal{E}}_i$. The following properties hold:*

1) For every nonempty $M_i \subseteq N_i$, $r \in \{2, ..., R\}$ and $a_{M_i} \in \{-1, 1\}^{|M_i|}$,

$$\begin{pmatrix} f_i^E(a_{M_i}, -1|M_i \cup \{i\}, r) \\ f_i^E(a_{M_i}, 1|M_i \cup \{i\}, r) \end{pmatrix} = \begin{pmatrix} \Pr^E(s_i = -1, \theta = -1) & \Pr^E(s_i = -1, \theta = 1) \\ \Pr^E(s_i = 1, \theta = -1) & \Pr^E(s_i = 1, \theta = 1) \end{pmatrix} \times \begin{pmatrix} f_i^E(a_{M_i}|\theta = -1, M_i \cup \{i\}, r) \\ f_i^E(a_{M_i}|\theta = 1, M_i \cup \{i\}, r) \end{pmatrix}.$$

Furthermore, matrix $(\Pr^E(s_i, \theta))_{s_i, \theta \in \{-1,1\}}$ is full rank.

2) For every nonempty set $M_i \subseteq N_i$, $r \in \{2, .., R\}$ and $(a_{M_i}, s_i) \in \{-1, 1\}^{|M_i|+1}$,

$$\begin{cases} f_i^E(a_{M_i}|\theta, M_i, r) = f_i^E(a_{M_i}|\theta, M_i \cup \{i\}, r) \\ f_i^E(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r) = \Pr^E(s_i|\theta) f_i^E(a_{M_i}|\theta, M_i \cup \{i\}, r). \end{cases}$$

Proof sketch. The second equation in part 2 holds because private signals are independent of the diffusion process and across individuals, conditional on the state of the world. This

[5] Note also that the conditional distribution of an individual's observation when she receives only her private signal is known by her. For ease of notation, it is thus excluded from her local model.

implies the system of linear equations in part 1. Moreover, matrix $(\Pr^E(s_i,\theta))_{s_i,\theta\in\{-1,1\}}$ is full rank because i's private signal is informative. For the first equation in part 2, notice that the arrivals of private signals are independent across individuals. This implies that conditional on i acting later than her neighbors, the way information has diffused to her neighbors is not affected by whether i receives her private signal. See Appendix A.2 for the details. □

Since i knows the quality of her informative private signal, part 1 of Lemma 1 implies that the conditional distribution of the neighbor actions observed by i when she receives a private signal can be derived from the unconditional distribution of her observation when she receives a private signal. Then part 2 of the lemma completes a one-to-one mapping between i's local model and the tuple of unconditional distributions of i's observation when she receives a private signal.

Corollary 1. *For every E and $E' \in \bar{\mathcal{E}}_i$, if the respective local models f_i^E and $f_i^{E'}$ of individual i satisfy that $f_i^E(.|M_i\cup\{i\},r) = f_i^{E'}(.|M_i\cup\{i\},r)$ for all nonempty set $M_i \subseteq N_i$ and $r \in \{2,...,R\}$, then $f_i^E = f_i^{E'}$.*

3.3 Empiricist learning rules

Empiricist learning rules build on two ideas. First, if an individual believes that her learning environment is stationary, she does not need to learn the underlying environment but only the local model it induces. Second, this local model can be inferred from the unconditional distribution of her observation, which in turn could be learned from its empirical counterpart.

Learning about a stationary environment is connected to learning about a local model. Denote by $\mathcal{F}_i = \{f_i^E \text{ for some } E \in \bar{\mathcal{E}}_i\}$ the set of all local models induced by feasible environments that are stationary to i. Let Γ_i be i's learning rule. If $\Gamma_i(\tilde{h}_i^{t-1}, h_{it}) \in \Delta(\bar{\mathcal{E}}_i)$ for all histories $(\tilde{h}_i^{t-1}, h_{it})$, then Γ_i induces a *learning rule over local models*, $\gamma_i : h_i^t \mapsto p \in \Delta(\mathcal{F}_i)$, that satisfies for all $F_i \subseteq \mathcal{F}_i$,

$$\gamma_i(h_i^t)(F_i) = \int_{E\in\bar{\mathcal{E}}_i} \mathbf{1}\{f_i^E \in F_i\} d\Gamma_i(\tilde{h}_i^{t-1}, h_{it})(E).$$

Notice that to update her belief about the true local model, individual i only need to keep track of past observed actions of her neighbors and not her own. The reason is that in stationary environments, an individual's own action has no influence on the actions of other individuals in future periods, and thus does not matter for her learning.

Let β_i be the belief-updating rule founded by such Γ_i. It follows that for all histories $(\tilde{h}_i^{t-1}, h_{it})$ and state of the world θ,

$$\beta_i(\tilde{h}_i^{t-1}, h_{it})(\theta) = \int_{f_i \in \mathcal{F}_i} \left(\frac{f_i(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r)}{f_i(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r) + f_i(a_{M_i}, s_i|-\theta, M_i \cup \{i\}, r)} \right) d\gamma_i(h_i^t)(f_i)$$

if $h_{it} = (a_{M_i}, s_i, r)$, and

$$\beta_i(\tilde{h}_i^{t-1}, h_{it})(\theta) = \int_{f_i \in \mathcal{F}_i} \left(\frac{f_i(a_{M_i}|\theta, M_i, r)}{f_i(a_{M_i}|\theta, M_i, r) + f_i(a_{M_i}|-\theta, M_i, r)} \right) d\gamma_i(h_i^t)(f_i)$$

if $h_{it} = (a_{M_i}, r)$.

Next is to connect an individual's learning of her local model to the empirical distribution of her observation. Denote by $\hat{f}_i(.|M_i \cup \{i\}, r)(h_i^t)$ the empirical distribution of i's observation conditional on observation set $M_i \cup \{i\}$ and action round r given history h_i^t. Formally, for each nonempty set $M_i \subseteq N_i, r \in \{2, ..., R\}$ and $(a_{M_i}, s_i) \in \{-1, 1\}^{|M_i|+1}$, let

$$\hat{f}_i(a_{M_i}, s_i|M_i \cup \{i\}, r)(h_i^t) = \frac{\sum_{t'=1}^{t} \mathbf{1}\{h_{it'} = (a_{M_i}, s_i, r)\}}{\sum_{t'=1}^{t} \sum_{(a'_{M_i}, s'_i) \in \{-1,1\}^{|M_i|+1}} \mathbf{1}\{h_{it'} = (a'_{M_i}, s'_i, r)\}}$$

if the denominator is positive, otherwise set $\hat{f}_i(a_{M_i}, s_i|M_i \cup \{i\}, r)(h_i^t) = 1/2^{|M_i|+1}$. For every $\epsilon > 0$ and history h_i^t, define $F_i^\epsilon(h_i^t) \subseteq \mathcal{F}_i$ as the set of all local models f_i such that for all nonempty set $M_i \subseteq N_i$ and $r \in \{2, ..., R\}$,

$$\|f_i(.|M_i \cup \{i\}, r) - \hat{f}_i(.|M_i \cup \{i\}, r)(h_i^t)\| < \epsilon.$$

That is, $F_i^\epsilon(h_i^t)$ is the set of local models that induce unconditional distributions within ϵ-distance from their empirical counterparts.

Along every history, empiricist learning rules put probability one on stationary environments, so learning about the underlying environment reduces to learning about the local model. Furthermore, empiricist learning rules asymptotically put probability one on local models consistent with the empirical distribution of individuals' observations.

Definition 4. *(Empiricist learning rules) Learning rule* $\Gamma_i : (\tilde{h}_i^{t-1}, h_{it}) \mapsto \tilde{p} \in \Delta(\mathcal{E}_i)$ *is an empiricist learning rule if*

1) for all history $(\tilde{h}_i^{t-1}, h_{it})$*,* $\Gamma_i(\tilde{h}_i^{t-1}, h_{it}) \in \Delta(\bar{\mathcal{E}}_i)$*,*

2) and the map $\gamma_i : h_i^t \mapsto p \in \Delta(\mathcal{F}_i)$ *induced by* Γ_i *satisfies that for all* $\epsilon > 0$ *and* h_i^t*,*

$$\lim_{t\to\infty} \gamma_i(h_i^t)(F_i^\epsilon(h_i^t)) = 1.$$

A belief-updating rule is empiricist if it is founded by an empiricist learning rule.

Empiricist learning rules have a high-level connection with the literature on fictitious play, in particular, the concept of asymptotically empirical assessment rules by Fudenberg and Kreps (1993). In a fictitious play, each player in a repeated game believes that their opponents employ stationary mixed strategies. Asymptotically empirical assessment rules require that a player's assessment of others' mixed strategies asymptotically agrees with their empirical distribution. Analogously, in my model, individuals that use empiricist learning rules believe that their environments are stationary and that they should eventually take the empirical distribution of observed play as the true stationary play. The difference is that in my model, the distribution of play is not the final object of learning, but how it correlates with the unobserved state of the world. To learn the latter, empiricist learning rules exploit Lemma 1, which connects the distribution of observed neighbors' actions unconditional on the state of the world to that conditional on the state of the world.

Another intuition that carries analogously from fictitious play to my model is that fictitious play would be asymptotically optimal if the opponents indeed used stationary strategies. The reason is that by the Strong Law of Large Numbers, the empirical distribution of a stationary random variable converges to its true stationary distribution almost surely. An individual that asymptotically believed in the empirical distribution would, therefore, asymptotically learn the true stationary distribution of opponents' play and respond optimally. In my model, when the underlying learning environment is stationary, empiricist learning rules ensure asymptotic learning of the true local model and thus, asymptotic optimality. More specifically, part 2 of the following proposition says that conditional on sufficiently long histories, a pair of an empiricist learning rule and a smooth decision rule approximately optimizes the probability that an individual's action matches the true state of the world in the current period and in all future periods. Note that in an environment stationary to i, the distribution over histories h_i^∞ depends only on the objective environment and the profile of belief-updating rules and decision rules of individuals other than i. Since asymptotic learning holds at almost every history h_i^∞, so does asymptotic optimality.

Proposition 1. *Take any $i \in N, E \in \bar{\mathcal{E}}_i$ and let f_i^E be i's local model induced by E. Suppose that i adopts an empiricist learning rule Γ_i, which induces learning rule γ_i of local models and empiricist belief-updating rule β_i. The following hold:*

1) (Asymptotic Learning) For all $\epsilon > 0$,

$$\Pr^E\left(\lim_{t\to\infty} \gamma_i(h_i^t)\left(B^\epsilon\left(f_i^E\right)\right) = 1\right) = 1,$$

where $B^\epsilon\left(f_i^E\right) = \{f_i \in \mathcal{F}_i : \|f_i - f_i^E\| < \epsilon\}$ is the ϵ-ball around f_i^E.

2) (Asymptotic Optimality) For all belief-updating rule β_i', decision rule σ_i' and $\epsilon > 0$,

$$\Pr^E\Big(\exists T : \mathbb{E}^E(\sigma_i^\eta(\beta_i(\tilde{h}_i^{t-1}, h_{it}))(\theta_t)|h_i^t) \\ \geq \mathbb{E}^E(\sigma_i'(\beta_i'(\tilde{h}_i^{t-1}, h_{it}))(\theta_t)|h_i^t) - \epsilon, \forall t > T\Big) = 1,$$

where $\mathbb{E}^E$ is the expectation with respect to probability measure $\Pr^E$.

Proof sketch. In every stationary environment, it holds by the Strong Law of Large Numbers that the empirical distribution of an individual's observation converges to the true unconditional distribution almost surely. Empiricist learning rules, which asymptotically put probability one on local models consistent with the empirical distribution of observations, thus asymptotically put probability one on local models consistent with the true unconditional distribution. By Corollary 1, it follows that beliefs under empiricist learning rules asymptotically concentrate on the true local model. Part 2 follows from part 1 and the asymptotic optimality of smooth decision rules. See Appendix A.3 for the details. □

4 Convergence of learning

Section 3 showed that empiricist learning rules are individually asymptotically optimal in stationary environments. But how do they perform in nonstationary environments, where individuals actively learn? This section shows the first main theorem of my paper: convergence of beliefs and of play is achieved in every environment where all individuals adopt empiricist learning rules and smooth decision rules. Subsection 4.1 sketches the proof of this theorem. Subsection 4.2 discusses how the convergence result can be extended to more general settings.

4.1 Convergence on all networks

When all individuals adopt empiricist learning rules and smooth decision rules, individuals' play converges to a stationary play, where only the current observation affects an individual's belief about the current state of the world. Moreover, each individual asymptotically learns the true local model induced by the long-run play.

Theorem 1. *Fix any environment $E = (G, H, (q_i)_{i\in N}, (\beta_i)_{i\in N}, (\sigma_i^\eta)_{i\in N})$ where for each i, β_i is an empiricist belief-updating rule and σ_i^η is a smooth decision rule. Let γ_i be individual i's learning rule over local models and ψ be the function of play induced by E. The following hold.*

1) *For each $i \in N$, β_i converges to some stationary belief-updating rule $\bar{\beta}_i^*$.*

2) *ψ converges to some stationary function of play $\bar{\psi}^*$.*

3) *Let f_i denote the local model induced by $\bar{\psi}^*$. Then for each $i \in N$, γ_i asymptotically learns f_i. That is, for every ϵ-ball $B^\epsilon(f_i)$ around f_i,*

$$\Pr^E\left(\lim_{t\to\infty} \gamma_i(h_i^t)\left(B^\epsilon(f_i)\right) = 1\right) = 1.$$

Moreover, $\bar{\psi}^*$ is the symmetric and individually optimal stationary function of play that is uniquely defined by the objective environment $(G, H, (q_i)_{i\in N})$. For each i, $\bar{\beta}_i^*$ gives the correct likelihoods of the state of the world at each current observation, according to local model f_i.

Proof sketch. The proof of parts 1 and 2 proceeds inductively on action rounds, showing for each $r \in \{1, ..., R\}$ that for every individual i,

$$\Pr^E\left(\lim_{t\to\infty} \mathbf{1}\{r_i(\tau_t) = r\}\|\beta_i(\tilde{h}_i^{t-1}, h_{it}) - \bar{\beta}_i^*(h_{it})\| = 0\right) = 1$$

and

$$\Pr^E\left(\lim_{t\to\infty} \mathbf{1}\{r_i(\tau_t) = r\}\|\psi_i(\omega_t, \tilde{h}_{t-1}, t) - \bar{\psi}_i^*(\omega_t)\| = 0\right) = 1.$$

Consider the base case with $r = 1$ and take an arbitrary i. For every τ_t such that $r_i(\tau_t) = 1$, it must hold that $\tau_{it} = 1$ and $M_i(\tau_t) = \{i\}$. That is, i must receive her private signal in round 1 and choose an action using only her private signal. For all $\omega = (\tau, s, \zeta)$ such that $r_i(\tau) = 1$, construct

$$\bar{\beta}_i^*(s_i, 1)(\theta) = q_i^{\mathbf{1}\{s_i=\theta\}}(1 - q_i)^{\mathbf{1}\{s_i\neq\theta\}} \text{ for all } \theta \in \{-1, 1\}, \quad \text{and} \quad \bar{\psi}_i^*(\omega) = s_i.$$

Because individual i knows the quality of her private signal, $\beta_i(\tilde{h}_i^{t-1}, h_{it}) = \bar{\beta}_i^*(h_{it})$ for all observation $h_{it} = (s_{it}, 1)$. Moreover, since $q_i > 1/2$, her optimal action at such h_{it} is to choose $a_{it} = s_{it}$. It then follows from the definition of smooth decision rules that ψ_i asymptotically puts probability one on the action that agrees with her private signal when she receives only her private signal. This completes the proof of the base case.

The inductive step from r to $r + 1$ relies on showing that as an individual's environment becomes approximately stationary, her belief converges and thus so does her play. Empiricist learning rules ensure convergence of learning of approximately stationary environments, and smooth decision rules ensure that convergence of beliefs leads to convergence of play. This argument is formalized by Lemma A2 in Appendix A.4. Finally, the proof of part 3 extends the proof of part 1 of Proposition 1 to approximately stationary environments. □

Theorem 1 demonstrates a form of learning externalities: if an individual learns the informativeness of her neighbors' actions well, then the informativeness of her action to other neighbors will be learned well by these other neighbors. As a general pattern of learning, convergence spreads from early receivers of news to late receivers of news. This pattern is inherent in the inductive construction of the profile of stationary belief-updating rules $(\bar{\beta}_i^*)_{i\in N}$ and the stationary function of play $\bar{\psi}^*$ (see Appendix A.4).

While the way information diffuses across rounds suggests a natural order for an inductive argument, the inductive step relies crucially on the fact that empiricist learning rules do not allow (mis)learning in later rounds to contaminate learning in earlier rounds. The latter holds because under empiricist learning rules, learning in a round depends only on the empirical frequency of neighbors' actions observed in that round, which in turn depends only on others' learning in earlier rounds. Note that neighbors' play across periods is in fact linked through the fundamentals of the environment, such as the network structure and the informativeness of private signals. By forgoing information from such linkages, empiricist learning rules may slow down learning but they avoid potential failures due to misspecification.

That convergence of play under empiricist learning rules holds in all networks has important implications. First, it lends support to the assumption of empiricist learning rules that individuals believe in stationary learning environments. Second and most importantly, it allows the long-run performance of empiricist learning rules to be evaluated at the stationary play to which play under empiricist learning rules and smooth decision rules converges. This stationary play is the focus of Sections 5 and 6.

Another implication of the general convergence result is that my model can provide a learning foundation for the common knowledge assumptions in a direct analysis of the one-period game. Consider the one-period game of my model where the objective environment is common knowledge and each individual aims to maximize the probability that her action matches the true state of the world. In this game, the unique Perfect Bayesian equilibrium that breaks ties equally between the two states of the world induces the same stationary play as the long-run play in my model when individuals use empiricist learning rules and smooth decision rules.

4.2 Convergence in more general settings

In fact, the convergence result can be generalized to richer settings.

First, the objective environment can allow any finite number of states of the world and general diffusion processes that are not necessarily tied to a permanent underlying network. The crucial assumption to maintain is the stationarity of the objective environment. This assumption ensures that at each inductive step, the local environment becomes approximately stationary as others' play converges.

Second, individuals can have multiple signals with arbitrary correlation, as long as each individual has a private signal that satisfies *independence* and *informativeness*. Here independence means that the arrival and realization of this private signal is independent of other components of the individual's observation set. This allows for a mapping between the unconditional distribution of an individual's observation when she receives this private signal and the corresponding conditional distributions, analogous to property 1 of Lemma 1. The independence of this private signal also allows a direct connection between the conditional distribution of an individual's observation when she receives and when she does not receive this signal, analogous to property 2 of Lemma 1. If the private signal is informative, these two mappings result in a one-to-one mapping between the unconditional distribution of one's observation and her local model. Here, informativeness means that the matrix $(\Pr^E(s_i, \theta))_{s_i,\theta}$ has full column rank.

Finally, as in the main setting, individuals that use empiricist learning rules believe in stationary environments and asymptotically believe in the empirical distribution of their observation. Furthermore, they use asymptotically optimal decision rules that are smoothed to ensure that convergence of beliefs leads to convergence of play. As in the main setting, one way of smoothing is to derive these rules from the choice probabilities of a random utility model that takes subjective beliefs as mean utilities and has decaying logistic errors.

5 Strong efficiency on clique trees

This section builds on the convergence result of the previous section to prove a positive efficiency result: on clique trees, the long-run play induced by empiricist learning rules and smooth decision rules achieves strong efficiency within the class of symmetric feasible stationary plays. That is, no other symmetric feasible stationary play can make some individual strictly better off, even at the expense of some other individuals.

Definition 5. *(Strong efficiency) Fix an objective environment, which induces probability measure* $\Pr$ *on* (θ, ω)*. A stationary function of play* $\bar{\psi}$ *is said to dominate another stationary function of play* $\bar{\psi}'$ *if for every* i*,*

$$\Pr(\bar{\psi}_i(\omega) = \theta) \geq \Pr(\bar{\psi}'_i(\omega) = \theta).$$

*A stationary function of play is **strongly efficient** within a class of stationary functions of play if it belongs to that class and it dominates all stationary functions of play in that class.*

Subsection 5.1 defines the class of symmetric feasible stationary functions of play. Subsection 5.2 derives properties of clique trees that are key to the efficiency result, which is shown in Subsection 5.3.

5.1 Symmetric feasible stationary functions of play

Recall from Theorem 1 that if all individuals adopt empiricist learning rules and smooth decision rules, the induced play converges to a stationary function of play $\bar{\psi}^*$ that is well-defined by the underlying objective environment. More specifically, $\bar{\psi}^*$ is associated with the profile $(\bar{\beta}_i^*)_{i\in N}$ of stationary belief-updating rules that correctly learn the local environments and the profile $(\bar{\sigma}_i^*)_{i\in N}$ of optimal stationary decision rules that treat the two states symmetrically.

For efficiency comparison, it is only meaningful to compare $\bar{\psi}^*$ to stationary functions of play that respect local observability. That is, an individual's play can only condition on her observation.

Definition 6. *(Feasibility) A stationary function of play $\bar{\psi}$ is feasible if and only if it can be induced by a profile of stationary belief-updating rules $(\bar{\beta}_i)_{i\in N}$ and a profile of stationary decision rules $(\bar{\sigma}_i)_{i\in N}$.*

Moreover, to study environments without biases towards either state of the world, one can focus on the class of symmetric stationary functions of play. Since the objective environment is symmetric with respect to the state of the world, play should also be symmetric with respect to the state of the world if individuals care equally about false positives and false negatives.

A stationary belief-updating rule $\bar{\beta}_i$ is *symmetric* if for every observation (a_{M_i}, s_i, r),

$$\bar{\beta}_i(a_{M_i}, s_i, r)(\theta) = \bar{\beta}_i(-a_{M_i}, -s_i, r)(-\theta),$$

and for every observation (a_{M_i}, r) of i,

$$\bar{\beta}_i(a_{M_i}, r)(\theta) = \bar{\beta}_i(-a_{M_i}, r)(-\theta).$$

A stationary decision rule $\bar{\sigma}_i$ is symmetric if for all $b \in \Delta(\{-1,1\})$, $\bar{\sigma}_i(\mathbf{1}-b) = \mathbf{1} - \bar{\sigma}_i(b)$. That is, a symmetric belief-updating rule flips one's belief between the two states of the world if the observed actions and private signal flip signs; a symmetric decision rule flips the probability of each action when one's belief flips between the two states of the world.

Definition 7. *(Symmetry) A feasible stationary function of play $\bar{\psi}$ is symmetric if*

$$\bar{\psi}(\tau, s, \zeta) = -\bar{\psi}(\tau, -s, \mathbf{1} - \zeta)$$

for all τ, s and ζ. Equivalently, $\bar{\psi}$ is symmetric if and only if it can be induced by some profile $(\bar{\beta}_i)_{i \in N}$ of symmetric belief-updating rules and some profile $(\bar{\sigma}_i)_{i \in N}$ of symmetric stationary decision rules.[6]

5.2 Clique trees

The main efficiency theorem of this paper establishes that individuals on clique trees have their interests perfectly aligned. This subsection describes clique trees and the two properties of clique trees that are sufficient for strong efficiency.

Definition 8. *(Clique trees) A network G is a clique tree if any two individuals i and j that belong to the same cycle are linked.*

Essentially, clique trees are extensions of trees, where a node can be replaced by a clique, that is, a set of fully linked individuals. Qualitatively, clique trees are stylized examples of networks of different groups with dense within-links and sparse between-links.

Recall that information travels from one individual to another individual only through the shortest paths connecting them. Thus, every feasible stationary function of play $\bar{\psi}$ must satisfy that for every i, $\bar{\psi}_i$ depends on the system state ω only through:

- the signal-timing vector τ,
- the private signal s_j of each $j \in \arg\min_{k \in N}\{\tau_k + l(k,i)\}$,
- the instrumental draws ζ_j of each $j \in \arg\min_{k \in N}\{r_k(\tau) + l(k,i)\}$ and ζ_i.

[6] The equivalence relation is shown in Appendix A.5.

This means that the role of the underlying network in information diffusion is exactly summarized by the structure it imposes on the shortest paths connecting individuals. Let $IF_{j\setminus i} = \{k \in N : l(k,j) + l(j,i) = l(k,i)\}$ denote the set of individuals whose shortest connecting paths to i go through j. This notion is useful for establishing two technical properties of clique trees. Note that on geodetic networks, any two individuals are connected through a unique shortest path, and that clique trees are geodetic.

Lemma 2. *Take any connected network G.*

1) If G is geodetic then for all $i \in N$, $\{IF_{j\setminus i}\}_{j \in N_i}$ is a partition of $N\setminus\{i\}$.

2) If G is a clique tree then for all $i \in N$, $j \in N_i$ and $k \in N_j\setminus(N_i \cup \{i\})$, $IF_{k\setminus j} \subseteq IF_{j\setminus i}$.

Proof. See Appendix A.6. □

The above lemma establishes two technical properties of clique trees. First, an individual's neighbors act as separate channels through which signals and actions of other individuals can influence her. Second, sources that are not jointly shared between two neighbors i and j and influence j must also influence i exactly through j. In Lemma 3 and Lemma 4, these two technical properties of clique trees are translated into two properties crucial for establishing strong efficiency. First, the separation of sources that can influence an individual is exploited to show that for every individual on a geodetic network, the observed actions of her neighbors are always independent conditional on the state of the world. As a result, the overall informativeness of these neighbors' actions is increasing in the informativeness of each of their actions, captured by a sufficient statistic under symmetry of play. Then, the relation between the set of those individuals that can influence an individual and the set of those influencing her neighbors is shown to align neighbors' interests. That is, an individual maximizes the informativeness of her action to her neighbors exactly by optimally aggregating information at each of her conditioning sets.

Lemma 3 establishes the conditional independence of observed neighbors' actions. The proof involves two steps, each of which exploits part 1 of Lemma 2. The first step further uses the independence of the realizations of private signals to show that the actions of i's neighbors are independent conditional on the state of the world and the signal-timing vector. The second step further uses the independence of signal arrivals to show that the realized path of information that reaches an individual through each neighbor of hers is independent across these neighbors.

Lemma 3. *(Conditional independence of observed neighbors' actions) Fix an objective environment* $(G, H, (q_i)_{i\in N})$ *where* G *is geodetic and a feasible stationary function of play* $\bar{\psi}$*. For every* i*, nonempty set* $M_i \subseteq N_i \cup \{i\}$ *and* $r \in \{2, ..., R\}$*,*

$$\Pr(\bar{\psi}_{M_i\setminus\{i\}}(\omega)|\theta, M_i(\tau) = M_i, r_i(\tau) = r) = \prod_{j\in M_i\setminus\{i\}} \Pr(\bar{\psi}_j(\omega)|\theta, M_i(\tau) = M_i, r_i(\tau) = r).$$

Proof. See Appendix A.7. □

An implication of Lemma 3 together with symmetry is that for each conditioning set of an individual, what matters for her prediction is a set of sufficient statistics, each capturing the quality of each source that she hears from. Specifically, fix an objective environment where the network is a clique tree and consider a symmetric feasible stationary function of play $\bar{\psi}$. Take an individual i with observation set $M_i \subseteq N_i$ and action round $r \in \{2, ..., R\}$. For each $j \in M_i$, define the quality of j's signal given i's conditioning set being (M_i, r) by

$$\begin{aligned}\tilde{q}_j^{M_i,r} &= \Pr(\bar{\psi}_j(\omega) = -1|\theta = -1, M_i(\tau) = M_i, r_i(\tau) = r)\\ &= \Pr(\bar{\psi}_j(\omega) = 1|\theta = 1, M_i(\tau) = M_i, r_i(\tau) = r).\end{aligned}$$

The second equality follows from the symmetry of $\bar{\psi}$ and the assumption that $\Pr(s|\theta) = \Pr(-s| - \theta)$ for all $s \in \{-1, 1\}^N$. Then the likelihood ratio of the state of the world θ conditional on i observing a_{M_i} in round r is

$$\frac{\mathcal{L}(\theta = 1|a_{M_i}, r)}{\mathcal{L}(\theta = -1|a_{M_i}, r)} = \prod_{j\in M_i} \left(\frac{\tilde{q}_j^{M_i,r}}{1 - \tilde{q}_j^{M_i,r}}\right)^{\mathbf{1}\{a_j=1\}-\mathbf{1}\{a_j=-1\}}.$$

When individual i also observes her private signal, so i's observation set is $M_i \cup \{i\}$, the likelihood ratio becomes

$$\frac{\mathcal{L}(\theta = 1|a_{M_i}, s_i, , r)}{\mathcal{L}(\theta = -1|a_{M_i}, s_i, r)} = \left(\frac{q_i}{1 - q_i}\right)^{\mathbf{1}\{s_i=1\}-\mathbf{1}\{s_i=-1\}} \prod_{j\in M_i} \left(\frac{\tilde{q}_j^{M_i,r}}{1 - \tilde{q}_j^{M_i,r}}\right)^{\mathbf{1}\{a_j=1\}-\mathbf{1}\{a_j=-1\}}.$$

The optimal predictor for individual i is to choose $a_i = 1$ when the likelihood ratio is larger than 1 and to choose $a_i = -1$ otherwise. Given that the different sources that i hears from are conditionally independent, the probability that her optimal predictor matches the state of the world is weakly increasing in the probability that each source matches the state of the world. The reason is that were she to have sources of higher quality, she could always distort their quality downward by adding symmetric independent noise and use the optimal predictor of the lower quality sources. Note that under $\bar{\psi}^*$ induced in the long run by all individuals using empiricist learning rules and smooth decision rules, each individual uses the optimal predictor at each of her conditioning sets. The next lemma shows that on clique trees, when a neighbor j of i optimizes the probability that her action matches the state of the world, she simultaneously maximizes the informativeness of her action to i. This is an important property for selfish learning to be socially efficient.

Lemma 4. *(Local alignment of interests) Fix an objective environment $(G, H, (q_i)_{i\in N})$ where G is a clique tree. Let $\bar{\psi}$ be a symmetric feasible stationary function of play of this objective environment. Take an individual i, a nonempty set $M_i \subseteq N_i$ and an action round $r \in \{2, ..., R\}$. For every $j \in M_i$,*

$$\tilde{q}_j^{M_i,r} = \sum_{M_j \subseteq (N_j \setminus (N_i \cup \{i\})) \cup \{j\}} w_{M_j} \times \Pr(\bar{\psi}_j(\omega) = \theta | \theta, M_j(\tau) = M_j, r_j(\tau) = r - 1),$$

where each $w_{M_j} \geq 0$ depends only on G and H, and $\sum_{M_j \subseteq (N_j \setminus (N_i \cup \{i\})) \cup \{j\}} w_{M_j} = 1$.

Proof. See Appendix A.8. □

The proof of this lemma relies on part 2 of Lemma 2, which ensures that any neighbor j of i, given her observation set and action round, can distinguish between diffusions that affect her but not i and diffusions that reach i through her. Thus, the informativeness of j's action to i can be written as a weighted average of the probabilities that j's action matches the state of the world at different conditioning sets of j. This establishes local alignment of interests.

5.3 Strong efficiency

Theorem 2 shows that under conditional independence of observed neighbors' actions, local alignment of interests implies global alignment of interests. Formally, let $\bar{\psi}^*$ be the long-run play when all individuals in a network adopt empiricist learning rules and smooth decision rules. Strong efficiency on clique trees is established by inductively applying the following implication of Lemma 3 and Lemma 4: for an individual to aggregate information better under an alternative symmetric feasible stationary function of play $\bar{\psi}$ than she does under $\bar{\psi}^*$, some of her neighbors must aggregate information better under $\bar{\psi}$ than they do under $\bar{\psi}^*$.

Theorem 2. *Fix an objective environment $(G, H, (q_i)_{i\in N})$. Let $\bar{\psi}^*$ be the stationary function of play induced in the long run by all individuals adopting empiricist learning rules and smooth decision rules. If G is a clique tree then $\bar{\psi}^*$ is strongly efficient within the class of symmetric feasible stationary functions of play.*

Proof. Suppose to the contrary that there exists some symmetric feasible stationary function of play $\bar{\psi}$ such that for some i,

$$\Pr(\bar{\psi}_i(\omega) = \theta) > \Pr(\bar{\psi}_i^*(\omega) = \theta).$$

Then there must exist some $M_i \subseteq N_i \cup \{i\}$ and $r \in \{1, ..., R\}$ such that

$$\Pr(\bar{\psi}_i(\omega) = \theta|\theta, M_i(\tau) = M_i, r_i(\tau) = r) > \Pr(\bar{\psi}_i^*(\omega) = \theta|\theta, M_i(\tau) = M_i, r_i(\tau) = r).$$

Take r to be the smallest round such that the above inequality holds for some i and some $M_i \subseteq N_i \cup \{i\}$. If $M_i = \{i\}$ then

$$q_i = \Pr(\bar{\psi}_i^*(\omega) = \theta|\theta, M_i(\tau) = \{i\}, r_i(\tau) = r) \geq \Pr(\bar{\psi}_i(\omega) = \theta|\theta, M_i(\tau) = \{i\}, r_i(\tau) = r),$$

that is, the above strict inequality cannot hold. Thus M_i includes some of i's neighbors, which also implies that $r \geq 2$.

Since individuals use the optimal predictor under $\bar{\psi}^*$, by Lemma 3 and symmetry of play, $\bar{\psi}_i$ is strictly more likely to match the state of the world at conditioning set (M_i, r) than does $\bar{\psi}_i^*$ only if the action of some neighbor of i is more informative to i under $\bar{\psi}$ than under $\bar{\psi}^*$. That is, there exists some $j \in M_i \backslash \{i\}$ such that

$$\Pr(\bar{\psi}_j(\omega) = \theta|\theta, M_i(\tau) = M_i, r_i(\tau) = r) > \Pr(\bar{\psi}_j^*(\omega) = \theta|\theta, M_i(\tau) = M_i, r_i(\tau) = r).$$

By Lemma 4, this means that there must exist some $M_j \subseteq (N_j \backslash (N_i \cup \{i\})) \cup \{j\}$ such that

$$\begin{aligned} &\Pr(\bar{\psi}_j(\omega) = \theta | \theta, M_j(\tau) = M_j, r_j(\tau) = r-1) \\ &> \Pr(\bar{\psi}_j^*(\omega) = \theta | \theta, M_j(\tau) = M_j, r_j(\tau) = r-1). \end{aligned}$$

This contradicts that r is the smallest round in which someone can be made strictly better off at some conditioning set. □

Theorem 2 identifies a condition on the underlying network that ensures strong efficiency regardless of the timing of signal arrivals and the quality of private signals. On a tree, the informativeness of a neighbor's action is summarized by a sufficient statistic and the overall informativeness of the actions of one's neighbors is increasing in each of these sufficient statistics. Inductively, better information aggregation starting from the leaves of a tree helps information aggregation towards the root of that tree. Since cliques, or complete subgraphs, do not add confoundedness of sources, the above intuition generalizes to clique trees.

6 Challenges to efficiency

This section explains the challenges to generalizing Theorem 2. It includes a formal converse, Theorem 3, showing that on any network that is not a clique tree, strong efficiency fails for some diffusion process and signal structure. Recall from Lemma 3 and Lemma 4 that clique trees possess two key properties. First, observed neighbors' actions are always conditionally independent. Second, an individual optimizes the informativeness of her action to her neighbors exactly by optimizing her information aggregation at each of her conditioning sets. Theorem 3 generalizes the intuition of two examples: one fails the first property of clique trees and not the second; the other fails the second property and not the first. These examples are presented in Subsections 6.1 and 6.2. The converse is presented in Subsection 6.3. Subsection 6.4 explores whether it is the strong notion of efficiency that demands a restrictive class of networks. The short answer is no. The general forces that lead to failure of strong efficiency as illustrated in Subsections 6.1 and 6.2, when coupled with opportunities for favor trading across rounds, can result in failure of Pareto efficiency. Finally, Subsection 6.5 illustrates the implications of asymmetric environments, where asymmetry comes from either the underlying environment or individuals' play.

In each example of this section, $\bar{\psi}^*$ denotes the stationary function of play induced in the long run by all individuals using empiricist learning rules and smooth decision rules.

6.1 Failure of conditional independence of neighbors' actions

The following example shows that when neighbors' actions correlate, it is possible to distort downward the individual informativeness of some neighbor's action to break the correlation in a way that improves the overall informativeness of all neighbors' actions.

Example 1. *An objective environment has* $G = \{\{1,2\},\{1,3\},\{2,4\},\{3,4\}\}$*,* $R = 3$*,* $\Pr(\tau_1 = 1) \approx 1, \Pr(\tau_2 = 2) = \Pr(\tau_2 = \infty) = \Pr(\tau_3 = 2) = \Pr(\tau_3 = \infty) \approx 0.5, \Pr(\tau_4 = 3) \approx 1$*, and* $(q_1, q_2, q_3, q_4) = (0.5 + \epsilon, 0.75, 0.9, 0.75)$ *for* $\epsilon > 0$ *small.*

There exists a symmetric feasible stationary function of play where individual 2 plays a suboptimal strategy that makes individual 4 strictly better off compared to $\bar{\psi}^*$*.*

Proof. The first panel of Figure 2 plots the network, which is not a clique tree. The second panel illustrates the flow of information when $\tau = (1, \infty, \infty, 3)$, which occurs with probability close to 0.25. The third panel illustrates the flow of information when $\tau \in \{(1,2,2,3),(1,2,\infty,3),(1,\infty,2,3)\}$, which occurs with probability close to 0.75. In both cases, information flows from individual 1 to individuals 2 and 3 and then to individual 4, who chooses an action based on what she observes from 2 and 3 and her private signal. The difference is that in the second panel neither 2 nor 3 receives their private signal while in the third panel at least one of them does.

Consider $\bar{\psi}^*$ first. Since $q_2 > q_1$, individual 2 passes onto individual 4 her private signal if she receives it, $a_2 = s_2$, and otherwise passes on the action of individual 1, $a_2 = a_1 = s_1$. Similarly individual 3 chooses $a_3 = s_3$ if he receives his private signal and otherwise chooses $a_3 = a_1 = s_1$. Thus, a_2 and a_3 are perfectly correlated when neither individual 2 nor 3 receives their private signal, and a_2 and a_3 are conditionally independent otherwise.

Here optimally, individual 4 chooses $a_4 = s_4$ when $a_2 \neq a_3$ or $a_2 = a_3 = s_4$. The key question is whether individual 4 would follow the consensus actions of individuals 2 and 3

Figure 2: The network and key diffusions in Example 1

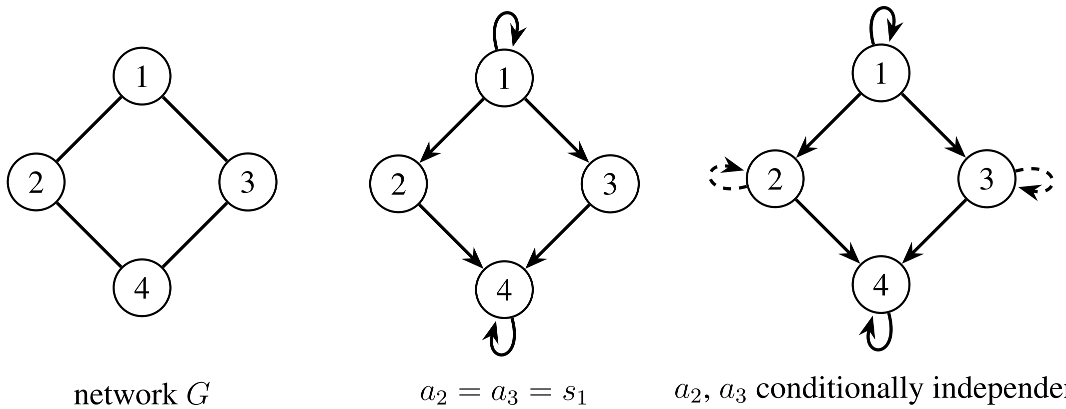

when $a_2 = a_3 \neq s_4$. The informativeness of $a_2 = a_3$ is summarized by the following ratio

$$\begin{aligned}&\frac{\Pr(a_2 = a_3 = \theta|\theta, M_4(\tau) = \{2,3,4\}, r_4(\tau) = 3)}{\Pr(a_2 = a_3 = -\theta|\theta, M_4(\tau) = \{2,3,4\}, r_4(\tau) = 3)}\\ \approx&\frac{0.25(0.5+\epsilon) + 0.25(0.5+\epsilon)(0.75) + 0.25(0.5+\epsilon)(0.9) + 0.25(0.75)(0.9)}{0.25(0.5-\epsilon) + 0.25(0.5-\epsilon)(0.25) + 0.25(0.5-\epsilon)(0.1) + 0.25(0.25)(0.1)}\\ \approx&2.86 < 3 = \frac{\Pr(s_4 = \theta|\theta)}{\Pr(s_4 = -\theta|\theta)}.\end{aligned}$$

This means that under $\bar{\psi}^*$, individual 4 always chooses $a_4 = s_4$.

Now consider an alternative stationary function of play $\bar{\psi}$ where individual 2 randomizes with equal probabilities between 1 and -1 when she does not receive her private signal. This change breaks the correlation between a_2 and a_3. Then from $\Pr(a_2 = \theta|\theta, M_4(\tau) = \{2,3,4\}, r_4(\tau) = 3) \approx 0.625$ and $\Pr(a_3 = \theta|\theta, M_4(\tau) = \{2,3,4\}, r_4(\tau) = 3) \approx 0.7 + 0.5\epsilon$,

$$\begin{aligned}&\frac{\Pr(a_2 = a_3 = \theta|\theta, M_4(\tau) = \{2,3,4\}, r_4(\tau) = 3)}{\Pr(a_2 = a_3 = -\theta|\theta, M_4(\tau) = \{2,3,4\}, r_4(\tau) = 3)}\\ \approx&\frac{(0.625)(0.7+0.5\epsilon)}{(0.375)(0.3-0.5\epsilon)} \approx 3.89 > \frac{\Pr(s_4 = \theta|\theta)}{\Pr(s_4 = -\theta|\theta)}.\end{aligned}$$

This means that when $a_2 = a_3 \neq s_4$, individual 4 is now strictly better-off choosing $a_4 = a_2 = a_3$ than choosing $a_4 = s_4$. □

The key intuition of this example is that the correlation between the actions of individuals 2 and 3 when they jointly receive low-quality information from individual 1 downgrades the average informativeness of their consensus actions to individual 4. Note that correlation between different information sources does not hurt learning. Individual 4 correctly learns the correlation structure of the actions of individuals 2 and 3, but it is because she cannot distinguish whether the second panel or the third panel of Figure 2 takes place that she disregards her neighbors' actions altogether. Furthermore, since the private signal of individual 1 is of low quality, the modification made in the alternative function of play hurts individual 2 little but lends noticeable help to individual 4.

Finally, notice that by setting $R = 3$, I ensure that this example does not fail the second property of clique trees. In this example, individual informativeness of each neighbor's action is indeed optimized by the neighbors' selfish learning. Intuitively, failure of the second property requires that an individual hear from some neighbor who himself faces confoundedness of sources, that is, he heard from another neighbor who had heard from some other neighbor. This chain is not possible when $R = 3$. The purpose of this technical trick is to separate the importance of each property, showing that the failure of one property does not imply the failure of the other.

6.2 Failure of local alignment of interests

By local observabiliity, the local model of each individual is averaged over different diffusion realizations. This means that an individual may play optimally to an average distribution that pools some diffusions irrelevant to her neighbors. In other words, the informativeness of an individual's action to her neighbors is not necessarily maximized by her playing optimally at each of her conditioning sets, as is the case in the following example.

Example 2. *An objective environment has* $G = \{\{1, 2\}, \{1, 3\}, \{2, 4\}, \{3, 5\}, \{4, 5\}\}$*,* $R = 4$*,* $\Pr(\tau_1 = 1) = \Pr(\tau_1 = \infty) \approx 0.5, \Pr(\tau_2 = 2) = \Pr(\tau_4 = 3) \approx 1, \Pr(\tau_3 = \infty) = \Pr(\tau_5 = \infty) \approx 1$*, and* $(q_1, q_2, q_3, q_4, q_5) = (0.9, 0.5 + \epsilon, 0.5 + \epsilon, 0.7, 0.7)$.

There exists a symmetric feasible stationary function of play where individual 4 plays a suboptimal strategy that makes individual 5 strictly better off compared to $\bar{\psi}^*$.

Proof. The first panel of Figure 3 plots the network, which is geodetic but not a clique tree. The second and the third panels of this figure illustrate the two main realized diffusions of

Figure 3: The network and key diffusions in Example 2

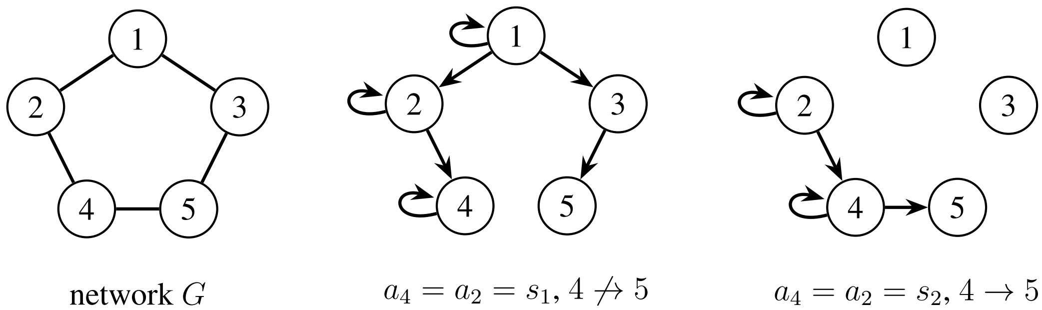

this objective environment, each occurring with probability approximately 0.5. On the second panel where individual 1 receives his private signal in round 1, individual 5 takes action in round 3 after observing the action of individual 3. On the third panel where individual 1 does not receive his private signal, individual 5 takes action in round 4 after observing the action of individual 4.

Under $\bar{\psi}^*$, individual 2 passes onto individual 4 the private signal of individual 1 when she hears from him and passes on her own private signal otherwise. Thus, individual 4 learns that

$$\Pr(a_2 = \theta | \theta, M_4(\tau) = \{2, 4\}, r_4(\tau) = 3) \approx 0.5(0.9) + 0.5(0.5 + \epsilon) = 0.7 + 0.5\epsilon > q_4.$$

This means that under $\bar{\psi}^*$, individual 4 chooses $a_4 = a_2$ when he hears from individual 2 and receives a private signal himself in round 3. Note that while individual 4 plays optimally against the distribution of individual 2's action averaged over both cases illustrated in Figure 3, only in the case in the third panel of this figure does individual 4's action matter to individual 5. Thus,

$$\Pr(a_4 = \theta | \theta, M_5 = \{4\}, r_5(\tau) = 4) = \Pr(s_2 = \theta | \theta) = 0.5 + \epsilon.$$

Consider an alternative symmetric feasible stationary function of play $\bar{\psi}$ where individual 4 chooses $a_4 = s_4$ upon hearing from individual 2 in round 3. Then

$$\Pr(a_4 = \theta | \theta, M_5 = \{4\}, r_5(\tau) = 4) = \Pr(s_4 = \theta | \theta) = 0.7.$$

Thus $\bar{\psi}$ makes individual 5 strictly better off, at a (small) expense of individual 4. □

6.3 A converse of Theorem 2

The following theorem generalizes Example 1 and Example 2 to show that on any network that is not a clique tree, strong efficiency fails for some diffusion process and some set of private signals.

Theorem 3. *Suppose that G is not a clique tree. Then there exist a distribution H of the signal-timing vector and a vector of signal qualities $(q_i)_{i \in N}$ such that the stationary function of play $\bar{\psi}^*$ induced in the long-run by all individuals adopting empiricist learning rules and smooth decision rules is not strongly efficient, even within the class of symmetric stationary functions of play.*

Proof sketch. Note that a cycle is incomplete if some individuals in the cycle are not linked. A cycle of length at least four is chordless if any two non-consecutive individuals in the cycle are not linked. There are three cases based on the smallest incomplete cycle G' of G.

Case 1: G' is a chordless cycle of size $2m$ for $m \geq 2$.

Case 2: G' is a chordles cycle of size $2m+1$ for $m \geq 2$.

Case 3: G' is a chordal graph, that is, it does not have any chordless cycle.

The construction of H and $(q_i)_{i \in N}$ for Case 1 is a generalization of Example 1, where the correlation in the actions of two neighbors influenced by a common source makes the individual disregards their consensus actions and breaking such correlation could benefit the individual. The construction for Case 2 is a generalization of Example 2 where an individual follows a neighbor's action that is less informative than her own private signal exactly when her action will be observed by another neighbor. If she used her private signal, the latter neighbor would strictly benefit. In Case 3, it can be shown that G' must be a cycle of size four with exactly one missing link. Then the construction for Case 1 applies. See Appendix A.9 for the details. □

6.4 Pareto improvement from favor trading

This section explores whether it is the strong notion of efficiency that demands a restrictive class of networks. First, I show that a weak form of Pareto efficiency holds on all networks:

compared to the long-run play induced by empiricist learning rules and smooth decision rules, there is no symmetric play that strictly improves the probability of correct prediction by an individual without hurting the probability of correct prediction by some other individual in some round. This weak efficiency result does not extend to the standard notion of Pareto efficiency where an individual's welfare is measured by the probability of correct prediction pooling over all action rounds. The reason is that compensations can be made across rounds. An individual may forgo some benefits for her neighbor's sake when she receives information early if her neighbor forgoes some of his benefits for her sake when she receives information late. Examples 3 and 4 build on Examples 1 and 2 and the idea of compensations across rounds to illustrate how Pareto efficiency may not extend much beyond clique trees.

Proposition 2. *Fix an objective environment. Let* $\bar{\psi}^*$ *be the stationary function of play induced as the long-run play when all individuals on the network adopt empiricist learning rules and smooth decision rules. There does not exist any symmetric feasible stationary function of play* $\bar{\psi}$ *such that for each* $i \in N$ *and each* $r \in \{1, ..., R\}$,

$$\Pr(\bar{\psi}_i(\omega) = \theta | r_i(\tau) = r) \geq \Pr(\bar{\psi}^*_i(\omega) = \theta | r_i(\tau) = r),$$

with the inequality holding strictly for some i *and some* r.

Proof. See Appendix A.10. □

This proposition shows that in environments without biases, failure of Pareto efficiency, as defined below, must be due to some form of favor-trading across rounds.

Definition 9. *(Pareto efficiency) Fix an objective environment. A feasible stationary function of play* $\bar{\psi}$ *is* ***Pareto efficient*** *within some class of stationary functions of play if it belongs to that class and there does not exist any stationary function of play* $\bar{\psi}'$ *in that class such that for all* i,

$$\Pr(\bar{\psi}'_i(\omega) = \theta) \geq \Pr(\bar{\psi}_i(\omega) = \theta),$$

with a strict inequality holding for some i.

Following I present two examples showing that forces that present challenges to strong efficiency, correlation of neighbors' actions and pooling over diffusions irrelevant to one's

Figure 4: The network and key diffusions in Example 3

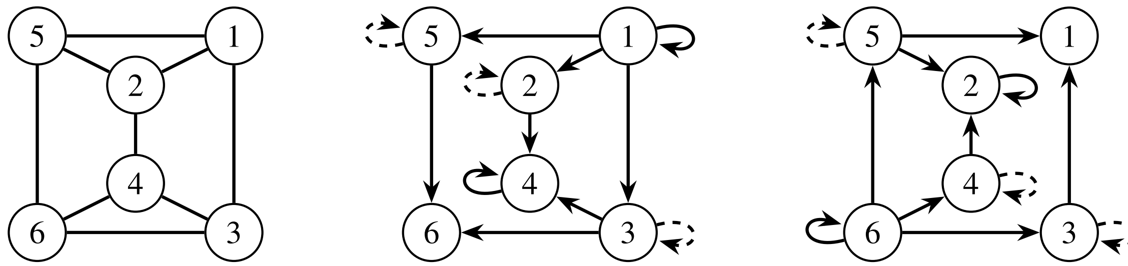

network G

externalities $2 \to 4$

externalities $4 \to 2$

neighbors, similarly present challenges to Pareto efficiency when individuals can trade favors across rounds.

Example 3. *An objective environment has* $G = \{\{1,2\},\{1,3\},\{1,5\},\{2,4\},\{2,5\},\{3,4\},\{3,6\},\{4,6\},\{5,6\}\}$*,* $R = 3$*,* $\Pr(\tau_1 = 1) = \Pr(\tau_1 = \infty) = \Pr(\tau_6 = 1) = \Pr(\tau_6 = \infty) \approx 0.5, \Pr(\tau_2 = 2) = \Pr(\tau_2 = 3) = \Pr(\tau_4 = 2) = \Pr(\tau_4 = 3) \approx 0.5, \Pr(\tau_3 = 2) = \Pr(\tau_3 = \infty) = \Pr(\tau_5 = 2) = \Pr(\tau_5 = \infty) \approx 0.5$*, and* $(q_1, q_2, q_3, q_4, q_5, q_6) = (0.5 + \epsilon, 0.75, 0.9, 0.75, 0.9, 0.5 + \epsilon)$.

Compared to $\bar{\psi}^*$*, there exists an alternative symmetric stationary function of play where individuals 2 and 4 play suboptimally in round 2 to make the other strictly better off in round 3 in a way that both of them are overall better off.*

Proof. The first panel of Figure 4 presents the network and the next two panels illustrate the two key diffusions where an opportunity for mutually beneficial favor exchange exists for individuals 2 and 4. The second panel illustrates the case when $\tau_1 = 1, \tau_6 = \infty$ and $\tau_4 = 3$, which occurs with probability close to 0.125. It replicates the flow of information in Example 1, where under $\bar{\psi}^*$ the induced correlation between the actions of individuals 2 and 3 causes individual 4 to downgrade their consensus actions and disregard their actions altogether. The third panel is a mirrored case of the second panel, where individual 6 plays the role of individual 1, individual 5 plays the role of individual 3 and individuals 2 and 4 swap roles. This occurs with probability close to 0.125, when $\tau_6 = 1, \tau_1 = \infty, \tau_2 = 3$.

Similar to Example 1, individual 2 can make individual 4 strictly better off in round 3 by randomizing between 1 and -1 with equal probabilities when she only hears from individual 1 in round 2. Similarly, individual 4 can make individual 2 strictly better off in round 3 by randomizing between 1 and -1 with equal probabilities when he only hears from individual 6 in round 2. The network structure and the diffusion process are chosen so that such changes do not affect any other individuals. Moreover, as ϵ gets closer to zero, the losses in round 2 incurred by individuals 2 and 4 for randomizing rather than respectively passing on the action of individual 1 and the action of individual 6 shrink to zero. However, their gains from breaking the correlations remain. This means that for sufficiently small ϵ, such changes in the play of individuals 2 and 4 lead to a Pareto improvement over $\bar{\psi}^*$. $\square$

Example 4. *An objective environment has* $G = \{\{1,2\},\{1,3\},\{2,4\},\{3,5\},\{4,5\}\}$, $R = 4$, $\Pr(\tau_1 = 1) \approx 1/3, \Pr(\tau_1 = \infty) \approx 2/3, \Pr(\tau_2 = 2) = \Pr(\tau_2 = \infty) = \Pr(\tau_3 = 2) = \Pr(\tau_3 = \infty) \approx 0.5, \Pr(\tau_4 = 3) = \Pr(\tau_4 = \infty) = \Pr(\tau_5 = 3) = \Pr(\tau_5 = \infty) \approx 0.5$ *and* $(q_1, q_2, q_3, q_4, q_5) = (0.9, 0.5 + \epsilon, 0.5 + \epsilon, 0.7, 0.7)$.

Compared to $\bar{\psi}^*$*, there exists an alternative symmetric stationary function of play where individuals 4 and 5 play suboptimally in round 3 to make the other strictly better off in round 4 in a way that both of them are overall better off.*

Proof. The network is a cycle of length five, as in Example 2. Figure 5 illustrates the key diffusions. The first panel occurs with probability close to 1/3, when $\tau_1 = 1$. In this case both individuals 4 and 5 take action in round 3 without observing the action of the other. In the second panel, information flows from individual 2 to individual 4 and then to individual 5 as in Example 2. In the third panel, information flows from individual 3 to individual 5 and then to individual 4.

The distribution H is chosen so that under $\bar{\psi}^*$,

$$\Pr(a_2 = \theta | \theta, M_4(\tau) = \{2,4\}, r_4(\tau) = 3) \approx 0.5q_1 + 0.5q_2 = 0.7 + 0.5\epsilon > q_4,$$

which means that $a_4 = a_2$ when $M_4(\tau) = \{2,4\}$ and $r_4(\tau) = 3$. If instead individual 4 chooses $a_4 = s_4$ at this conditioning set, individual 5 will do strictly better in the second panel of Figure 5. Symmetrically, under $\bar{\psi}^*$, $a_5 = a_3$ when $M_5(\tau) = \{3,5\}$ and $r_5(\tau) = 3$. If instead individual 5 chooses $a_5 = s_5$ at this conditioning set, then individual 4 will do strictly better in the third panel of Figure 5. The losses to individuals 4 and 5 in round 3 vanish as ϵ gets close to zero while their gains in round 4 do not. Moreover, such changes have no effects on other individuals. This completes the proof that a symmetric feasible

Figure 5: Key diffusions in Example 4

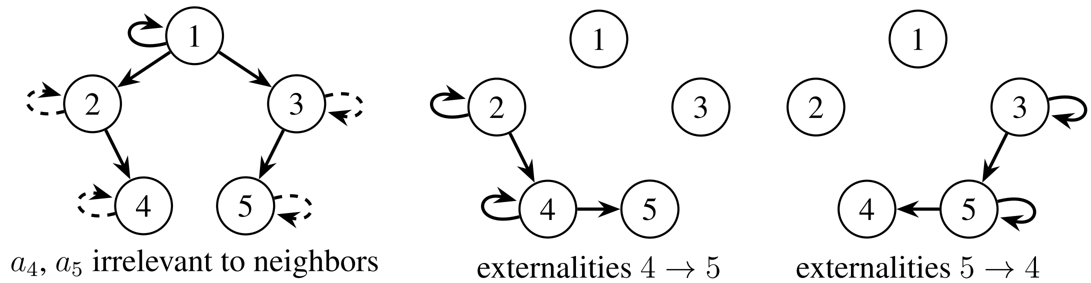

stationary function of play that adopts these changes is a Pareto improvement over $\bar{\psi}^*$. $\square$

6.5 Asymmetric environments

In general, the overall informativeness of multiple sources of information depends on the quality of each source, their correlation, and their biases towards either false positives or false negatives. Intuitively speaking, the restriction of Theorem 2 to clique trees shuts down the second channel of externalities and the restriction to symmetric environments shuts down the third channel. The following two examples illustrate how perfect alignment of interests is hard to achieve in asymmetric environments, even on special networks.

First is an example that shows how strong efficiency could fail even on the simplest tree if the two states are asymmetric.

Example 5. *Consider a setting where* $\Pr(\theta = 1) = 3/4$ *and* $\Pr(\theta = -1) = 1/4$*. An objective environment* $(G, H, (q_i)_{i \in N})$ *has* $G = \{\{1, 2\}\}$*,* $R = 2$*,* $\Pr(\tau_1 = 1) = \Pr(\tau_2 = 1) = \Pr(\tau_1 = 2) = \Pr(\tau_2 = 2) \approx 0.5$ *and* $(q_1, q_2) = (2/3, 2/3)$*. Then there does not exist a feasible stationary function of play that is strongly efficient within the class of feasible stationary functions of play.*

Proof. First, notice that one negative signal is not sufficient to overturn the positive state, since

$$\frac{\Pr(\theta = -1 | s_1 = -1)}{\Pr(\theta = 1 | s_1 = -1)} = \frac{\Pr(\theta = -1 | s_2 = -1)}{\Pr(\theta = 1 | s_2 = -1)} = \frac{(1/4)(2/3)}{(3/4)(1/3)} = \frac{2}{3} < 1,$$

This means that whoever receives only her own private signal optimally chooses the positive action, regardless of the signal realization. This renders her action completely uninformative to the individual who moves next. As a result, individual optimality implies that $a_1 = a_2 = 1$ regardless of the realized diffusion and signals.

However, two negative signals would indicate that the negative state is more likely,

$$\frac{\Pr(\theta = -1|s_1 = s_2 = -1)}{\Pr(\theta = 1|s_1 = s_2 = -1)} = \frac{(1/4)(2/3)(2/3)}{(3/4)(1/3)(1/3)} = \frac{4}{3} > 1.$$

Thus, the second mover would strictly benefit from the first mover reporting her private signal rather than her optimal action. This completes the proof that in every feasible stationary play, there exists an individual that can be made strictly better off at the expense of the other individual. □

Second, even in symmetric objective environments, some individual might benefit from asymmetric play of others, as illustrated by the following example.

Example 6. *An objective environment has* $G = \{\{1, 2\}, \{1, 3\}, \{2, 4\}, \{3, 5\}\}$*,* $\Pr(\tau_1 = 3) \approx 1, \Pr(\tau_2 = 2) = \Pr(\tau_2 = \infty) = \Pr(\tau_3 = 2) = \Pr(\tau_3 = \infty) \approx 0.5, \Pr(\tau_4 = 1) = \Pr(\tau_5 = 1) \approx 1$*, and* $(q_1, q_2, q_3, q_4, q_5) = (0.75, 1, 1, 0.5 + \epsilon, 0.5 + \epsilon)$*.*

There exists a stationary function of play, where the actions of individuals 2 and 3 are biased towards opposite states of the world, that makes individual 1 strictly better off compared to $\bar{\psi}^*$*.*

Proof. Figure 6 illustrates the diffusions: individuals 4 and 5 receive their private signals in round 1 and pass on their private signals to individuals 2 and 3 respectively, who may or may not receive their private signals in round 2; in round 3, individual 1 receives her private signal and observes the actions of individuals 2 and 3. Since G is a tree, individuals 2 and 3's actions, as observed by individual 1, are always conditionally independent. Under $\bar{\psi}^*$, $a_2 = s_4$ when individual 2 does not receive her private signal and $a_2 = s_2$ when she does. Similarly, $a_3 = s_5$ when individual 3 does not receive his private signal and $a_3 = s_3$ when he does. Overall,

$$\begin{aligned}\Pr(a_2 = \theta|\theta, M_1(\tau) = \{1, 2, 3\}, r_1(\tau) = 3) &= \Pr(a_3 = \theta|\theta, M_1(\tau) = \{1, 2, 3\}, r_1(\tau) = 3)\\ &= 0.5(0.5 + \epsilon) + 0.5(1) = 0.75 + 0.5\epsilon.\end{aligned}$$

Figure 6: Key diffusions in Example 6

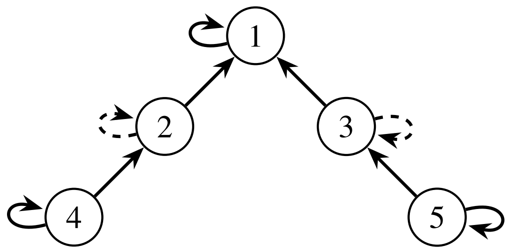

For $\epsilon > 0$ small, the best predictor for individual 1 given (s_1, a_2, a_3) agrees with the majority, thus matching the true state of the world with probability approximately equal to $(3)(0.75)^2(0.25) + (0.75)^3 = \frac{27}{32}$.

Consider an alternative feasible stationary function of play $\bar{\psi}$ where $a_2 = 1$ if individual 2 does not receive her private signal and $a_2 = s_2$ if she does, and $a_3 = -1$ if individual 3 does not receive his private signal and $a_3 = s_3$ if he does. These modifications give

$$\begin{aligned}
&\Pr(a_2 = 1|\theta = 1, M_1(\tau) = \{2,3\}, r_1(\tau) = 3) = 1, \\
&\Pr(a_2 = 1|\theta = -1, M_1(\tau) = \{2,3\}, r_1(\tau) = 3) = 0.5, \\
&\Pr(a_3 = 1|\theta = 1, M_1(\tau) = \{2,3\}, r_1(\tau) = 3) = 0.5, \\
&\Pr(a_3 = 1|\theta = -1, M_1(\tau) = \{2,3\}, r_1(\tau) = 3) = 0.
\end{aligned}$$

That is, from the perspective of individual 1, a_2 is now conclusive of the negative state and a_3 is now conclusive of the positive state. In this case the best predictor for individual 1 agrees with the conclusive news if there is one and agrees with her private signal if the actions of individuals 2 and 3 are inconclusive. This best predictor matches the state of the world with probability $0.5 + (0.5)(0.75) = \frac{7}{8} > \frac{27}{32}$. □

This example illustrates that exposure to biases in opposite directions could help with making better predictions. If the biases were instead in the same direction, then an individual who has no intrinsic bias in either state of the world would mechanically appear as biased in that same direction. Note that two key properties of empiricist learning rules, convergence of learning and long-run individual optimality, do not rely on the environments being symmetric. To capture time-invariant biases, one can modify the smooth decision rules accordingly, and given that convergence always holds, one can then focus on studying the stationary function of play induced by everyone correctly learning her local model. While it is potentially interesting to study environments with biases, how they spread or get neutralized, this direction is deferred to future research.

7 Literature review

The majority of papers about learning on networks focus on learning about a single state of the world. My paper belongs to a relatively small group of papers that model varying states across time. Sethi and Yildiz (2016, 2019) study endogenous communication networks with independently drawn states, where plays across periods connect through individuals' learning about the persistent types of other individuals. Alatas et al. (2016) and Dasaratha et al. (2018) model states that evolve according to an AR(1) process. Individuals take action simultaneously in each period after observing their neighbors' past actions, which are informative about past states and thus informative about the current state. The key feature distinguishing my model from these models is the diffusion process within each period, which randomly decides the order of play and the observation sets.

This per-period feature relates to the social learning literature, where each agent takes action exactly once and sequentially. Classical papers of this literature are Banerjee (1992), Bikhchandani et al. (1992) and Smith and Sørensen (2000); recent papers include, for example, Acemoglu et al. (2011) and Lobel and Sadler (2015). The modeling of each period in my model is more related to these recent papers, but further allows randomness in the order of play. Moreover, the efficiency metrics of interest in this literature concerns the n-th individual in large populations, while my efficiency analysis concerns everyone on the network. The crucial difference between my model and the social learning literature goes beyond these per-period features. While most papers in this literature study perfect Bayesian equilibria under common knowledge of the network topology, my model lets individuals learn relevant information about their environments from repeated interactions.[7]

The construction of empiricist learning rules relates to the literature on fictitious play (Brown, 1951; Fudenberg and Kreps, 1993), where individuals asymptotically believe that the empirical frequency of others' play is their true stationary play. My paper brings this

[7]Some exceptions in the social learning literature include, for example, Wiseman (2009), Eyster and Rabin (2010), Guarino and Jehiel (2013), and Bohren (2016). However, individuals' behaviors in these models are also not founded by learning.

idea of the learning literature to a network game of information aggregation.

Local observability, and to a lesser extent, local knowledge of the network structure, are not new assumptions in the network literature. For example, these assumptions are implicit in DeGroot-style learning models, where each individual takes a weighted average of their neighbors' actions. Li and Tan (2020) takes the localness of knowledge to a model of misspecified learning, where each individual assumes her local network is the full network. McBride (2006, 2008) studies network formation models with imperfect monitoring of other' network relationships, allowing incorrect perceptions of others' links to arise as an equilibrium phenomenon. Breza et al. (2018), from their empirical finding that network knowledge is limited and local, suggest using models with incomplete information of the network structure. Rather than turning to heuristics, misspecifications or richer information structure, I build a model with repeated interactions for robust learning.

Since my model is not closely related to any previous work, its concepts of convergence and quality of information aggregation are not direct analogs of similarly named concepts in other papers. Convergence in my model means that asymptotically individuals' actions depend only on their current observations. Each individual asymptotically learns the long-run conditional distribution of her observation, based on which she forms a belief about the current state of the world given her current observation. Knowledge of the distribution of private observation conditional on the state of the world is exactly what is needed for the analysis of social learning models by, for example, Acemoglu et al. (2011) and Lobel and Sadler (2015). In these models, such local knowledge is immediately implied by common knowledge of the network topology and perfect Bayesian equilibria. On the technical side, the per-period play of these models satisfies two key assumptions for the convergence result in my paper. First is the sequential nature of moves, which is crucial to the iterative argument over action rounds. Second is access to private signals, which allows individuals to back out the conditional distribution of their observation from the unconditional distribution. This means that with some small modifications, empiricist learning rules can be used to provide a learning foundation for these social learning models, relaxing the common knowledge assumption of the network topology and of opponents' play.

Using a model very different from mine, with only one state of the world and individuals communicating their beliefs over time as more information about that state arrives, Li and Tan (2020) identify clique trees as the necessary and sufficient condition for strong efficiency.[8] While this result sounds similar to my result on quality of information aggregation, failure of efficiency in their model must come from inferring mistakes induced by misspecified beliefs about the network structure. In contrast, individuals in my model always correctly learn the conditional distribution of their observations. It is local observability and coarse communication that give rise to an endogenous form of misalignment of interest, challenging social optimality.

8 Conclusion

This paper proposes a novel model of information diffusion and aggregation that allows a separation between learning about the environment and aggregating information about the current state. Empiricist learning rules abstract from the details of the environment, focusing directly on the relevant object of learning, that is, the distribution of private observation conditional on the true state of the world. In stationary environments, empiricist learning rules achieve asymptotic learning of the local environment, and thus asymptotic individual optimality. I study convergence of play and efficiency of information aggregation when all individuals use empiricist learning rules and smooth decision rules, which are technical modifications of optimal decision rules.

In this paper, convergence of play means that asymptotically individuals' actions depend only on their current observations. If convergence holds then the long-run quality of information aggregation, including for example, how likely individuals' actions agree with the true state of the world or agree with each other, can be evaluated at a stationary play. I show that if all individuals adopt empiricist learning rules and smooth decision rules, then play converges for all objective environments. The proof uses an inductive argument on action rounds. Convergence of play in previous rounds means that the conditional distribution of an individual's observation in the current round converges to a stationary distribution. Empiricist learning rules then ensures asymptotic learning of this stationary distribution. Finally, convergence of beliefs leads to convergence of play under smooth decision rules. Note that the conditional distributions of an individual's observation at different pairs of an observation set and an action round are all linked together by the fundamentals of the underlying environment. However, empiricist learning rules essentially

[8] See Proposition 2 of Li and Tan (2020).

treat these conditional distributions separately. While forgoing some information, this feature avoids contamination of learning from later rounds to earlier rounds. This is the exact reason why the inductive argument works.

My analysis on quality of information aggregation focuses on the long-run probability that each individual's action matches the true state of the world. Despite individuals' limited knowledge of the network structure, local observability, selfish and myopic motives, the long-run play induced by empiricist learning rules and smooth decision rules achieves strong efficiency on clique trees. Specifically, it maximizes each individual's probability of matching the state of the world within the class of stationary plays that treat the two states symmetrically. Two key properties of clique trees ensure this result. First, neighbors' actions are conditionally independent, so their overall informativeness is maximized by maximizing the individual informativeness of each neighbor's action. Second, an individual's action is most informative to her neighbors when it is optimal to herself.

I then identify several distinct reasons for why efficiency of information aggregation is likely to fail in general circumstances, including a weak converse to the positive efficiency result. On any network that is not a clique tree, there exist some diffusion process and private signals such that the long-run play induced by empiricist learning rules and smooth decision rules is strictly dominated by some symmetric stationary play. This result generalizes the intuition of two examples that speak to two key properties of clique trees: conditional independence of neighbor's actions and local alignment of interests. In one example when the actions of two neighbors are conditionally correlated, a suboptimal play by one neighbor might break this correlation in a way that improves their overall informativeness. In another example when the flow of information to an individual affects both the quality of an observed neighbor's action and whether her action is observed by another neighbor, a suboptimal decision given the average quality of the action of the former neighbor may be optimal to the latter neighbor. Moreover, these forces similarly present challenges to Pareto efficiency when individuals can trade favors across rounds. While efficiency comparison in my paper focuses on environments without biases, it is also noted that exposure to biases in opposite directions may help with information aggregation.

There are several directions for future research. One direction is to formalize the extension of my convergence result to richer settings, thus providing a learning foundation for a rich class of games of information aggregation. More specifically, my model presents a generalizable two-step framework. In the first step, one would build a model where the game of interest is played repeatedly, and construct learning rules that ensure convergence of learning about relevant elements of the environment. In the second step, one would then analyze the asymptotic play of the multi-period model given individuals' asymptotic learning, thus providing a prediction for the outcome of the one-period game. For example, if the two-step framework is adapted to the context of social learning, it would predict the same outcome as the perfect Bayesian equilibria studied in papers that assume common knowledge of the underlying objective environment. That is, the two-step framework would achieve the same predictions as these papers, while making only minimal assumptions on individuals' knowledge of the environment.

Finally, if my model is taken seriously as a model of how people interact on social platforms, it has the following implications. First, access to independent private news is important for an individual to learn the reliability of information from her neighbors. Second, since social networks are unlikely to be clique trees, social platforms are unlikely to achieve social optimality of information aggregation. Future research could build on these baseline findings and the various challenges to efficiency illustrated in this paper to study, for example, policies that improve observability of information paths and the implications of persistent biases in the context of social platforms.

A Appendix

A.1 Proof that local models are well-defined

To show that every local model is well defined, I show that when i acts in round 2 or later, she could hear from any subset of her neighbors and possibly also receive a private signal.

Lemma A1. *Take any* $E \in \bar{\mathcal{E}}_i$. *For every nonempty set* $M_i \subseteq N_i$ *and* $r \in \{2, ..., R\}$, $\Pr^E(M_i(\tau) = M_i \cup \{i\}, r_i(\tau) = r) > 0$ *and* $\Pr^E(M_i(\tau) = M_i, r_i(\tau) = r) > 0$.

Proof. Construct $\tau \in \{1, ..., R, \infty\}^n$ with $\tau_i = r, \tau_j = r - 1$ for all $j \in M_i$ and $\tau_j = \infty$ for all $j \notin M_i \cup \{i\}$. Then $M_i(\tau) = M_i \cup \{i\}$ and $r_i(\tau) = r$. Similarly, construct $\tau' \in \{1, ..., R, \infty\}^n$ with $\tau'_j = r - 1$ for all $j \in M_i$ and $\tau'_j = \infty$ for all $j \notin M_i$. Then $M_i(\tau') = M_i$ and $r_i(\tau') = r$. The claim follows from the assumption that the distribution of the signal-timing vector has full support. □

A.2 Proof of Lemma 1

To see the second equation of part 2, recall that private signals are independent of the diffusion process and across individuals, conditional on the state of the world. Take any nonempty set $M_i \subseteq N_i$ and $r \in \{2, ..., R\}$. For any $\tau \in \{1, ..., R, \infty\}^n$ such that $M_i(\tau) = M_i \cup \{i\}$ and $r_i(\tau) = r$, the actions of i's neighbors in round $r - 1$ are not affected by i's private signal. That is, for such τ,

$$\mathrm{Pr}^E(a_{M_i}, s_i|\theta, \tau) = \mathrm{Pr}^E(a_{M_i}|\theta, \tau)\mathrm{Pr}^E(s_i|\theta, \tau) = \mathrm{Pr}^E(a_{M_i}|\theta, \tau)\mathrm{Pr}^E(s_i|\theta).$$

Then,

$$\begin{aligned}
&f_i^E(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r) \\
=& \sum_{\tau: M_i(\tau)=M_i\cup\{i\}, r_i(\tau)=r} \mathrm{Pr}^E(a_{M_i}, s_i|\theta, \tau)\mathrm{Pr}^E(\tau|M_i(\tau) = M_i \cup \{i\}, r_i(\tau) = r) \\
=& \mathrm{Pr}(s_i|\theta) f_i^E(a_{M_i}|\theta, M_i \cup \{i\}, r).
\end{aligned}$$

It follows that

$$\begin{aligned}
f_i^E(a_{M_i}, s_i|M_i \cup \{i\}, r) &= \sum_{\theta\in\{-1,1\}} \mathrm{Pr}(\theta) \mathrm{Pr}(s_i|\theta) f_i^E(a_{M_i}|\theta, M_i \cup \{i\}, r) \\
&= \sum_{\theta\in\{-1,1\}} \mathrm{Pr}(s_i, \theta) f_i^E(a_{M_i}|\theta, M_i \cup \{i\}, r).
\end{aligned}$$

Furthermore, the matrix $(\mathrm{Pr}^E(s_i, \theta))_{s_i,\theta\in\{-1,1\}} = \frac{1}{2}(\mathrm{Pr}^E(s_i|\theta))_{s_i,\theta\in\{-1,1\}}$ is full rank since $q_i > 1/2$. This completes the proof of part 1 of the lemma.

Finally, the first equation of part 2 relies on the assumption that the arrivals of private signals are independent across individuals. Conditional on i not having received her private signal by round $r-1$, the actions of those neighbors of i that act in round $r-1$ depend only on the arrivals and realizations of the private signals of individuals other than i. Formally, take any two signal-timing vectors τ and τ' such that $\tau_j = \tau'_j$ for all $j \neq i$, $\tau_i \geq r, \tau'_i \geq r$ and $M_i = M_i(\tau)\backslash\{i\}$, it holds that $\mathrm{Pr}^E(a_{M_i}|\theta, \tau, s) = \mathrm{Pr}^E(a_{M_i}|\theta, \tau', s)$ for all $a_{M_i} \in$

$\{-1,1\}^{|M_i|}$. This implies

$$
\begin{aligned}
& f_i^E(a_{M_i}|\theta, M_i, r) \\
=& \frac{\sum_{\tau: M_i(\tau)=M_i} \sum_s H_{-i}(\tau_{-i}) H_i(\tau_i) \mathrm{Pr}^E(s|\theta) \mathrm{Pr}^E(a_{M_i}|\theta, \tau, s)}{\sum_{\tau: M_i(\tau)=M_i} H_{-i}(\tau_{-i}) H_i(\tau_i)} \\
=& \frac{\mathrm{Pr}^E(\tau_i > r) \sum_{\tau_{-i}: M_i(\tau_{-i}, \infty)=M_i} \sum_s H_{-i}(\tau_{-i}) \mathrm{Pr}^E(s|\theta) \mathrm{Pr}^E(a_{M_i}|\theta, (\tau_{-i}, \infty), s)}{\mathrm{Pr}^E(\tau_i > r) \sum_{\tau_{-i}: M_i(\tau_{-i}, \infty)=M_i} H_{-i}(\tau_{-i})} \\
=& \frac{\mathrm{Pr}^E(\tau_i = r) \sum_{\tau_{-i}: M_i(\tau_{-i}, r)\setminus\{i\}=M_i} \sum_s H_{-i}(\tau_{-i}) \mathrm{Pr}^E(s|\theta) \mathrm{Pr}^E(a_{M_i}|\theta, (\tau_{-i}, r), s)}{\mathrm{Pr}^E(\tau_i = r) \sum_{\tau_{-i}: M_i(\tau_{-i}, r)\setminus\{i\}=M_i} H_{-i}(\tau_{-i})} \\
=& f_i^E(a_{M_i}|M_i \cup \{i\}, r, \theta).
\end{aligned}
$$

A.3 Proof of Proposition 1

Notice that by Lemma 1, there is a one-to-one mapping $\Phi : (f_i^E(.|M_i \cup \{i\}, r))_{M_i,r} \mapsto f_i^E$ between the stationary unconditional distributions of an individual's observation (when she receives her private signal) and her local model. Moreover, this mapping is linear.

By the Strong Law of Large Numbers,

$$
\mathrm{Pr}^E\left(\lim_{t\to\infty} \|(\hat{f}_i(.|M_i \cup \{i\}, r)(h_i^t))_{M_i,r} - (f_i^E(.|M_i \cup \{i\}, r))_{M_i,r}\| = 0\right) = 1.
$$

The linearity of Φ then implies that for every $\epsilon > 0$, there exists $\epsilon' > 0$ such that

$$
\begin{aligned}
1 =& \mathrm{Pr}^E\left(\lim_{t\to\infty} \mathbf{1}\left\{\Phi(B^{\epsilon'}((\hat{f}_i(.|M_i \cup \{i\}, r)(h_i^t))_{M_i,r})) \subseteq B^{\epsilon}(\Phi((f_i^E(.|M_i \cup \{i\}, r))_{M_i,r}))\right\} = 1\right) \\
=& \mathrm{Pr}^E\left(\lim_{t\to\infty} \mathbf{1}\left\{F_i^{\epsilon'}(h_i^t) \subseteq B^{\epsilon}(f_i^E)\right\} = 1\right).
\end{aligned}
$$

Since $\lim_{t\to\infty} \gamma_i(h_i^t)(F_i^{\epsilon'}(h_i^t)) = 1$ by the definition of empiricist learning rules, part 1 of the proposition follows.

For part 2, construct a stationary belief-updating rule $\bar{\beta}_i^* : h_{it} \to b_i \in \Delta(\{-1,1\})$ such that for all observation h_{it} and all $\theta \in \{-1,1\}$,

$$
\bar{\beta}_i^*(h_{it})(\theta) = \frac{\mathrm{Pr}^E(h_{it}|\theta_t = \theta)}{\mathrm{Pr}^E(h_{it}|\theta_t = \theta) + \mathrm{Pr}^E(h_{it}|\theta_t = -\theta)}.
$$

It follows from part 1 that

$$\Pr^E\left(\lim_{t\to\infty}\|\beta_i(\tilde{h}_i^{t-1}, h_{it})(\theta_t) - \bar{\beta}_i^*(h_{it})\| = 0\right) = 1.$$

Recall that σ_i^η is a technical modification of the symmetric optimal stationary decision rule $\bar{\sigma}_i^*$, converging to $\bar{\sigma}_i^*$ as $t \to \infty$. Moreover, the decaying smoothing parameter of σ_i^η ensures that convergence of beliefs leads to convergence of play. Thus,

$$\Pr^E\left(\lim_{t\to\infty}\left|\sigma_i^\eta(\beta_i(\tilde{h}_i^{t-1}, h_{it}))(\theta_t) - \bar{\sigma}_i^*(\bar{\beta}_i^*(h_{it}))(\theta_t)\right| = 0\right) = 1.$$

This implies part 2 of the proposition.

A.4 The inductive step in Theorem 1

Lemma A2. *Fix environment* $E = (G, H, (q_i)_{i\in N}, (\beta_i)_{i\in N}, (\sigma_i^\eta)_{i\in N})$ *where for each* i, β_i *is an empiricist belief-updating rule and* σ_i^η *is a smooth decision rule. Consider an individual* i*, a nonempty subset* $M_i \subseteq N_i$ *and a round* $r \in \{2, ..., R\}$*. If there exists some stationary function* $\bar{\psi}$ *such that for every* $j \in M_i$*,*

$$\Pr^E\left(\lim_{t\to\infty}\mathbf{1}\{M_i(\tau_t) = M_i \cup \{i\}, r_i(\tau_t) = r\}\|\psi_j(\omega_t, \tilde{h}^{t-1}, t) - \bar{\psi}_j(\omega_t)\| = 0\right) = 1,$$

then whenever i*'s observation set is* M_i *or* $M_i \cup \{i\}$ *and* i*'s action round is* r*,* i*'s belief and play converge. That is,*

1) for some unconditional distribution $f_i(.|M_i \cup \{i\}, r)$*, it holds for all* $\epsilon > 0$ *that*

$$\Pr^E\left(\lim_{t\to\infty}\gamma_i(h_i^t)\left(\{f_i' \in \mathcal{F}_i : \|f_i'(.|M_i \cup \{i\}, r) - f_i(.|M_i \cup \{i\}, r)\| > \epsilon\}\right) = 0\right) = 1;$$

2) for some stationary belief-updating rule $\bar{\beta}_i^*$*,*

$$\Pr^E\left(\lim_{t\to\infty}\sum_{a_{M_i}, s_i}\left[\begin{array}{c}\mathbf{1}\{h_{it} = (a_{M_i}, s_i, r) \text{ or } h_{it} = (a_{M_i}, r)\} \\ \times\|\beta_i(\tilde{h}_i^{t-1}, h_{it}) - \bar{\beta}_i^*(h_{it})\|\end{array}\right] = 0\right) = 1;$$

3) and for some stationary function $\bar{\psi}_i^*$ of i's play,

$$\Pr^E\left(\lim_{t\to\infty}\sum_{a_{M_i},s_i}\left[\begin{array}{c}\mathbf{1}\{M_i(\tau_t)=M_i \textit{ or } M_i(\tau_t)=M_i\cup\{i\},\ \textit{and } r_i(\tau_t)=r\}\\ \times\|\psi_i(\omega_t,\tilde{h}^{t-1},t)-\bar{\psi}_i^*(\omega_t)\|\end{array}\right]=0\right)=1.$$

Proof. From $\bar{\psi}_{M_i}$, construct for each $a_{M_i}\in\{-1,1\}^{|M_i|}$ and $s_i\in\{-1,1\}$,

$$f_i(a_{M_i},s_i,r|M_i\cup\{i\},r)=\Pr^E(\bar{\psi}_{M_i}(\omega)=a_{M_i},s_i|M_i(\tau)=M_i\cup\{i\},r_i(\tau)=r).$$

By Lemma 1, for each nonempty set $M_i\subseteq N_i$ and action round $r\in\{2,...,R\}$, there is a one-to-one linear mapping $\phi^{M_i,r}:f_i(.|M_i\cup\{i\},r)\mapsto(f_i(.|\theta,M_i\cup\{i\},r),f_i(.|\theta,M_i,r))_{\theta\in\{-1,1\}}$.

Construct a hypothetical empirical distribution that takes draws from $f_i(.|M_i\cup\{i\},r)$ such that for all $a_{M_i}\in\{-1,1\}^{|M_i|}$, $s_i\in\{-1,1\}$ and history ω^t of system states,

$$\tilde{f}_i(a_{M_i},s_i|M_i\cup\{i\},r)(\omega^t)=\frac{\sum_{t'=1}^t\mathbf{1}\{\bar{\psi}_{M_i}(\omega_{t'})=a_{M_i},s_{it'}=s_i,r_i(\tau_{t'})=r\}}{\sum_{t'=1}^t\mathbf{1}\{M_i(\tau_{t'})=M_i\cup\{i\},r_i(\tau_{t'})=r\}}$$

if the denominator is positive, otherwise set $\tilde{f}_i(a_{M_i},s_i|M_i\cup\{i\},r)(\omega^t)=1/(2^{|M_i|+1})$.

By the inductive hypothesis,

$$\Pr^E\left(\lim_{t\to\infty}\|\hat{f}_i(.|M_i\cup\{i\},r)(h_i^t)-\tilde{f}_i(.|M_i\cup\{i\},r)(\omega^t)\|=0\right)=1.$$

Moreover, by the Strong Law of Large Numbers,

$$\Pr^E\left(\lim_{t\to\infty}\|\tilde{f}_i(.|M_i\cup\{i\},r)(\omega^t)-f_i(.|M_i\cup\{i\},r)\|=0\right)=1.$$

It follows that

$$\Pr^E\left(\lim_{t\to\infty}\|\hat{f}_i(.|M_i\cup\{i\},r)(h_i^t)-f_i(.|M_i\cup\{i\},r)\|=0\right)=1.$$

By an argument analogous to the proof of Proposition 1 and using the linearity of the mapping $\phi^{M_i,r}$, it can then be shown that γ_i asymptotically puts zero probability on local models that induce an unconditional distribution at observation set $M_i\cup\{i\}$ and action round r of some ϵ-distance away from that induced by f_i. That is, part 1 of this lemma holds.

Construct $\bar{\beta}_i^*$ such that for every $a_{M_i} \in \{-1,1\}^{|M_i|}$, $s_i \in \{-1,1\}$ and $\theta \in \{-1,1\}$,

$$\bar{\beta}_i^*(a_{M_i}, s_i, r)(\theta) = \frac{f_i(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r)}{f_i(a_{M_i}, s_i|\theta, M_i \cup \{i\}, r) + f_i(a_{M_i}, s_i|-\theta, M_i \cup \{i\}, r)}$$

and

$$\bar{\beta}_i^*(a_{M_i}, r)(\theta) = \frac{f_i(a_{M_i}|\theta, M_i, r)}{f_i(a_{M_i}|\theta, M_i, r) + f_i(a_{M_i}|-\theta, M_i, r)}.$$

Such a stationary belief-updating rule satisfies part 2 of the lemma.

To see part 3, recall that smooth decision rule σ_i^η converges to the symmetric optimal stationary decision rule $\bar{\sigma}_i^*$ and ensures that convergence of beliefs leads to convergence of play. Construct a stationary function of i's play, $\bar{\psi}_i^*$, as following. If ω induces for i observation (a_{M_i}, s_i, r), set

$$\bar{\psi}_i^*(\omega) = \begin{cases} -1 & \text{if } \zeta_i \leq \bar{\sigma}_i^*(\bar{\beta}_i^*(a_{M_i}, s_i, r))(-1) \\ 1 & \text{otherwise.} \end{cases}$$

If ω induces for i observation (a_{M_i}, r), set

$$\bar{\psi}_i^*(\omega) = \begin{cases} -1 & \text{if } \zeta_i \leq \bar{\sigma}_i^*(\bar{\beta}_i^*(a_{M_i}, r))(-1) \\ 1 & \text{otherwise.} \end{cases}$$

Then $\bar{\psi}^*$ satisfies part 3 of the lemma. □

Finally, notice that the case when i only receives her private signal is similar to the base case of $r = 1$. In this case when $h_{it} = (s_{it}, r)$, the individual holds correct belief about the current state, that is, for all $\theta \in \{-1, 1\}$,

$$\beta_i(\tilde{h}_i^{t-1}, h_{it})(\theta) = \bar{\beta}_i^*(s_{it}, r)(\theta) = q_i^{\mathbf{1}\{s_i=\theta\}}(1-q_i)^{\mathbf{1}\{s_i\neq\theta\}},$$

regardless of history $\tilde{h}_i^{t-1}$. Moreover, smooth decision rules imply that asymptotically, an individual i facing such history reports her private signal. For all ω that induces observation (s_i, r) for individual i, set $\bar{\psi}_i^*(\omega) = s_i$.

Together with parts 2 and 3 of Lemma A2, this final argument completes the inductive step in the proof of parts 1 and 2 of Theorem 1. To see part 3 of Theorem 1, collect part 1 of Lemma $A2$ over all possible observation sets and action rounds of each individual i.

A.5 Proof of the equivalence in Definition 7

It is easy to see that for any given objective environment, a profile of symmetric belief-updating rules and a profile of symmetric decision rules induce a symmetric feasible stationary function of play. For the other direction, fix an objective environment and a symmetric feasible stationary function of play $\bar{\psi}$. Since $\bar{\psi}$ is feasible, it can be induced by some profile $(\bar{\beta}'_i)_{i\in N}$ of stationary belief-updating rules and some profile $(\bar{\sigma}'_i)_{i\in N}$ of stationary decision rules. Consider an observation (a_{M_i}, r) of individual i, under symmetry of $\bar{\psi}$,

$$\begin{aligned}&\{(\tau, s, \zeta) : (\tau, s, \zeta) \text{ induces observation } (a_{M_i}, r) \text{ for } i\}\\ =&\{(\tau, s, \zeta) : (\tau, -s, \mathbf{1} - \zeta) \text{ induces observation } (-a_{M_i}, r) \text{ for } i\}.\end{aligned}$$

Call this set $\tilde{\Omega}$. By the symmetry of $\bar{\psi}_i$,

$$\begin{aligned}\bar{\sigma}'_i(\bar{\beta}'_i(a_{M_i}, r))(1) &= \Pr(\bar{\psi}_i(\tau, s, \zeta) = 1 | (\tau, s, \zeta) \in \tilde{\Omega})\\ &= \Pr(\bar{\psi}_i(\tau, -s, \mathbf{1} - \zeta) = -1 | (\tau, s, \zeta) \in \tilde{\Omega})\\ &= \bar{\sigma}'_i(\bar{\beta}'_i(-a_{M_i}, r))(-1).\end{aligned}$$

Similarly, $\bar{\sigma}'_i(\bar{\beta}'_i(a_{M_i}, s_i, r))(1) = \bar{\sigma}'_i(\bar{\beta}'_i(-a_{M_i}, -s_i, , r))(-1)$ for every history (a_{M_i}, s_i, r).

Now, construct symmetric functions $(\bar{\beta}_i)_{i\in N}$ and $(\bar{\sigma}_i)_{i\in N}$ from $(\bar{\beta}'_i)_{i\in N}$ and $(\bar{\sigma}'_i)_{i\in N}$. For each pair $(a_{M_i}, -a_{M_i})$, set $\bar{\beta}_i(a_{M_i}, r) = \bar{\beta}'_i(a_{M_i}, r)$ and $\bar{\sigma}_i(\bar{\beta}_i(a_{M_i}, r)) = \bar{\sigma}'_i(\bar{\beta}'_i(a_{M_i}, r))$; then set $\bar{\beta}_i(-a_{M_i}, r) = \mathbf{1} - \bar{\beta}_i(a_{M_i}, r)$ and $\bar{\sigma}_i(\bar{\beta}_i(-a_{M_i}, r)) = \bar{\sigma}'_i(\bar{\beta}'_i(-a_{M_i}, r))$. Perform similar construction for cases when i receives her private signal. For observations that never arise, the beliefs induced by the belief-updating rules at such observations can be determined arbitrarily and symmetrically. Similarly for beliefs that never arise, the decision rules can be determined arbitrarily and symmetrically at such beliefs. By construction $(\bar{\beta}_i)_{i\in N}$ and $(\bar{\sigma}_i)_{i\in N}$ induce $\bar{\psi}$ and they are symmetric.

A.6 Proof of two technical properties of clique trees

The first property follows immediately from the existence of a unique shortest path connecting any two individuals in G. To see the second property, suppose by contradiction that for some $i \in N, j \in N_i$ and $k \in N_j \backslash (N_i \cup \{i\})$, there exists some $k' \in IF_{k\backslash j}$ but $k' \notin IF_{j\backslash i}$. This means that $l(k', i) < l(k', j) + l(j, i)$, that is, the shortest path connecting k' to i does not go through j. This path and the shortest path that goes from k' to i through j form a cycle. This cycle includes k, but it is not complete because k and i are not linked. This contradicts that G is a clique tree.

A.7 Proof of Lemma 3

By property 1 of Lemma 2, $\{IF_{j\backslash i}\}_{j\in N_i}$ are mutually exclusive. Consider the case that $i \notin M_i$ and enumerate individuals in M_i by $j_1, ..., j_{|M_i|}$. By the conditional independence of private signals and the independence of the instrumental variables, $\Pr(\bar{\psi}_{M_i}|\theta, \tau) = \prod_{j\in M_i} \Pr(\bar{\psi}_j|\theta, \tau)$ for any τ such that $M_i(\tau) = M_i$ and $r_i(\tau) = r$.

Next, notice that the event that i observes from M_i and acts in round r can be rewritten as a join of independent events

$$\{\tau : M_i(\tau) = M_i, r_i(\tau) = r\} = \left(\cap_{j\in M_i} \left\{\tau : \min_{k\in IF_{j\backslash i}} l(j,k) + \tau_k = r-1\right\}\right) \cap \left(\cap_{j\in N_i\backslash M_i} \left\{\tau : \min_{k\in IF_{j\backslash i}} l(j,k) + \tau_k \geq r\right\}\right) \cap \{\tau : \tau_i > r\}.$$

For all $j \in M_i$, let $\mathcal{A}_{ij} = \{\tau_{IF_{j\backslash i}} : \min_{k\in IF_{j\backslash i}} l(j,k) + \tau_k = r - 1\}$. Also, let

$$\mathcal{A}_i^- = \{\tau_{N\backslash(\cup_{j\in M_i} IF_{j\backslash i})} : \tau_i > r \text{ and } \min_{k\in IF_{j\backslash i}} l(j,k) + \tau_k \geq r \text{ for all } j \in N_i\backslash M_i\}.$$

It follows from the above decomposition that

$$\begin{aligned} &\Pr(\tau | M_i(\tau) = M_i, r_i(\tau) = r) \\ =&\Pr\left(\tau_{N\backslash(\cup_{j\in M_i} IF_{j\backslash i})} | \tau_{N\backslash(\cup_{j\in M_i} IF_{j\backslash i})} \in \mathcal{A}_i^-\right) \prod_{j\in M_i} \Pr(\tau_{IF_{j\backslash i}} | \tau_{IF_{j\backslash i}} \in \mathcal{A}_{ij}). \end{aligned}$$

Thus,

$$
\begin{aligned}
&\Pr(\bar{\psi}_{M_i}(\omega)|M_i(\tau)=M_i, r_i(\tau)=r) \\
=& \sum_{\tau: M_i(\tau)=M_i, r_i(\tau)=r} \Pr(\bar{\psi}_{M_i}(\omega)|\theta,\tau)\Pr(\tau|M_i(\tau)=M_i, r_i(\tau)=r) \\
=& \sum_{\tau_{IF_{j_1}\backslash i}\in\mathcal{A}_{ij_1}} \cdots \sum_{\tau_{IF_{j_{|M_i|}}\backslash i}\in\mathcal{A}_{ij_{|M_i|}}} \\
& \sum_{\tau_{N\backslash(\cup_{j\in M_i}IF_{j\backslash i})}\in\mathcal{A}_i^-} \left(\begin{array}{c} \prod_{j\in M_i}\Pr(\bar{\psi}_j(\omega)|\theta,\tau_{IF_{j\backslash i}}, M_i(\tau)=M_i, r_i(\tau)=r) \\ \times\prod_{j\in M_i}\Pr(\tau_{IF_{j\backslash i}}|\tau_{IF_{j\backslash i}}\in\mathcal{A}_{ij}) \\ \times\Pr\left(\tau_{N\backslash(\cup_{j\in M_i}IF_{j\backslash i})}|\tau_{N\backslash(\cup_{j\in M_i}IF_{j\backslash i})}\in\mathcal{A}_i^-\right) \end{array} \right) \\
=& \prod_{j\in M_i}\Pr(\bar{\psi}_j(\omega)|M_i(\tau)=M_i, r_i(\tau)=r).
\end{aligned}
$$

The proof for the case when $i \in M_i$ is similar.

A.8 Proof of Lemma 4

When $r = 2$, it must be that $M_j(\tau) = \{j\}$ and thus $\tilde{q}_j^{M_i,r} = \Pr(\bar{\psi}_j(\omega) = \theta|\theta, M_j(\tau) = \{j\}, r_j(\tau) = 1) = q_j$. Now consider the case that $r \geq 3$. Since $j \in M_i(\tau)$, $(M_j(\tau)\backslash\{j\}) \cap (N_i \cup \{i\}) = \varnothing$. By the second property of Lemma 2 and the full support assumption of H, any observation set of j that satisfies the above condition is feasible, that is, $\Pr(M_j(\tau) = M_j, M_i(\tau) = M_i, r_i(\tau) = r) > 0$ for all $M_j \subseteq (N_j\backslash(N_i \cup \{i\})) \cup \{j\}$ and $r \geq 3$. Then the quality of j's action observed by i at observation set (M_i, r) is

$$
\tilde{q}^{M_i,r} = \sum_{M_j\subseteq(N_j\backslash(N_i\cup\{i\}))\cup\{j\}} \left(\begin{array}{c} \Pr(M_j(\tau)=M_j|M_i(\tau)=M_i, r_i(\tau)=r) \\ \times\Pr(\bar{\psi}_j(\omega)=\theta|\theta, M_j(\tau)=M_j, M_i(\tau)=M_i, r_i(\tau)=r) \end{array} \right).
$$

It remains to show that for all $M_j \subseteq (N_j\backslash(N_i \cup \{i\})) \cup \{j\}$,

$$
\begin{aligned}
&\Pr(\bar{\psi}_j(\omega)=\theta|\theta, M_j(\tau)=M_j, M_i(\tau)=M_i, r_i(\tau)=r) \\
=&\Pr(\bar{\psi}_j(\omega)=\theta|\theta, M_j(\tau)=M_j, r_j(\tau)=r-1).
\end{aligned}
$$

For such M_j, let

$$\tilde{\mathcal{A}}_{ij} = \{\tau_{\cup_{k\in M_j} IF_{k\setminus j}} : \exists \tau' \text{ s.t. } \tau'_{\cup_{k\in M_j} IF_{k\setminus j}} = \tau_{\cup_{k\in M_j} IF_{k\setminus j}}, M_j(\tau') = M_j, M_i(\tau') = M_i, r_i(\tau') = r\},$$
$$\tilde{\mathcal{A}}_{j} = \{\tau_{\cup_{k\in M_j} IF_{k\setminus j}} : \exists \tau' \text{ s.t. } \tau'_{\cup_{k\in M_j} IF_{k\setminus j}} = \tau_{\cup_{k\in M_j} IF_{k\setminus j}}, M_j(\tau') = M_j, r_j(\tau') = r-1\}.$$

By Lemma 2, $\cup_{k\in M_j} IF_{k\setminus j} \subseteq IF_{j\setminus i}$. Therefore, conditional on $M_j(\tau) = M_j$ and $r_j(\tau) = r-1$, the observation set and action round of i depends further only on $\tau_{IF_{j'\setminus i}}$ for $j' \in (N_i \cup \{i\})\setminus\{j\}$. This means that $\tilde{\mathcal{A}}_{ij} = \tilde{\mathcal{A}}_j$. Furthermore, conditional on j hearing from M_j in round $r-1$, $\bar{\psi}_j$ depends only on $(\tau_{j'}, s_{j'}, \zeta_{j'})$ of $j' \in \cup_{k\in M_j} IF_{k\setminus j}$ and on ζ_j. Then,

$$\begin{aligned}
&\Pr(\bar{\psi}_j(\omega)|M_i(\tau) = M_i, M_j(\tau) = M_j, r_i(\tau) = r) \\
=& \sum_{\tau_{\cup_{k\in M_j} IF_{k\setminus j}} \in \tilde{\mathcal{A}}_{ij}} \begin{pmatrix} \Pr(\tau_{\cup_{k\in M_j} IF_{k\setminus j}} | M_j(\tau) = M_j, M_i(\tau) = M_i, r_i(\tau) = r) \\ \times \Pr(\bar{\psi}_j(\omega)|\tau_{\cup_{k\in M_j} IF_{k\setminus j}}, M_j(\tau) = M_j, M_i(\tau) = M_i, r_i(\tau) = r) \end{pmatrix} \\
=& \sum_{\tau_{\cup_{k\in M_j} IF_{k\setminus j}} \in \tilde{\mathcal{A}}_{j}} \begin{pmatrix} \Pr(\tau_{\cup_{k\in M_j} IF_{k\setminus j}} | M_j(\tau) = M_j, r_j(\tau) = r-1) \\ \times \Pr(\bar{\psi}_j(\omega)|\tau_{\cup_{k\in M_j} IF_{k\setminus j}}, M_j(\tau) = M_j, r_j(\tau) = r-1) \end{pmatrix} \\
=& \Pr(\bar{\psi}_j(\omega)|M_j(\tau) = M_j, r_j(\tau) = r-1).
\end{aligned}$$

This completes the proof of Lemma 4.

A.9 Proof of Theorem 3

Take a graph G. Denote by $\mathcal{C}(i_1, ..., i_L)$ a cycle of size L where $\{i_l, i_{l+1}\} \in G$ for all $l = 1, ..., L$ and $i_{L+1} = i_1$. This cycle is complete if any two individuals of this cycle are linked. A chord of a cycle is a link between two non-consecutive individuals. A cycle of size at least four is chordless if any two non-consecutive individuals are not linked. Furthermore, a graph is chordal if all cycles of size at least four have a link between two non-consecutive individuals.

Suppose that G is not a clique tree. Then G must have at least a cycle that is not complete. Take a smallest incomplete cycle $\mathcal{C}(i_1, ..., i_L)$ and refer to this subgraph of G as G'. There are three cases.

Case 1: G' is a chordless cycle of size $L = 2m$ for $m \geq 2$.

Figure 7: The key diffusion in a generalization of Example 1

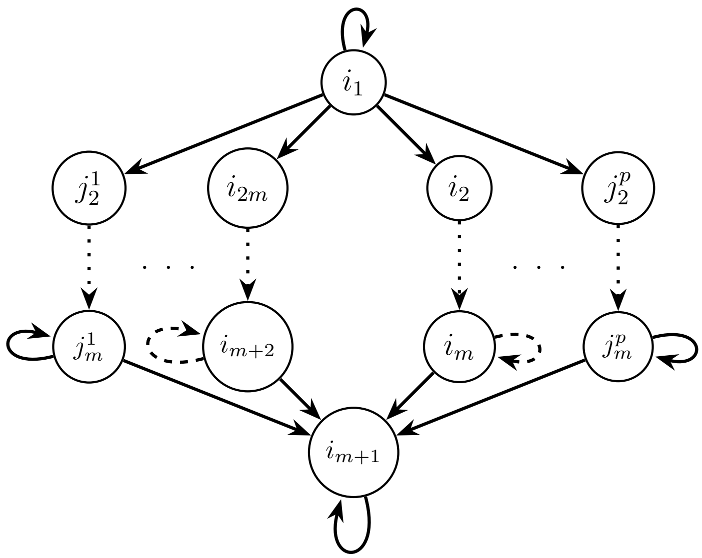

$\Pr(\tau_{i_{m+1}} = m+1) \approx 1$ when p is even

Consider a potential path $(i_1, j_2, ..., j_{m'}, i_{m+1})$ from i_1 to i_{m+1} different from the paths $(i_1, i_2, ..., i_m, i_{m+1})$ and $(i_1, i_{2m}, ..., i_{m+2}, i_{m+1})$. Let I^* collect individuals in $\{i_1, ..., i_{m+1}\}$ that appear in the first path. If $m' < m$, there must exist $i_{m_1}, i_{m_2} \in I^*$ such that $m_2 \geq m_1 + 2$. Then for some m_1' and m_2' such that $2 \leq m_1', m_2' \leq m'$, there is a cycle of size strictly less than L, $\mathcal{C}(i_{m_1}, i_{m_1+1}, ..., i_{m_2-1}, i_{m_2}, j_{m_2'}, j_{m_2'-1}, ..., j_{m_1'+1}, j_{m_1'})$. Moreover, this cycle is incomplete because i_{m_1} and i_{m_2} are not linked. This contradicts that G' is the smallest incomplete cycle of G. If $m' = m$ and I^* contains individuals other than i_1 and i_{m+1}, then a similar argument provides a contradiction to G' being the smallest incomplete cycle. This concludes that the shortest paths connecting i_1 and i_{m+1} include only paths of length m with no joint elements other than i_1 and i_{m+1}. Suppose that there are $p \geq 0$ such paths other than $(i_1, i_2, ..., i_m, i_{m+1})$ and $(i_1, i_{2m}, ..., i_{m+2}, i_{m+1})$. Denote such a path by $(i_1, j_2^{p'}, ..., j_m^{p'}, i_{m+1})$ for $p' = 1, ..., p$.

Consider first the case when p is even. Construct a vector of signal qualities with $q_{i_1} = 0.5 + \epsilon$, $q_{i_m} = q_{i_{m+2}} = 0.8$ and $q_{i_{m+1}} = q_{j_m^{p'}} = 0.75$ for all $p' = 1, ..., p$. Construct the distribution of the signal-timing vector so that $\Pr(\tau_{i_1} = 1) \approx 1$, $\Pr(\tau_{i_m} = m) = \Pr(\tau_{i_m} = \infty) = \Pr(\tau_{i_{m+2}} = m) = \Pr(\tau_{i_{m+2}} = \infty) \approx 0.5$, $\Pr(\tau_{i_{m+1}} = m+1) = \Pr(\tau_{j_m^{p'}} = m) \approx 1$

for all $p' = 1, ..., p$, and $\Pr(\tau_i = \infty) \approx 1$ for all other i.

Figure 7 illustrates the subgraph of all shortest paths connecting i_1 and i_{m+1}, and plots the key diffusion for this example when p is even. In this diffusion, i_1 receives his private signal in round 1. He then passes this signal, $a_{i_1} = s_{i_1}$, onto i_2, i_{2m} and $j_2^{p'}$ for $p' = 1, ..., p$. In round 2, these individuals then simply pass a_1 onto i_3, i_{2m-1} and $j_3^{p'}$ for $p' = 1, ..., p$ respectively. This goes on until round m when $j_m^{p'}$ for $p' = 1, ..., p$ receive their private signals, which are more informative than a_{i_1}. They thus choose $a_{j_m^{p'}} = s_{j_m^{p'}}$. In round m, i_{m+2} and i_m may or may not receive their private signals. If i_{m+2} receives his private signal then he chooses $a_{i_{m+2}} = s_{i_{m+2}}$, otherwise $a_{i_{m+2}} = a_{i_{m+3}} = s_{i_1}$. Similarly, if i_m receives his private signal then he chooses $a_{i_m} = s_{i_m}$, otherwise $a_{i_m} = a_{i_{m-1}} = s_{i_1}$. In round $m+1$, i_{m+1} receives his private signal and hears from $J \cup \{i_{m+2}, i_m\}$, where $J = \{j_m^{p'} : p' = 1, ..., p\}$. Let $M_i = J \cup \{i_{m+2}, i_m, i_{m+1}\}$.

Notice that when $a_{i_{m+2}} \neq a_{i_m}$, the optimal decision of i_{m+1} depends only on his private signal and the actions of individuals in J. The key is whether consensus posts by i_{m+2} and i_m could overturn a decision based solely on i_{m+1}'s private signal and actions of her neighbors in J. The informativeness of $(s_{i_{m+1}}, a_J) = (s_{i_{m+1}}, s_J)$ is summarized by

$$\frac{\Pr(s_{i_{m+1}}, a_J|\theta, M_i(\tau) = M_i, r_i(\tau) = m+1)}{\Pr(-s_{i_{m+1}}, -a_J|\theta, M_i(\tau) = M_i, r_i(\tau) = m+1)} = \frac{\Pr(s_{i_{m+1}}, s_J|\theta)}{\Pr(-s_{i_{m+1}}, -s_J|\theta)}$$
$$= \left(\frac{0.75}{0.25}\right)^{\sum_{j \in J \cup \{i_{m+1}\}} \mathbf{1}(s_j = \theta) - \mathbf{1}(s_j = -\theta)}.$$

If i_{m+2} and i_m follow the actions of i_{m+3} and i_{m-1} respectively when they do not receive private signals, the informativeness of the consensus posts by i_{m+2} and i_m as they reach i_{m+1} is summarized by

$$\frac{\Pr(a_{i_{m+2}} = a_{i_m} = \theta|\theta, M_i(\tau) = M_i, r_i(\tau) = m+1)}{\Pr(a_{i_{m+2}} = a_{i_m} = -\theta|\theta, M_i(\tau) = M_i, r_i(\tau) = m+1)}$$
$$= \frac{(0.25)(0.5+\epsilon) + (0.5)(0.5+\epsilon)(0.8) + (0.25)(0.8)^2}{(0.25)(0.5-\epsilon) + (0.5)(0.5-\epsilon)(0.2) + (0.25)(0.2)^2} \approx 2.62 < \frac{0.75}{0.25}.$$

Thus, i_{m+1} optimally follows the majority of $\{s_{i_{m+1}}, a_{j_m^1}, ..., a_{j_m^p}\}$, ignoring the posts of i_{m+2} and i_m. If instead, i_{m+2} and i_m randomize equally between 1 and -1 when not receiving private signals, the informativeness of their consensus posts to i_{m+1} is summarized by

$$\frac{\Pr(a_{i_{m+2}} = a_{i_m} = \theta | \theta, M_i(\tau) = M_i, r_i(\tau) = m+1)}{\Pr(a_{i_{m+2}} = a_{i_m} = -\theta | \theta, M_i(\tau) = M_i, r_i(\tau) = m+1)}$$
$$= \frac{((0.5)(0.5) + (0.5)(0.8))^2}{((0.5)(0.5) + (0.5)(0.2))^2} = 3.45 > \frac{0.75}{0.25}.$$

In this alternative play, their consensus posts can overturn the majority of $\{s_{i_{m+1}}, a_{j^1_m}, ..., a_{j^p_m}\}$ when the winning margin is one. Thus, this change strictly benefits i_{m+1}.

When p is odd, modify the above construction by letting $\Pr(\tau_{i_{m+1}} = \infty) \approx 1$. In the individually optimal play, i_{m+1} follows the majority of $\{a_{j^1_m}, ..., a_{j^p_m}\}$. In the alternative play where i_{m+2} and i_m randomize equally between 1 and -1 when not receiving private signals, i_{m+1} strictly benefits from following these neighbors' consensus posts when the majority of $\{a_{j^1_m}, ..., a_{j^p_m}\}$ disagree with them but the majority is won by a margin of one.

Case 2: G' is a chordless cycle of size $L = 2m + 1$ for $m \geq 2$.

Notice that $(i_1, i_2, ..., i_m, i_{m+1})$ is the unique shortest path connecting i_1 and i_{m+1}. The reason is that if there was an alternative path of length less than or equal to m connecting i_1 and i_{m+1}, an incomplete cycle of size at most $2m$ would be formed by elements of these two paths. This would contradict that G' is the smallest incomplete cycle. Similarly, $(i_1, i_{2m+1}, ..., i_{m+3}, i_{m+2})$ is the unique shortest path connecting i_1 and i_{m+2}.

Construct a vector of signal qualities with $q_{i_1} = 0.9$, $q_{i_{m+3}} = 0.5 + \epsilon$ and $q_{i_{m+2}} = 0.7$. Construct a distribution H over the signal-timing vector such that $\Pr(\tau_{i_1} = 1) = \Pr(\tau_{i_1} = \infty) \approx 0.5$, $\Pr(\tau_{i_{m+3}} = m) \approx 1$, $\Pr(\tau_{i_{m+2}} = m+1) \approx 1$ and $\Pr(\tau_i = \infty) \approx 1$ for all $i \in N \backslash \{i_1, i_{m+3}, i_{m+2}\}$. Figure 8 plots two key diffusions of this example, with each occurring with probability approximately 0.5. In the left panel, i_1 receives his private signal in round 1 and has it spread to i_m and i_{m+3} in round m. Both of these individuals pass on this information to i_{m+1} and i_{m+2} respectively. In round $m+1$, individual i_{m+2} would consider between the action of i_{m+3} and her own private signal, but her decision will not be observed by i_{m+1}, who already takes an action in the same round. In the right panel, i_{m+3} receives his private signal in round m and passes that onto i_{m+2}, who will consider this piece of information with his own private signal to make a decision later observed by i_{m+1}.

Figure 8: The key diffusions in a generalization of Example 2

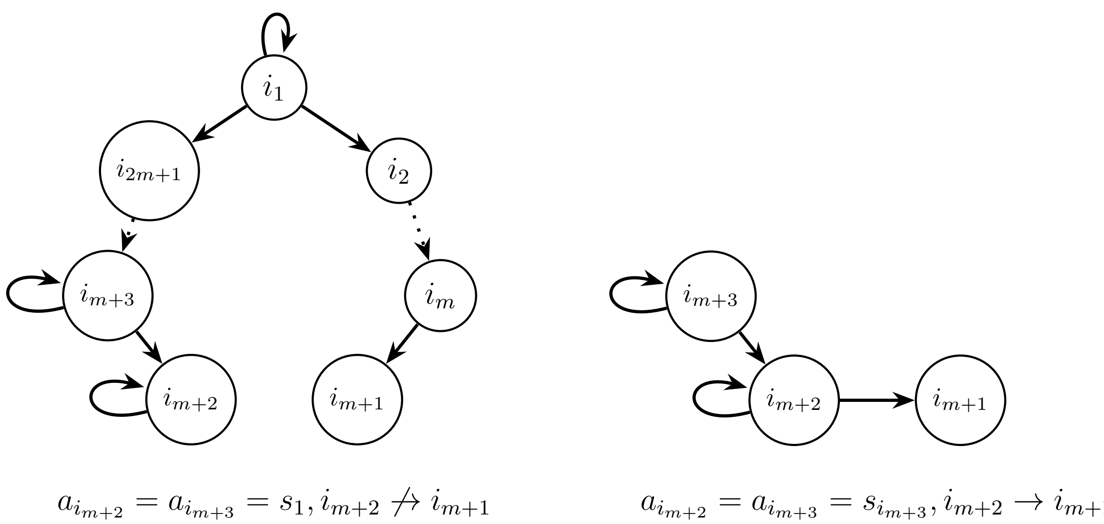

Under the constructed diffusion process,

$$\begin{aligned}&\Pr(a_{i_{m+3}} = \theta|\theta, M_{i_{m+2}}(\tau) = \{i_{m+3}, i_{m+2}\}, r_{i_{m+2}}(\tau) = m+1)\\ \approx&0.5(0.9) + 0.5(0.5+\epsilon) = 0.7 + 0.5\epsilon > q_{i_{m+2}}.\end{aligned}$$

It follows that at observation set $M_{i_{m+2}}(\tau) = \{i_{m+3}, i_{m+2}\}$ and action round $r_{i_{m+2}}$, individual i_{m+2} optimally chooses $a_{i_{m+2}} = a_{i_{m+3}}$. This gives

$$\Pr(a_{i_{m+2}} = \theta|\theta, M_{i_{m+1}}(\tau) = \{i_{m+2}\}, r_{i_{m+1}}(\tau) = m+2) = \Pr(s_{i_{m+3}} = \theta|\theta) = 0.5 + \epsilon.$$

If instead i_{m+2} chooses $a_{i_{m+2}} = s_{i_{m+2}}$, then

$$\Pr(a_{i_{m+2}} = \theta|\theta, M_{i_{m+1}}(\tau) = \{i_{m+2}\}, r_{i_{m+1}}(\tau) = m+2) = \Pr(s_{i_{m+2}} = \theta|\theta) = 0.7.$$

Such change strictly benefits i_{m+1} at a small expense of i_{m+2}.

Case 3: G' is a chordal graph.

Since G' is an incomplete chordal graph, it must be of size at least four and have a chord. Without loss of generality, suppose that i_1 and i_l are linked for some $l \in \{3, ..., L-1\}$. Then $\mathcal{C}(i_1, i_2, ..., i_{l-1}, i_l)$ and $\mathcal{C}(i_l, i_{l+1}, ..., i_L, i_1)$ are two cycles smaller than G'. These two cycles must be complete because otherwise it contradicts that G' is a smallest incomplete cycle. Since G' is incomplete, this means that there exist $j \in \{i_2, ..., i_{l-1}\}$ and $k \in \{i_{l+1}, ..., i_L\}$ that are not linked. In fact, $\mathcal{C}(i_1, j, i_l, k)$ is a smallest incomplete cycle of G. Consider the subgraph of G consisting of all shortest paths connecting j and k. Denote by $J = \{j_1, ..., j_p\}$ the set of individuals in this subgraph other than i_1, j, i_l and k. Construct a vector of signal qualities with $q_j = 0.5 + \epsilon$, $q_{i_1} = q_{i_l} = 0.8$ and $q_k = q_{j_{p'}} = 0.75$ for all $p' = 1, ..., p$. Construct the distribution of the signal-timing vector so that $\Pr(\tau_j = 1) \approx 1$, $\Pr(\tau_{i_1} = 2) = \Pr(\tau_{i_1} = \infty) = \Pr(\tau_{i_l} = 2) = \Pr(\tau_{i_l} = \infty) \approx 0.5$ and $\Pr(\tau_{j_{p'}} = 2) \approx 1$ for all $p' = 1, ..., p$. Furthermore, let $\Pr(\tau_k = 3) \approx 1$ if p is even and let $\Pr(\tau_k = \infty) \approx 1$ if p is odd. Similarly to the general example constructed in Case 1, this environment fails strong efficiency because under the individually optimal play, the correlation in the actions of i_1 and i_l makes k disregard their consensus actions. The alternative play where either i_1 or i_l randomizes equally between 1 and -1 when not receiving their private signals increases the informativeness of their consensus actions sufficiently to strictly benefit k.

A.10 Proof of Proposition 2

Suppose to the contrary that such $\bar{\psi}$ exists. Take r to be the smallest round such that for some $i \in N$, $\Pr(\bar{\psi}_i(\omega) \neq \bar{\psi}_i^*(\omega) | r_i(\tau) = r) > 0$. Since the local environment of i does not change at any conditioning set (M_i, r) for $M_i \subseteq N_i \cup \{i\}$, i can not do strictly better at these conditioning sets than she does under $\bar{\psi}^*$. That is,

$$\Pr(\bar{\psi}_i^*(\omega) = \theta | \theta, M_i(\tau) = M_i, r_i(\tau) = r) \geq \Pr(\bar{\psi}_i(\omega) = \theta | \theta, M_i(\tau) = M_i, r_i(\tau) = r)$$

for all $M_i \subseteq N_i \cup \{i\}$. The equality holds if and only if $\Pr(\bar{\psi}_i^*(\omega) = \bar{\psi}_i(\omega) | M_i(\tau) = M_i, r_i(\tau)) = 1$. The reason is that when i is indifferent between $a_i = 1$ and $a_i = -1$, the only symmetric response is to randomize between the two actions with equal probabilities. The inequality holding as an equality for all $M_i \subseteq N_i \cup \{i\}$ would contradict the choice of r as the smallest round where play differs between $\bar{\psi}^*$ and $\bar{\psi}$. The inequality holding strictly for some $M_i \subseteq N_i \cup \{i\}$ would mean that i is strictly worse off in round r under $\bar{\psi}$, a contradiction.

References

ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian learning in social networks," *The Review of Economic Studies*, 78, 1201–1236.

ALATAS, V., A. BANERJEE, A. G. CHANDRASEKHAR, R. HANNA, AND B. A. OLKEN (2016): "Network structure and the aggregation of information: Theory and evidence from Indonesia," *American Economic Review*, 106, 1663–1704.

BANERJEE, A. V. (1992): "A simple model of herd behavior," *The quarterly journal of economics*, 107, 797–817.

BATTIGALLI, P., A. FRANCETICH, G. LANZANI, AND M. MARINACCI (2019): "Learning and self-confirming long-run biases," *Journal of Economic Theory*, 183, 740–785.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, 100, 992–1026.

BOHREN, J. A. (2016): "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247.

BREZA, E., A. G. CHANDRASEKHAR, AND A. TAHBAZ-SALEHI (2018): "Seeing the forest for the trees? An investigation of network knowledge," Tech. rep., National Bureau of Economic Research.

BROWN, G. W. (1951): "Iterative solution of games by fictitious play," *Activity analysis of production and allocation*, 13, 374–376.

DASARATHA, K., B. GOLUB, AND N. HAK (2018): "Social learning in a dynamic environment," *arXiv preprint arXiv:1801.02042*.

EYSTER, E. AND M. RABIN (2010): "Naive herding in rich-information settings," *American economic journal: microeconomics*, 2, 221–43.

FUDENBERG, D. AND D. M. KREPS (1993): "Learning mixed equilibria," *Games and Economic Behavior*, 5, 320–367.

GUARINO, A. AND P. JEHIEL (2013): "Social learning with coarse inference," *American Economic Journal: Microeconomics*, 5, 147–74.

KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849–1892.

LI, W. AND X. TAN (2020): "Locally Bayesian learning in networks," *Theoretical Economics*, 15, 239–278.

LOBEL, I. AND E. SADLER (2015): "Information diffusion in networks through social learning," *Theoretical Economics*, 10, 807–851.

MCBRIDE, M. (2006): "Imperfect monitoring in communication networks," *Journal of Economic Theory*, 126, 97–119.

——— (2008): "Position-specific information in social networks: Are you connected?" *Mathematical Social Sciences*, 56, 283–295.

SETHI, R. AND M. YILDIZ (2016): "Communication with unknown perspectives," *Econometrica*, 84, 2029–2069.

——— (2019): "Culture and communication," *Available at SSRN 3263560*.

SMITH, L. AND P. SØRENSEN (2000): "Pathological outcomes of observational learning," *Econometrica*, 68, 371–398.

WISEMAN, T. (2009): "Sequential choice and non-bayesian observational learning," *International Game Theory Review*, 11, 285–300.