Echo Chambers: Social Learning under Unobserved Heterogeneity

Cole Randall Williams*

September 9, 2019

Abstract

People are often more amenable to opinions that agree with their own and even seek information from those with whom they expect to most agree—behaviors often attributed to irrational biases. In this paper, I argue that these behaviors can be understood within the context of rational social learning by accounting for the presence of unobserved heterogeneity in preferences or priors. Individuals display local learning by placing greater weight on opinions or behavior that are closer to their own. When individuals choose who to learn from, local learning leads to the development of echo chambers.

People rely on social learning to navigate the world: Consumers read product reviews, firms gauge a technology's quality by its level of adoption, and citizens consult each other when forming their political opinions. Throughout the course of learning, people routinely assess opinions that agree with their own to be more reliable (Nickerson, 1998) and actively

*Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, cole.williams@univie.ac.at. I am grateful to Jean-Paul Carvalho and Stergios Skaperdas for guidance and support. This work benefited from comments by John Duffy, Ami Glazer, Maarten Janssen, Daniel Garcia, Eeva Mauring, Omer Tamuz, Laura Doval, Juuso Välimäki and seminar participants at UC Irvine, University of Vienna, Caltech, Chapman, Bristol, Cambridge, Higher School of Economics, and the 30th International Conference on Game Theory at Stony Brook. I acknowledge financial support from the Kassouf Foundation and the Austrian Science Fund FWF under project number I 3487.

seek information from those with whom they tend to agree (Del Vicario et al., 2016). Behaviors like these appear to conflict with classic models of rational learning and have been attributed to the presence of biases in information processing. Examples include motivated reasoning (Kunda, 1990), persuasion bias (Jasny, Waggle and Fisher, 2015), and a desire for certainty (Boutyline and Willer, 2017).

In this paper, I show that these seemingly anomalous patterns can be understood within the context of rational learning. Specifically, in a population with unobserved heterogeneity in preferences or prior beliefs, individuals engaged in social learning rationally exhibit a type of confirmation bias and naturally form into echo chambers. Social learning in a context of unobserved heterogeneity becomes a process of dual learning: by observing the behavior of others an individual learns both about his parameter of interest and the structure of the heterogeneity in the population. Encountering someone with divergent behavior could now mean that this individual differs from himself in important ways. As such, learning is local: individuals place greater weight on behavior closer to their own. When choosing who to learn from, individuals select those with whom they expect to find the most agreement.

There are many real-world examples of dual learning processes. The following presents four such examples.

Restaurant. Restaurant choice is a canonical example of social learning.[1] Under unobserved heterogeneity, the more positive experiences one has of a restaurant, the more likely one is to discount a negative reviewer's opinion as proceeding from different preferences.

Politics. Individuals may form their political beliefs by sharing opinions, but may also retain distinct preferences over policies due to different normative values and/or interpretation of evidence. The larger the divergence in opinions, the more likely one is to attribute disagreement to intrinsic differences.

Technological Innovation. Observed patterns in innovation diffusion are

[1] Examples include Becker (1991), Banerjee (1992), Kirman (1993), Smith and Sorensen (2000), Ellison and Fudenberg (1995),Chamley (2004), and Eyster and Rabin (2014).

often explained as the result of heterogeneity among an innovation's potential adopters.[2] An example is Munshi (2004) who studies the diffusion of high yield varieties of rice and wheat during the Indian Green Revolution. Rice yields were particularly sensitive to variations in factors like soil characteristics and managerial inputs that are not easy to observe. Munshi finds evidence that growers came to place less weight on their neighbors' rice-growing decisions and outcomes than they do in the case of wheat.

Scientific Theories. Experts equally fluent in a scientific discipline often disagree.[3] One possible source of disagreement is the diversity in inferences drawn from evidence. Bayesian econometricians focus less on statistically significant p-values, and people may be convinced to different degrees of an instrumental variable's excludability or a theoretical model's assumptions. The different lenses through which we filter empirical observations, including scientific research, can lead to a diversity of opinion. Hence, experts may attribute disagreement to different dispositions to evidence.

I explore how social learning influences behavior and how individuals seek information from others in a simple observational learning environment. A population of individuals each face a sequence of actions which yield uncertain payoffs. Individuals take each action on the basis of two sources of information. First, each receives a private signal. Second, each selects one other individual from the population and observes his action. A challenge to learning from others is *unobserved heterogeneity*: individuals are uncertain about the *similarity* (correlation) between their own preferences or priors and that of the others in the population.

Social learning becomes dual learning by inducing belief revision along two dimensions. Firstly, holding fixed an individual's beliefs over the structure of heterogeneity in the population, social learning influences an indi-

[2]See Jensen (1982), Mahajan and Peterson (1985), Jeuland (1987), and Young (2009).

[3]Galileo battled with the Catholic church and fellow scientists alike over the heliocentric model of the solar system, the germ theory of disease was contested for centuries, and there was longstanding dissent over theories of continental drift. Contemporary science hosts disagreements over the fundamental roles of randomness and measurement in quantum mechanics (Schlosshauer, Kofler and Zeilinger, 2013) and the plausibility of group selection in evolutionary biology (et al. Abbot, 2011). There is evidence that ideological leanings influence the research conducted by economists (Jelveh, Kogut and Naidu, 2018).

vidual's beliefs in qualitatively the same way as in a homogeneous environment. Along this dimension, larger disagreement tends to yield a larger shift in the individual's beliefs and thus actions. Secondly, social learning induces an individual to update his beliefs over the structure of heterogeneity in the population. This dimension determines how much *weight* an individual places on social information when choosing an action.

Our first result establishes *local learning*: individuals tend to place greater weight on social information that is in closer agreement with their private beliefs. Thus, in the face of disagreement, the two dimensions of dual learning conflict with each other. This can lead to non-monotonicities in belief formation, whereby, moderate disagreement can produce a greater shift in behavior than larger disagreement. Moreover, as an individual continues learning from someone, the degree to which this other individual influences his actions becomes decreasing in their disagreement.

Local learning further drives the evolution of an individual's choice of who to observe. In the baseline model with binary-valued preference types, individuals observe those with whom they expect to find the least disagreement. This gives rise to a characteristic feature of *echo chambers*: the disagreement between those who share information will tend to be less than the disagreement in the general population and is also supermartingale, diminishing in expectation over time.

The observational network dynamics becomes more complicated when we consider richer forms of heterogeneity. However, if we specify that individuals occasionally see the past actions of members of their society, have imperfect control over who they learn from, or incur a cost to observe others, then the limit of the process is an echo chamber. The basic idea is simple. Under each of these extensions, an individual occasionally observes the actions of someone who is not necessarily his first choice to learn from. If he comes to find that this individual has exhibited a sufficiently strong record of agreement, then he will eventually choose to learn from this individual more frequently.

Previous work in social learning largely focuses on characterizing the types of observational networks that allow a society to effectively learn the optimal action; for example, whether the members of society learn to adopt a new technology in the event that it is superior to its predecessor.[4] They find that preference heterogeneity can impede learning (Bala and Goyal, 1995; Smith and Sorensen, 2000; Lobel and Sadler, 2016), or if the heterogeneity is rich enough to prohibit an information cascade, enhance learning (Goeree, Palfrey and Rogers, 2006). In these models, the specific form of heterogeneity is common knowledge. I find that unobserved heterogeneity further constraints what can be learned from others via local learning: information that strays too far from one's beliefs will be largely discounted.

Closer to this paper's formulation is Sethi and Yildiz (2016) who also study a model in which individuals sequentially select someone to learn from, but they consider a distinct formulation of heterogeneity. In their model, individuals have fixed, heterogeneous beliefs (perspectives) about the distribution from which a sequence of states are drawn. They find that when outcomes are history-dependent, individuals eventually only choose to learn from a small subset of individuals whose perspectives they have come to best understand. This stands in contrast with my results where individuals choose to learn from those with whom they tend to most agree.

Che and Mierendorff (2019) study a related problem of choosing a media source with publicly known bias. Individuals are faced with a stopping problem, and can allocate their attention between a media source that is either strongly biased in favor or strongly biased against their prior beliefs. In some cases, it is optimal for those with strong priors to attend to own-biased media while those with weaker priors attend to opposite-biased media.

Jann and Schottmüller (2018) analyze a setting where individuals desire for both their own and others' actions to closely match the uncertain state plus their commonly known idiosyncratic bias. One important difference is that I focus on the endogenous formation of echo chambers while they

[4] Surveys include Chamley (2004) and Möbius and Rosenblat (2014). Early work identifies herding in a sequential observational network (Bikhchandani, Hirshleifer and Welch, 1992; Banerjee, 1992) and that a herd can form on a suboptimal action if and only if there is an upper bound on the strength of private information (Smith and Sorensen, 2000). More recent work considers alternative network structures (Gale and Kariv, 2003; Banerjee and Fudenberg, 2004; Rosenberg, Solan and Vieille, 2009; Acemoglu et al., 2011; Mossel, Sly and Tamuz, 2014) and costly information (Mueller-Frank and Pai, 2016; Ali, 2018).

identify cases in which welfare maximizing networks constitute an equilibrium. They find that segmenting the population into groups with shared biases can engender honest communication, thereby improving efficiency. Interestingly, they find cases in which fully segmenting the population is welfare maximizing, but does not constitute an equilibrium.

My results also connect with the literature on homophily in networks; wherein, a preference for interacting with those who share similar characteristics with oneself is shown to produce segregated communities (Schelling, 1969; Currarini, Jackson and Pin, 2009) and stifle the dissemination of information (Golub and Jackson, 2008; Halberstam and Knight, 2016). In light of my results, individuals may also form their social networks to include those who share similar opinions or take similar actions out of a desire for more useful information, even without an explicit preference for homophily.

The remainder of the paper is organized as follows. Section 1 presents the baseline model of dual learning in which preference types are binary-valued. Section 2 analyzes the model, characterizing how individuals act after observing others (2.1), how social influence relates to disagreement (2.2 and 2.3), and the emergence of echo chambers (2.4). Section 3 generalizes the results to settings with richer forms of heterogeneity and Section 4 concludes.

1 Baseline Model

There is a population of individuals $i \in \{1, 2, ..., n\}$ and a sequence of periods $t \in \{1, 2, \dots\}$. In period t, individual i seeks to learn $\theta_{it} \in \{a, b\}$. For example, each period the population is presented with two proposed political policies or two competing candidates and θ_{it} is the most beneficial option for individual i. Section 3 extends this to consider a larger class of preference types. At the beginning of each period, $\boldsymbol{\theta}_t = (\theta_{1t}, \theta_{2t}, ..., \theta_{nt})$ is drawn from a distribution $f(\boldsymbol{\theta}_t|\boldsymbol{\tau})$ that depends on a $n \times n$ correlation matrix $\boldsymbol{\tau}$ where element $\tau_{ij} = \tau_{ji} = corr(\theta_{it}, \theta_{jt}) \in [0, 1]$ is called the *similarity* of i and j. Assume the marginal distributions are the same across individuals $f(\theta_{it}) = f(\theta_{jt})$. Individuals do not directly observe $\boldsymbol{\tau}$, but it is commonly known to be drawn from a nondegenerate distribution $Q(\boldsymbol{\tau})$.

Upon entering period t, individual i receives a private signal $s_{it} \in S$ according to probability measure $\mu_{\theta_{it}}$. So that no signal perfectly reveals an individual's parameter and to guarantee that some signals are informative, assume μ_a and μ_b to be mutually absolutely continuous and not almost surely identical. With his signal in hand, individual i takes his first action of the period $x_{it} \in \mathbb{R}$ with quadratic loss payoffs

$$u(x_{it}; \theta_{it}) = -(x_{it} - \theta_{it})^2. \tag{1}$$

At this stage, i picks out one other individual $j \neq i$ and observes his first action for that period x_{jt}. Since their preferences may be correlated, observing x_{jt} allows i to revise his beliefs over θ_{it}. At the same time, observing how j acts allows i to update his *perceived similarity* of j which we simply denote by $Q(\tau_{ij}|\cdot)$. Finally, i takes his second action for the period $y_{it} \in \mathbb{R}$ which also yields quadratic loss payoffs

$$u(y_{it}; \theta_{it}) = -(y_{it} - \theta_{it})^2. \tag{2}$$

We assume that individuals do not observe the utility from their first action before taking their second action. In some cases, it is reasonable to assume a delay between the moment an action is taken and when its consequences are realized, as it would be with realizing the full effects of introducing a political policy. Following Sethi and Yildiz (2016), we could also think of x_{it} as i's *opinion*; it is the action i would take on the basis of only his private information and y_{it} is the action he ends up taking after consulting the opinion of j. Furthermore, individuals do not see who others chose to observe or any other actions before taking their second action so that the difference between y_{it} and x_{it} solely reflects what is learned from observing j.

After taking his action y_{it} the period concludes. At this point, individuals may observe their realized payoffs, see who others chose to observe, and how they acted in that or previous periods. For the moment, we remain agnostic about what individuals see. Later on, in Section 3, I show that if individuals are granted rich enough information, we obtain sharp echo chamber results under more general forms of heterogeneity.

As is common in models of social learning, individuals are assumed to be myopic, seeking to maximize the expected utility for the present period when taking their actions and choosing whom to observe. This enhances tractability, but more importantly, allows for the analysis to isolate the precise effects of unobserved heterogeneity.

The first objective of this paper is to characterize the degree to which an individual's actions are affected by observing others. Since x_{it} is how i acts with only private information and y_{it} is how he acts after observing j, we define j's *social influence* on i to be the absolute difference $|y_{it} - x_{it}|$. For example, this captures how a conversation with j shifts i's views on a political policy. A concept that proves useful for this characterization is that of *disagreement*, defined as the absolute difference between i and j's actions $d_{ij}(t) = |x_{jt} - x_{it}|$. Disagreement is also useful for the second objective of analyzing the communication networks that emerge from individuals choosing who to learn from.

2 Analysis

2.1 Local Learning

As a benchmark for comparison, note that there is a fundamental monotonicity for learning in a homogeneous society. If all citizens share the same preferences, then when an individual discovers that someone has an unexpectedly high opinion of a candidate, it must cause his own opinion of the candidate to improve.[5] A byproduct of this monotonic inference is that disagreement and social influence will tend to be positively related.

In this section, we shall see that the logic of this monotonic inference breaks down in the presence of unobserved heterogeneity as a result of dual

[5] In the language of Milgrom (1981), observing a higher action from j is "good news" for i. A higher action $x'_{jt} \geq x_{jt}$ yields a higher expectation $y_{it}(s_{it}, x'_{jt}) \geq y_{it}(s_{it}, x_{jt})$.

learning. When an individual discovers that someone has a much higher opinion of a candidate than himself, it is both evidence of the desirability of the candidate, but also that this person could have dissimilar preferences. Let us begin by identifying these two dimensions of inference.

Upon entering period t, individual i receives a private signal s_{it} and takes his first action x_{it} of the period. Since this choice has no bearing on the quality of information he receives later in the period, i's objective is simply to maximize

$$\max_{x_{it}} \mathbb{E}\left[u(x_{it};\theta_{it})|s_{it}\right] \tag{3}$$

which is achieved at $x_{it} = \mathbb{E}\left[\theta_{it}|s_{it}\right]$. Letting h_{it} denote the history of i's observations up to time t, i chooses an individual j, observes his action x_{jt}, and selects his second action y_{it}, both to maximize

$$\max_{j} \mathbb{E}\left[\max_{y_{it}} \mathbb{E}\left[u(y_{it};\theta_{it})|s_{it},x_{jt},h_{it}\right] \middle| h_{it}\right]. \tag{4}$$

Since preferences may be correlated, observing how j acts allows i to update his beliefs over θ_{it} and also the perceived similarity $Q(\tau_{ij}|s_{it},x_{jt},h_{it})$. The second action i takes is now the revised expectation of his parameter. In Lemma A.1 of the appendix, I show that the joint distribution for $(\theta_{it},\theta_{jt})$ and a given similarity τ_{ij} can be represented as

$$f(\theta_{it},\theta_{jt}|\tau_{ij}) = \tau_{ij} f(\theta_{it})\mathbf{1}_{\theta_{it}=\theta_{jt}}(\theta_{it},\theta_{jt}) + (1-\tau_{ij})f(\theta_{it})f(\theta_{jt}). \tag{5}$$

Intuitively, with probability τ_{ij} it is guaranteed that $\theta_{it} = \theta_{jt}$, and with the complementary probability θ_{it} and θ_{jt} are independent. This intuition carries through when computing the action i takes after observing j, which can be written

$$y_{it} = \alpha\hat{y}_{it} + (1-\alpha)x_{it} \tag{6}$$

where $\hat{y}_{it} = \mathbb{E}\left[\theta_{it}|s_{it},x_{jt},\tau_{ij}=1\right]$ and $\alpha = \alpha(s_{it},x_{jt},h_{it}) \in [0,1]$ is a positive affine transformation of the posterior expected similarity $\mathbb{E}\left[\tau_{ij}|s_{it},x_{jt},h_{it}\right]$.[6]

[6]Letting $\bar{\tau}$ be the prior and $\bar{\tau}^*$ the posterior expected similarity, $\alpha = -\frac{\bar{\tau}\left(\bar{\tau}-\mathbb{E}[\tau^2]\right)}{\mathbb{E}[\tau^2]-\bar{\tau}^2} +$

Thus i's action is a weighted average between how he would act under homogeneity $\hat{y}_{it}$ and how he would act with only his private information x_{it}, with the weight determined by the posterior perceived similarity.

Notice that j's action enters twice in the expectation (6), capturing the two dimensions of dual learning. First, if we fix α, i's response to observing j is qualitatively the same as if they were homogeneous, $\hat{y}_{it}$ is monotonically increasing in x_{jt}. Second, x_{jt} enters α, and thus the weight i places on j's action is itself determined by the action.

Our first result establishes that i adjusts the weight placed on j's action in accordance to their *disagreement* in a way that resembles confirmation bias (Nickerson, 1998). Specifically, i will (on average) assign more weight to j when he observes less disagreement.

Proposition 1. *(Local Learning) On average, the weight placed on j's action is larger when there is less disagreement. That is, for any δ in the support of disagreement*

$$\mathbb{E}\left[\alpha | d_{ij}(t) \leq \delta\right] \geq \mathbb{E}\left[\alpha | d_{ij}(t) \geq \delta\right].$$

Proof. See appendix. □

In other words, if we compare two action pairs (x_{it}, x_{jt}) and (x'_{it}, x'_{jt}) whereby $|x_{jt} - x_{it}| < \delta < |x'_{jt} - x'_{it}|$, on average, the weight placed on x_{jt} will exceed the weight placed on x'_{jt}. For example, suppose signals are binary-valued $s_{it} \in \{a, b\}$ with precision $\Pr[s_{it} = \theta_{it}] = 3/4$ and uniform priors over preference types and similarity. It will either be that individuals "agree" $d_{ij}(t) = 0$, in which case $\alpha = 5/9$, or they will "disagree" $d_{ij}(t) = 1/2$ and $\alpha = 3/7$.

2.2 Non-Monotonicity in Disagreement

We have seen that disagreement creates a tension between the two dimensions of learning. In this section, we shall see that when an individual learns

$\frac{\bar{\tau}(1-\bar{\tau})}{\mathbb{E}[\tau^2]-\bar{\tau}^2}\bar{\tau}^*$.

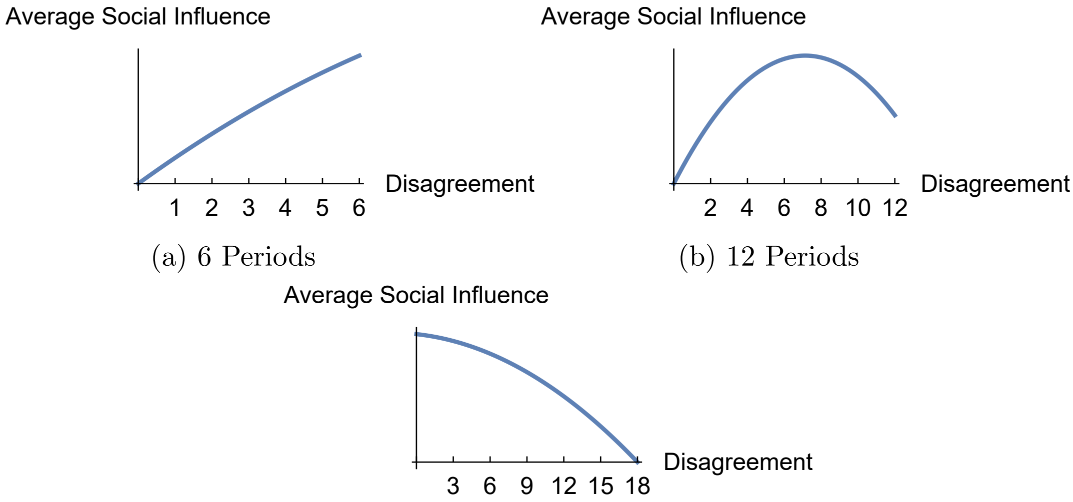

(a) 6 Periods

(b) 12 Periods

(c) 18 Periods

Figure 1: Non-Monotonicity in Disagreement. These figures plot the average social influence j has on i as a function of the number of their disagreements D when $r = 2/3$, $Q(\tau_{ji}) \sim U(0,1)$, and when i observes j in (a) $\hat{t} = 6$, (b) $\hat{t} = 12$, and (c) $\hat{t} = 18$ periods.

from someone over the course of several periods, this tension can produce a non-monotonic relationship between disagreement and social influence.

We shall characterize how an individual decides who to learn from later on in Section 2.4. For now, let T be a set of $\hat{t}$ periods wherein i observes j and let $(\mathbf{x}_i, \mathbf{y}_i)$ correspond to i's actions in these periods. Define $\hat{t}^{-\frac{1}{2}}\|\mathbf{x}_j - \mathbf{x}_i\|$ to be the *average disagreement* and $\hat{t}^{-\frac{1}{2}}\|\mathbf{y}_i - \mathbf{x}_i\|$ the *average social influence* where $\|\cdot\|$ is the Euclidean norm.

We illustrate the non-monotonicities that can emerge with the following example. Suppose there are two individuals $n = 2$, signals are binary-valued $s_{it} \in \{a, b\}$ with $\Pr[s_{it} = \theta_{it}] = r > \frac{1}{2}$, and the prior is uniform $\Pr[\theta_{it} = a] = \frac{1}{2}$. In this case, individuals "agree" ($d_{ij}(t) = 0$) if they receive the same signal and "disagree" ($d_{ij}(t) = 2r - 1$) if their signals are different. Holding fixed the perceived similarity, disagreement yields a higher social influence than agreement and local learning (Proposition 1) guarantees that agreement produces a posterior perceived similarity with first order stochastic dominance over that of disagreement $Q(\tau_{ij}|\text{agree}) \succeq_{FOSD} Q(\tau_{ij}|\text{disagree})$.

If i only observes j in a single period, then so long as the perceived similarity is not degenerate at $\tau_{ij} = 0$, the first dimension dominates and disagreement yields a higher social influence than agreement. Suppose that i observes j in $\hat{t}$ periods, of which D periods they disagree and the remaining $\hat{t} - D$ periods they agree. Figure 1 considers the relationship between (average) social influence and disagreement when $\hat{t} \in \{6, 12, 18\}$, $r = 2/3$, and $Q(\tau_{ij}) \sim U(0, 1)$. The horizontal axis increases the number of periods with disagreement D from 0 to $\hat{t}$. Since social influence depends on the order of observations, Figure 1 plots social influence averaged over all permutations of signals that give D disagreements in $\hat{t}$ observations.[7]

Figure 1a illustrates that, when there are only a few observations, larger disagreement yields more social influence. However, as i continues observing j (Figures 1b and 1c), the effect of local learning becomes increasingly strong so that mild or moderate disagreement leads to greater social influence than does larger disagreement. Local learning constrains how much can be learned from others. This is most evident in Figure 1c where an individual learns the most when his opinions are echoed back to him. As a result, being misinformed is doubly costly in a heterogeneous society—the very social information that is most useful in correcting an individual's mistaken beliefs will be most heavily discounted.

2.3 Social Influence Declines with Disagreement

We now generalize the insight that disagreement can diminish influence as the number of observations grow. For a given period, it is straightforward to show that i's action y_{it} is strictly increasing in the action he observes in that period x_{jt}. As a result, influence will typically be increasing in disagreement after the first period. As the periods progress, the relationship between influence and disagreement will depend on the specific prior $Q(\boldsymbol{\tau})$ and the resulting observational dynamics. Indeed, a simple extension of Proposition 1 finds that, holding everything else fixed, larger disagreement in one period reduces the weight placed on another's action in all subsequent periods. As the number of observations grows large, the effect of local learning becomes more pronounced and the following pattern emerges.

[7] For instance, if i and j only agree once in six periods, social influence will be higher if they agree in the first as opposed to the sixth period.

Proposition 2. *Suppose i observes j infinitely often. Then j's average social influence on i is almost surely decreasing in their average disagreement.*

To put this another way, consider two pairs of individuals, where i observes j and k observes ℓ in at least $\hat{t}$ periods. If i and j enjoy less disagreement than k and ℓ, then the chance that j exerts greater influence on i than ℓ does on k can be made arbitrarily close to one by making the number of observations $\hat{t}$ sufficiently large.

2.4 Echo Chambers

In this section, I show that social learning gives rise to echo chambers in the sense that individuals seek information from those with whom they expect to find the most agreement. As a consequence, the disagreement between those who share information will be less than that of those who do not share information, with the difference growing over time.

The expected payoff for i to observe j depends on history by way of the perceived similarity $Q(\tau_{ij}|h_{it})$; which, by Proposition 1 will tend to reflect how closely i has agreed with j in previous periods. In general, $Q(\tau_{ij}|h_{it})$ also depends on i's observations of individuals other than j since similarity is typically not independent across pairs of individuals.

Using the linear representation of the joint distribution (5), from i's vantage point at history h_{it}, the conditional distribution of j's parameter given his own only depends on their *expected similarity* $\bar{\tau}_{ij}(h_{it}) = \mathbb{E}\left[\tau_{ij}|h_{it}\right]$

$$f(\theta_{jt}|\theta_{it}, h_{it}) = \int f(\theta_{jt}|\theta_{it}, \tau_{ij}) \mathrm{d}Q(\tau_{ij}|h_{it}) \tag{7}$$

$$= \bar{\tau}_{ij}(h_{it}) \mathbf{1}_{\theta_{jt}=\theta_{it}}(\theta_{jt}) + (1 - \bar{\tau}_{ij}(h_{it}))\, f(\theta_{jt}). \tag{8}$$

Using this expression and letting $G(x_{jt}|\cdot)$ represent the CDF of j's action, we can write

$$G(x_{jt}|\theta_{it}, h_{it}) = \bar{\tau}_{ij}(h_{it}) G(x_{jt}|\theta_{it}) + (1 - \bar{\tau}_{ij}(h_{it}))\, G(x_{jt}). \tag{9}$$

If j has a higher expected similarity than k, then observing k is equivalent to observing j with probability $\frac{\bar{\tau}_{ik}}{\bar{\tau}_{ij}}(h_{it})$ and observing noise with the complementary probability. Hence, j's action is Blackwell more informative than k's action so that i always does best by observing the individual with the highest expected similarity (Blackwell, 1951). From local learning (Proposition 1), we know that this is synonymous with choosing the individual with the least expected disagreement.

Proposition 3. *Individuals observe the actions of those with whom they expect the least disagreement.*

This simple decision rule culminates in clear patterns for societal disagreement. Let j_{it} denote the individual i observes in period t and define *in-group* disagreement to be the average disagreement observed by an individual $d_t = \frac{1}{n}\sum_{i,j} d_{ij}(t)\mathbf{1}_{j=j_{it}}$. As its complement, let *out-group* disagreement denote the average disagreement that an individual does not observe $d_t' = \frac{1}{n(n-2)}\sum_{i,j} d_{ij}(t)\mathbf{1}_{j\neq i,j_{it}}$.

An echo chamber is characterized by higher rates of agreement between those linked within an observational network than between those not linked (Jasny, Waggle and Fisher, 2015). Due to the constraints placed on individuals' information and the random variability of realized signals, some periods may find in-group disagreement to exceed out-group disagreement. However, taking the average over all the possible trajectories for d_t and d_t', we obtain the following.

Proposition 4. *Expected in-group disagreement does not exceed expected out-group disagreement for any period. That is,*

$$\mathbb{E}\left[d_t\right] \leq \mathbb{E}\left[d_t'\right] \tag{10}$$

with a strict inequality in any period t with a positive probability that some individual i strictly prefers his choice j_{it} to some deviation $j \neq i, j_{it}$.

Consider how the disparity between in-group and out-group disagreement evolves over time. Notice that if observational networks are fixed, the difference between these two also remains fixed. Since individuals seek minimal disagreement, the expected disagreement encountered by an individual $\mathbb{E}\left[d_{ij_t(i)}|h_{it}\right]$ becomes a supermartingale process with respect to his history,

$$\mathbb{E}\left[\mathbb{E}\left[d_{ij_{t+1}(i)}(t+1)|h_{i,t+1}\right]|h_{it}\right] \leq \mathbb{E}\left[\mathbb{E}\left[d_{ij_t(i)}(t+1)|h_{i,t+1}\right]|h_{it}\right]$$
$$= \mathbb{E}\left[d_{ij_t(i)}|h_{it}\right]$$

where the inequality is strict at any history h_{it+1} where i strictly prefers his choice j_{it+1} to j_{it}. Taking the expectation over all histories, we can characterize the dynamics of disagreement.

Proposition 5. *Expected in-group disagreement is decreasing over time. That is,*

$$\mathbb{E}\left[d_{t+1}\right] \leq \mathbb{E}\left[d_t\right] \tag{11}$$

with a strict inequality in any period $t+1$ *with a positive probability that some individual* i *strictly prefers his choice* j_{it+1} *to* j_{it}.

Thus we find that when individuals are free to choose whom they learn from, the presence of heterogeneity inevitably leads to the development of echo chambers. In the present context, with binary-valued preference types, the result is particularly crisp—each period we should expect less disagreement within communication networks than across networks, with the difference growing over time. In the next section, I show how the essence of this result is much more general.

3 Generalizing

The main findings of the paper readily extend to settings with richer forms of heterogeneity. For example, an individual may seek to learn which income tax rate would offer him the greatest benefit if it were implemented, in which case, his preference type could take on a continuum of values. Individuals may instead hold identical preferences over a policy, but differ in their priors or subjective probability distributions.

3.1 Variations

We begin by introducing four variants on the baseline model. The first two extend preference heterogeneity to allow for more than two preference types (in these $\theta_{it} \in \mathbb{R}$). The second two suppose that individuals share identical preferences but differ on epistemic grounds.

I. Private Values. Private values and signals are both normally distributed $\boldsymbol{\theta}_t \sim \mathcal{N}(\boldsymbol{\theta}_0, \sigma_0^2 \boldsymbol{\tau})$ and $s_{it} \sim \mathcal{N}(\theta_{it}, \sigma^2)$.

II. Common and Private Values. As in Goeree, Palfrey and Rogers (2006), Hendricks, Sorensen and Wiseman (2012), and Lobel and Sadler (2016), preferences have two components $\theta_{it} = \theta_t + \eta_{it}$ where θ_t is unknown and common to everyone and η_{it} is specific to i and only known to him. Suppose both components are normally distributed $\theta_t \sim \mathcal{N}(\theta_0, \sigma_0^2)$ and $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\tau})$ and signals are iid conditional on the common component $s_{it} \sim \mathcal{N}(\theta_t, \sigma^2)$.

III. Subjective Probability Distributions. Everyone shares the same preferences $\theta_t \sim \mathcal{N}(\theta_0, \sigma_0^2)$. However, they disagree about how to interpret signals with individual i in period t believing that all signals are distributed $\mathcal{N}(\theta_t + \eta_{it}, \sigma^2)$ where η_{it} is i's privately known bias and $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\tau})$.

IV. Priors. Following Sethi and Yildiz (2012, 2016), individuals differ in their priors, with i's prior in period t being given by $\theta_t \sim \mathcal{N}(\eta_{it}, \sigma_0^2)$ and everyone agrees $s_{it} \sim \mathcal{N}(\theta_t, \sigma^2)$. However, rather than an individual's prior corresponding to a fixed "perspective" as in Sethi and Yildiz (2012, 2016), new priors are drawn each period $\boldsymbol{\eta}_t \sim \mathcal{N}(\boldsymbol{\eta}_0, \boldsymbol{\tau})$.

Each of these variations and also the baseline model have two related properties that drive the results. First, the distribution of actions satisfies the following condition

$$\Pr\left[x_{it} < \alpha, x_{jt} < \beta | \tau_{ij}\right] \leq \Pr\left[x_{it} < \alpha, x_{jt} < \beta | \tau'_{ij}\right] \tag{12}$$

when $\tau_{ij} \leq \tau'_{ij}$. Essentially, the chance that i and j's actions are simultaneously low (or high) is increasing in their similarity. Whenever the joint distribution of two individuals' actions satisfies the monotone likelihood ratio property,[8] condition (12) is equivalent to an outside observer finding x_{jt} to be more Lehmann informative of x_{it} when i and j are more similar (Lemma A.2). This condition turns out to be sufficient for individuals to gauge similarity on the basis of agreement.

Proposition 6. *When condition (12) is satsified, then for any δ in the support of disagreement*

$$\mathbb{E}\left[Q(\tau_{ij}|h_{it+1})\middle|d_{ij}(t) \leq \delta\right] \leq \mathbb{E}\left[Q(\tau_{ij}|h_{it+1})\middle|d_{ij}(t) \geq \delta\right]. \qquad (13)$$

In each of these variants of the model, the action i takes after observing j can be expressed as

$$y_{it} = \gamma x_{jt} + (1-\gamma) z_{it} \qquad (14)$$

where $\gamma = \int \gamma_\tau(\tau_{ij}) \mathrm{d}Q(\tau_{ij}|s_{it}, x_{jt}, h_{it})$ captures the weight placed on j's action with $\gamma_\tau(\tau_{ij})$ strictly increasing and $z_{it} = z_{it}(s_{it}, \eta_{it})$ a function of i's private information. This representation follows simply from iterated expectations $y_{it} = \mathbb{E}\left[\mathbb{E}\left[\theta_{it}|s_{it}, x_{jt}, \tau_{ij}\right]|s_{it}, x_{jt}, h_{it}\right]$ and noting that the inner expectation takes the familiar additive form since it is a posterior expectation in the normal-normal conjugate family. As a corollary to Proposition 6, we can extend local learning.

Corollary 1. *In variations I.-IV., on average, the weight placed on j's action is larger when there is less disagreement. That is, for $\delta > 0$*

$$\mathbb{E}\left[\gamma|d_{ij}(t) \leq \delta\right] > \mathbb{E}\left[\gamma|d_{ij}(t) > \delta\right].$$

Local learning suggests that one reason why individuals attend more to opinions and behaviors that are more in line with their own is simply that they are deemed more informative. In the baseline model and variation I., learning from someone who expresses closer agreement is likely to be directly more informative for how one should act. In variations II.-IV., closer agreement suggests one can better utilize knowledge about himself (η_{it}) to extract the private information contained in someone's action.

[8] The distribution of actions satisfies the *monotone likelihood ratio property* if $x_{jt} \leq x'_{jt}$ implies $\frac{g(x_{it}|x'_{jt}, \tau_{ij})}{g(x_{it}|x_{jt}, \tau_{ij})}$ is non-decreasing in x_{it} for all τ_{ij}.

By nearly the same intuition, in the event that similarity is known, an individual would find observing someone with higher similarity to be more Lehmann informative for how he should act

$$x_{jt}|\tau'_{ij} \succeq_i x_{jt}|\tau_{ij} \text{ when } \tau_{ij} \leq \tau'_{ij}. \tag{15}$$

In fact, from our previous work and Example 1 in Lehmann (1988), higher similarity increases the Blackwell informativeness of j's action in each version of model (see Lemma B.2), though this is stronger than we need for our results.

Proposition 7. *When conditions (12) and (15) are satisfied, if i observes j infinitely often, then j's average social influence on i is almost surely decreasing in their average disagreement.*

3.2 Extensions

In general, an individual's choice of whom to observe in a given period is not as straightforward as in the baseline model. The expected payoff to observing someone now depends on the particular shape of the perceived similarity, so that it is no longer always optimal to observe the individual with the highest expected similarity.

Nevertheless, a number of natural extensions on the model guarantee that the process will form echo chambers. For the remainder, we assume that, just as in variations I.-IV., the prior and distribution of signals are sufficiently well-behaved so that the variance over θ_{it}, given i's private information, is bounded from above by some value v.

(i) Additional Observations. Up until now, we have not specified what information is revealed at the end of a period. Suppose that after a period has concluded, individuals occasionally learn how others in the population have acted. Specifically, suppose that after period t, i sees the actions taken in that period by $j \in N_{it}$. The sequence $(N_{it})_t$ may be deterministic (e.g. $N_{it} = \emptyset$ when t is prime and $N_{it} = \{j\}$ otherwise) or stochastic (e.g. each period N_{it} contains two individuals, chosen randomly and with uniform probability). Assume N_{it} to not be empty infinitely often and $\Pr[j \in N_{it}]$

to be well-defined and independent of the realized similarity so that the fact that $j \in N_{it}$ is not itself informative of τ_{ij}.

(ii) Costly Observation. We may also suppose that observing someone's action requires costly effort, where this cost varies across individuals and across time. Specifically, if i selects $j_{it} = j$ then he incurs the cost c_{ij}^t where the vector of these costs $\mathbf{c}_{it}$ is drawn from a distribution with the convex hull of the support being $[0, \bar{c}]^{n-1}$ and $\bar{c} > v$.

(iii) Random Observation. Finally, consider the effect of individuals having imperfect control over who they learn from. Suppose that in each period there is a positive probability π that an individual can choose who he observes, but also a positive probability $1-\pi$ that nature will determine j_{it} according to some (stationary) exogenous process.

In each of these, individuals inevitably acquire information about the behavior of a wider number of individuals. As time progresses, this information provides a chance that individuals will shift away from those they have chosen to learn from and to observe those who have exhibited a stronger record of agreement. This results in the following.

Proposition 8. *When the prior Q is symmetric, conditions (12) and (15) hold, and assuming either (i), (ii), or (iii),*

(a) Limiting in-group disagreement is less than out-group disagreement

$$\lim_{t\to+\infty} \mathbb{E}[d_t] \leq \lim_{t\to+\infty} \mathbb{E}[d_t']. \tag{16}$$

(b) In-group disagreement decreases in the limit, in the sense that

$$\lim_{t\to+\infty} \mathbb{E}[d_t] \leq \mathbb{E}[d_1]. \tag{17}$$

Assuming Q to be symmetric is stronger than needed for part (a) of the proposition. What matters is that, in *(i)* and *(iii)*, nature is not biased towards linking those with low similarity. Specifically, if $N_i = \bigcap_{t=1}^{\infty} \bigcup_{k>t}^{\infty} N_{it}$ is the set individuals who enter N_{it} infinitely often , then it is sufficient that (on average) the expected disagreement with at least one $\hat{j} \in N_i$ is less than or equal to the average expected disagreement $\mathbb{E}\left[d_{i\hat{j}}\right] \leq \frac{1}{n-1}\sum_{j\neq i}\mathbb{E}\left[d_{ij}\right]$.

For example, this condition is satisfied if $n_i \geq 1$ individuals are randomly and with uniform probability selected to be members of N_i.

4 Concluding Remarks

The literature on social learning has demonstrated that the presence heterogeneity across individuals can complicate and stifle learning. This paper has shown that when heterogeneity is unobservable and itself an object of learning, individuals exhibit the ostensibly irrational behaviors of weighting opinions and actions that agree with their own more highly and seeking to learn from those with whom they expect to most agree.

In the model, learning takes place across a sequence of periods. In each period, an individual receives a private signal and also selects one other individual an observes his action. This allows me to track how behavior adjusts after seeing another individual act and to characterize the evolution of who individuals most prefer to observe. One direction to extend this analysis would be to study the evolution social networks when allowing for a larger variety of network topologies, say, with higher degree or undirected networks.

This paper sheds light on why people segment themselves on the basis of agreement, but does not focus on the welfare effects. Since individuals in the model are Bayesian and fully internalize the effects of their observational choices, there are no inefficiencies. Realistically, there are number of reasons why a segmented population would be marked by inefficiencies. For one, if individuals can choose who they share their information with and prefer to share with those with similar observable characteristics (for instance, due to homophily), then they could deprive those who are different from themselves of valuable information. Moreover, outcomes can clearly be inefficient if agreement-based observations are driven by bounds in rationality, such as motivated reasoning or a desire for certainty. What my results suggest is that attempts to assess the inefficiencies posed by echo chamber phenomena must also consider that these behaviors emerge naturally by factors as ubiquitous as unobserved heterogeneity.

References

Acemoglu, Daron, M. A. Dahleh, I. Lobel, and A. Ozdaglar. 2011. "Bayesian Learning in Social Networks." *The Review of Economic Studies*, 78(4): 1201–1236.

Ali, S. Nageeb. 2018. "Herding with costly information." *Journal of Economic Theory.*

Athey, Susan, and Jonathan Levin. 2018. "The value of information in monotone decision problems." *Research in Economics.*

Bala, V, and S Goyal. 1995. "A Theory of Learning with Heterogeneous Agents." *International Economic Review*, 36(2): 303–323.

Banerjee, Abhijit, and Drew Fudenberg. 2004. "Word-of-mouth learning." *Games and Economic Behavior*, 46(1): 1–22.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics*, 107(3): 797–817.

Becker, Gary S. 1991. "A Note on Restaurant Pricing and Other Examples of Social Influences on Price." *Journal of Political Economy.*

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, 100(5): 992.

Blackwell, David. 1951. "Comparison of Experiments."

Boutyline, Andrei, and Robb Willer. 2017. "The Social Structure of Political Echo Chambers: Variation in Ideological Homophily in Online Networks." *Political Psychology.*

Chamley, C. 2004. *Rational herds: Economic models of social learning.* Cambridge, U.K.:Cambridge University Press.

Che, Yeon-Koo, and Konrad Mierendorff. 2019. "Optimal Dynamic Allocation of Attention." *American Economic Review*, 109(8): 2993–3029.

Currarini, Sergio, Matthew O Jackson, and Paolo Pin. 2009. "An Economic Model of Friendship: Homophily, Minorities, and Segregation." *Econometrica*, 77(4): 1003–1045.

Del Vicario, Michela, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. 2016. "The spreading of misinformation online." *Proceedings of the National Academy of Sciences.*

Ellison, Glenn, and Drew Fudenberg. 1995. "Word-of-mouth communication and social learning." *The Quarterly Journal of Economics*, 110(1): 93–125.

Eyster, Erik, and Matthew Rabin. 2014. "Extensive imitation is irrational and harmful." *Quarterly Journal of Economics.*

Gale, Douglas, and Shachar Kariv. 2003. "Bayesian learning in social networks." *Games and Economic Behavior*, 45(2): 329–346.

Goeree, Jacob K., Thomas R. Palfrey, and Brian W. Rogers. 2006. "Social learning with private and common values." *Economic Theory*, 28(2): 245–264.

Golub, Benjamin, and M.O. Jackson. 2008. "How homophily affects diffusion and learning in networks." *Arxiv preprint arXiv:0811.4013*, 1–50.

Halberstam, Yosh, and Brian Knight. 2016. "Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter." *Journal of Public Economics.*

Hendricks, Kenneth, Alan Sorensen, and Thomas Wiseman. 2012. "Observational learning and demand for search goods." *American Economic Journal: Microeconomics.*

Jann, Ole, and Christoph Schottmüller. 2018. "Why Echo Chambers are Useful." *University of Oxford, Department of Economics Economics Series Working Papers, 2018.*

Jasny, Lorien, Joseph Waggle, and Dana R. Fisher. 2015. "An empirical examination of echo chambers in US climate policy networks." *Nature Climate Change.*

Jelveh, Zubin, Bruce Kogut, and Suresh Naidu. 2018. "Political language in economics." *Columbia Business School Research Paper*, , (14-57).

Jensen, Richard. 1982. "Adoption and diffusion of an innovation of uncertain profitability." *Journal of Economic Theory.*

Jeuland, Abel P. 1987. "Incorporating heterogeneity into parsimonious models of diffusion of innovation." *University of Chicago, Center for Research in Marketing.*

Kirman, A. 1993. "Ants, Rationality, and Recruitment." *The Quarterly Journal of Economics.*

Kunda, Ziva. 1990. "The case for motivated reasoning." *Psychological Bulletin.*

Lehmann, E. L. 1988. "Comparing Location Experiments." *The Annals of Statistics.*

Levin, Jonathan. 2001. "Information and the Market for Lemons." *The RAND Journal of Economics.*

Lobel, Ilan, and Evan Sadler. 2016. "Preferences, Homophily, and Social Learning." *Operations Research*, 64(3).

Mahajan, Vijay, and Robert A. Peterson. 1985. *Models for Innovation Diffusion.* Beverly Hills:Sage Publications.

Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics.*

Möbius, Markus, and Tanya Rosenblat. 2014. "Social Learning in Economics." *Annual Review of Economics*, 6(1): 140512170917005.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2014. "Asymptotic learning on Bayesian social networks." *Probability Theory and Related Fields.*

Mueller-Frank, Manuel, and Mallesh M. Pai. 2016. "Social learning with costly search." *American Economic Journal: Microeconomics.*

Munshi, Kaivan. 2004. "Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution." *Journal of Development Economics.*

Nickerson, Raymond S. 1998. "Confirmation bias: A ubiquitous phenomenon in many guises." *Review of General Psychology*, 2(2): 175–220.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille. 2009. "Informational externalities and emergence of consensus." *Games and Economic Behavior.*

Schelling, Thomas C. 1969. "Models of Segregation." *American Economic Review.*

Schlosshauer, Maximilian, Johannes Kofler, and Anton Zeilinger. 2013. "A snapshot of foundational attitudes toward quantum mechanics." *Studies in History and Philosophy of Science Part B - Studies in History and Philosophy of Modern Physics.*

Sethi, Rajiv, and Muhamet Yildiz. 2012. "Public disagreement." *American Economic Journal: Microeconomics.*

Sethi, Rajiv, and Muhamet Yildiz. 2016. "Communication With Unknown Perspectives." *Econometrica*, 84(6).

Smith, Lones, and Peter Sorensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica*, 68(2): 371–398.

Young, H Peyton. 2009. "Innovation diffusion in heterogeneous populations: Contagion, social influence, social learning." *American Economic Review*, 99(5): 1899–1924.

et al. Abbot, Patrick. 2011. "Responses to The Evolution of Eusociality, Nowak et al. Nature. 466." *Nature*, 471(7339): E1–4; author reply E9–10.

A Appendix: Preliminaries

Lemma A.1. *Let X and Y be random variables taking values in $\{a, b\}$, with symmetric joint pmf $f(x, y)$ and identical nondegenerate marginal pmf's $f(x)$ and $f(y)$. If $\rho = corr(X, Y)$ is the correlation coefficient of X and Y, then the joint pmf can be represented by*

$$f(x, y) = \rho f(x)\mathbf{1}_{x=y}(x, y) + (1-\rho)f(x)f(y). \tag{18}$$

Proof. Using the assumed symmetry, we can express the covariance

$$Cov(X, Y) = \sum_{x,y} f(x, y)(x - \mathbb{E}[x])(y - \mathbb{E}[y]) \tag{19}$$

$$= a^2\left(f(a, a) - f(a)^2\right) + 2ab\left(f(a, b) - f(a)f(b)\right) + b^2\left(f(b, b) - f(b)^2\right) \tag{20}$$

By symmetry as well as the law of total probability ($f(a) = f(a, a) + f(a, b)$ and $f(b) = f(b, b) + f(a, b)$), we obtain

$$f(a, a) - f(a) = f(b, b) - f(b) \tag{21}$$

which, subtracting $f(a)^2$ from both sides yields

$$f(a, a) - f(a)^2 = f(b, b) - f(b)^2. \tag{22}$$

The identity $f(a) = f(a, a) + f(a, b)$ rearranges to $f(a, b) - f(a)f(b) = f(a)^2 - f(a, a)$ and thus the covariance becomes

$$(a - b)^2(f(a, a) - f(a)^2). \tag{23}$$

Dividing by the product of the standard deviations

$$var[X]^{1/2}\, var[Y]^{1/2} = f(a)f(b)(b - a)^2$$

obtains the correlation coefficient

$$\rho = \frac{f(a, a) - f(a)^2}{f(a)f(b)} = \frac{f(b, b) - f(b)^2}{f(a)f(b)} = -\frac{f(a, b) - f(a)f(b)}{f(a)f(b)}. \tag{24}$$

Rearranging these equations obtains the desired representation. $\square$

Lemma A.2. *Let X and Y be random variables with identical marginal cdf's $F(x)$ and $F(y)$ and a joint cdf $G(x,y|\gamma)$ parameterized by $\gamma \in \mathbb{R}$. If $\gamma' \geq \gamma$ and $y' \geq y$, then*

(i) $G(x,y|\gamma) \leq G(x,y|\gamma')$ *and* $G(x|Y=y,\gamma) \preceq_{FOSD} G(x|Y=y',\gamma)$ *if and only if* $Y|\gamma'$ *is more Lehmann informative for* X *than* $Y|\gamma$.

(ii) $G(x,y|\gamma) \leq G(x,y|\gamma')$ *implies* $Cov(X,Y|\gamma) \leq Cov(X,Y|\gamma')$.

Proof. Part *(i)* follows by first noting that, since the marginal of Y does not depend on γ, $G(x,y|\gamma) \leq G(x,y|\gamma')$ is equivalent to what Athey and Levin (2018) term *monotone information order*: $G(x|F(y) \leq k,\gamma) \leq G(x|F(y) \leq k,\gamma')$ if and only $G(x|y \leq F^{-1}(k),\gamma) \leq G(x|y \leq F^{-1}(k),\gamma')$. By Athey and Levin (2018) and Levin (2001), $Y|\gamma'$ is more informative for X than $Y|\gamma$

Part (ii) follows from Hoeffding's well-known covariance identity

$$Cov(X,Y|\gamma) = \iint (G(x,y|\gamma) - F(x)F(y))\,\mathrm{d}x\mathrm{d}y. \tag{25}$$

Since the integrand is everywhere increasing, so too must the integral. $\square$

B Appendix: Proofs for Main Text

Proof of Proposition 6. Lemma A.2 guarantees that $G(A|B,\tau_{ij}) = \Pr(x_{jt} \leq A|x_{it} \leq B,\tau_{ij})$ and $\tilde{G}(A|B,\tau_{ij}) = \Pr(x_{jt} > A|x_{it} > B,\tau_{ij})$ are increasing in τ_{ij}.

$$\begin{aligned} E^C &= \{\mathbf{x} : |x_{jt} - x_{it}| \geq \delta\} = \{\mathbf{x} : x_{jt} \leq x_{it} - \delta\} \cup \{\mathbf{x} : x_{jt} \geq x_{it} + \delta\} && (26)\\ &= \{(\mathbf{x},c) : x_{jt} \leq c - \delta, x_{it} \geq c\} \cup \{(\mathbf{x},c) : x_{jt} \geq c + \delta, x_{it} \leq c\} && (27)\\ &= E_1^C \cup E_2^C. && (28) \end{aligned}$$

We can write

$$\Pr(E_1^C|\tau_{ij}) = \Pr(x_{jt} \leq c - \delta | x_{it} \geq c, \tau_{ij}) \Pr(x_{it} \geq c) \tag{29}$$

$$= 1 - \int_{-\infty}^{\infty} \tilde{G}\left(c - \delta | c, \tau_{ij}\right) \Pr(x \geq c) \mathrm{d}c \tag{30}$$

and

$$\Pr(E_2^C|\tau_{ij}) = \Pr(x_{jt} \geq c + \delta | x_{it} \leq c, \tau_{ij}) \Pr(x_{it} \leq c) \tag{31}$$

$$= 1 - \int_{-\infty}^{\infty} G\left(c + \delta | c, \tau_{ij}\right) \Pr(x \leq c) \mathrm{d}c \tag{32}$$

Notice that the integrands in (30) and (32) are both increasing in τ_{ij} and thus $Pr(E^C|\tau_{ij})$ is decreasing in τ_{ij}.

Notice that the ratio

$$\frac{q(\tau_{ij}|E)}{q(\tau_{ij}|E^C)} = \frac{1 - \Pr(E^C|\tau_{ij})}{\Pr(E^C|\tau_{ij})} \frac{\Pr(E^C)}{\Pr(E)} \tag{33}$$

is increasing in τ_{ij} and thus $Q(\cdot|E)$ has FOSD over $Q(\cdot|E^C)$. This is equivalent to (13) and directly implies Proposition 1 and Corollary 1. □

Proof of Propositions 2 and 7. By Doob's consistency theorem the law of large numbers, as i's observations of j grow large, j's social influence on i almost surely converges to

$$\mathcal{S}_{ij}^{\infty} = \mathbb{E}\left[\left(y_{it} - x_{it}\right)^2 | \tau_{ij}\right]^{\frac{1}{2}} = \sqrt{\mathrm{Var}\left(y_{it}|\tau_{ij}\right) - \mathrm{Var}\left(x_{it}\right)} \tag{34}$$

We can relate social influence to expected utility by the law of total variance

$$\mathrm{Var}\left(\theta_{it}\right) = \mathrm{Var}\left(\mathbb{E}\left[\theta_{it}|x_{it}, x_{jt}, \tau_{ij}\right]\right) + \mathbb{E}\left[\mathrm{Var}\left(\theta_{it}|x_{it}, x_{jt}, \tau_{ij}\right)\right] \tag{35}$$

$$= \mathrm{Var}\left(y_{it}|\tau_{ij}\right) - \mathbb{E}\left[U\left(y_{it}\right)|\tau_{ij}\right]. \tag{36}$$

Combining equations (34) and (36) and noting that the expected payoff from observing j is increasing in similarity (by Lehmann (1988)) reveals the limiting social influence to be almost surely increasing in similarity.

Asymptotic disagreement almost surely equals

$$d_{ij}^{\infty} = \mathbb{E}\left[(x_{it} - x_{jt})^2 | \tau_{ij}\right]^{1/2} \tag{37}$$
$$= (\text{Var}(x_{it}) + \text{Var}(x_{jt}) - 2Cov(x_{it}, x_{jt}|\tau_{ij}))^{1/2}. \tag{38}$$

Since Lemma A.2 guarantees the covariance term to be increasing in similarity, the asymptotic disagreement is thus decreasing in similarity. Hence, asymptotic social influence is decreasing in asymptotic disagreement. □

Lemma B.1. *$Pr(x(s_{it}) < \alpha | x(s_{jt}) < \beta, \tau_{ij})$ is increasing in similarity in the baseline model.*

Proof. By Lemma 1 in Smith and Sorensen (2000), we can, without loss in generality, identify each signal with the corresponding private belief that the parameter is $\theta_{it} = b$. As actions are monotonically increasing in signals, we obtain

$$Pr(x(s_{it}) < \alpha | x(s_{jt}) < \beta, \tau_{ij}) = Pr(s_{it} < \alpha' | s_{jt} < \beta', \tau_{ij}) \tag{39}$$

where $\alpha' = x^{-1}(\alpha)$ and $\beta' = x^{-1}(\beta)$. Using Lemma A.1 we can unpack this expression as

$$Pr(s_{it} < \alpha' | s_{jt} < \beta', \tau_{ij}) = \tag{40}$$
$$(F_a(\alpha')Pr(a|s_{jt} < \beta') + F_b(\alpha')Pr(b|s_{jt} < \beta'))\tau_{ij} \tag{41}$$
$$+ (F_a(\alpha')Pr(a|s_{jt} < \beta') + F_b(\alpha')Pr(b|s_{jt} < \beta'))(1 - \tau_{ij}). \tag{42}$$

From Smith and Sorensen (2000), we know that F_b has FOSD over F_a, immediately implying that $F_a(\alpha') \geq F_b(\alpha')$ and thus $Pr(a|s_{jt} < \beta') = (1 + \frac{p_b}{p_a}\frac{F_b}{F_a}(\beta'))^{-1} \geq p_a$. □

Lemma B.2. *In variations I.-IV., when similarity is known and $\tau_{ik} \leq \tau_{ij}$*

(i) Individual i finds x_{jt} to be more Lehmann informative than x_{kt}.

(ii) x_{jt} is more Lehmann informative than x_{kt} for learning i's action.

Proof. First, recall from Lehmann (1988), if $X \sim \mathcal{N}(a\theta + b, \sigma^2)$ and $Y \sim \mathcal{N}(c\theta + d, \gamma^2\sigma^2)$, then $\gamma \in (0, 1)$ implies that X is Blackwell more informative than Y for θ as

$$X = \left(\frac{a}{c}\right) Y + Z \tag{43}$$

where $Z \sim \mathcal{N}\left(b - ad/c, (1 - \gamma^2)\sigma^2 a^2\right)$.

(i) From i's vantage point, j's action is normally distributed with the following.

	Mean	Variance
Private Values	$\alpha\tau_{ij}\theta_{it} + (1 - \tau_{ij})\theta_0$	$\alpha^2\left(\left(1 - \tau_{ij}^2\right) + \sigma^2\right)$
Common and Private Values	$\alpha\theta_t + (1 - \alpha)\theta_0 + \tau_{ij}\eta_{it}$	$\alpha^2\sigma^2 + 1 - \tau_{ij}^2$
Subjective Probability Distributions	$\alpha\left(\theta_t + (1 - \tau_{ij})\eta_{it}\right) + (1 - \alpha)\theta_0$	$\alpha^2(\sigma^2 + 1 - \tau_{ij}^2)$
Priors	$\alpha\theta_t + (1 - \alpha)\tau_{ij}\eta_{it}$	$\alpha^2\sigma^2 + (1 - \alpha)^2(1 - \tau_{ij}^2)$

(ii) To the outside observer seeking to learn x_{it}, j's action normally distributed with

	Mean	Variance
Private Values	$\alpha\tau_{ij}x_{it} + (1 - \alpha\tau_{ij})\theta_0$	$\alpha\sigma_0^2\left(1 - \alpha^2\tau_{ij}^2\right)$
Common and Private Values	$\frac{\alpha^2\sigma_0^2+\tau_{ij}}{\alpha\sigma_0^2+1}x_{it} + \left(1 - \frac{\alpha^2\sigma_0^2+\tau_{ij}}{\alpha\sigma_0^2+1}\right)\theta_0$	$\left(1 - \left(\frac{\alpha^2\sigma_0^2+\tau_{ij}}{\alpha\sigma_0^2+1}\right)^2\right)(\alpha\sigma_0^2 + 1)$
Subjective Probability Distributions	$\frac{\sigma_0^2+\tau_{ij}}{\sigma_0^2+\sigma^2+1}x_{it} + \left(1 - \frac{\sigma_0^2+\tau_{ij}}{\sigma_0^2+\sigma^2+1}\right)\theta_0$	$\left(1 - \left(\frac{\sigma_0^2+\tau_{ij}}{\sigma_0^2+\sigma^2+1}\right)^2\right)\alpha^2(\sigma_0^2 + \sigma^2 + 1)$
Priors	$\eta_0 + \frac{\alpha^2\sigma_0^2+(1-\alpha)^2\tau_{ij}}{\alpha^2(\sigma_0^2+\sigma^2)+(1-\alpha)^2}(x_{it} - \eta_0)$	$(1 - (\frac{\alpha^2\sigma_0^2+(1-\alpha)^2\tau_{ij}}{\alpha^2(\sigma_0^2+\sigma^2)+(1-\alpha)^2})^2)(\alpha^2(\sigma_0^2 + \sigma^2) + (1 - \alpha)^2)$

□

Proof of Proposition 4. The result follows by using iterated expectations to express these terms as $\mathbb{E}\left[d_t\right] = \frac{1}{n}\sum_i \mathbb{E}\left[\mathbb{E}\left[d_{ij_{it}}|h_{it}\right]\right]$ and $\mathbb{E}\left[d_t'\right] = \frac{1}{n(n-2)}\sum_{i,j\neq i,j_{it}} \mathbb{E}\left[\mathbb{E}\left[d_{ij}(t)|h_{it}\right]\right]$ and noting that Proposition 3 guarantees $\mathbb{E}\left[d_{ij_{it}}|h_{it}\right] \leq \frac{1}{n-2}\sum_{j\neq i,j_{it}} \mathbb{E}\left[d_{ij}|h_{it}\right]$ for all $i = 1, 2, ..., n$. Notice that this inequality is strict at any history where an individual strictly prefers his choice to some deviation since this would imply that $\mathbb{E}\left[\tau_{ij_{it}}|h_{it}\right] > \mathbb{E}\left[\tau_{ij}|h_{it}\right]$ and hence $\mathbb{E}\left[d_{ij_{it}}|h_{it}\right] < \mathbb{E}\left[d_{ij}|h_{it}\right]$ for some $j \neq i, j_{it}$. □

Proof of Proposition 8. Given our previous work, the assumption that Q is symmetric makes the proof of this result straightforward; though, as I discuss in the text, a weaker assumption would suffice.

(i) By revealed preference, the similarity of any individual i who chooses to observe infinitely often (call him j^*) exceeds that of any $\hat{j}$ whom i observes infinitely often at the end of a period. En route to proving Proposition 6, I show that condition 12 implies that $\mathbb{E}[d_{ij^*}] \leq \mathbb{E}[d_{i\hat{j}}]$. Since $\Pr[j \in N_{it}]$ is independent of the realized similarity, the expected disagreement with $\hat{j}$ is equal to the prior expected disagreement $\mathbb{E}\left[d_{i\hat{j}}\right] = \mathbb{E}[d_{ij}]$. Since priors are symmetric,

$$\mathbb{E}\left[\sum_{j\neq i} d_{ij}\right] = \sum_{j\neq i} \mathbb{E}[d_{ij}] = (n-1)\mathbb{E}[d_{ij}] \tag{44}$$

and also

$$\mathbb{E}\left[\sum_{j\neq i} d_{ij}\right] = \mathbb{E}[d_{ij^*}] + (n-2)\mathbb{E}\left[d_{ij}\middle| j \neq i, j^*\right]. \tag{45}$$

Combining these observations finds $\mathbb{E}[d_{ij^*}] \leq \mathbb{E}[d_{ij}] \leq \mathbb{E}[d_{ij}|j \neq i, j^*]$ from which (a) and (b) directly follow.

(ii) In the first period, i will observe the individual with the lowest cost, with the expected disagreement only depending on the prior $\mathbb{E}[d_{ij}]$. Asymptotically, if i ignores what he learns about similarity and continues to only observe the individual with the lowest cost, then the expected disagreement remains $\mathbb{E}[d_{ij}]$. If he ever chooses to observe an individual with a higher cost, then it must be that they have a higher similarity; hence, $\lim_{t\to+\infty} \mathbb{E}[d_t] \leq \mathbb{E}[d_{ij}]$. Thus (a) and (b) follows by similar reasoning as before.

(iii) The proof is similar to (i), the only difference being that $\mathbb{E}[d_t] = \pi\mathbb{E}[d_{ij^*}(t)] + (1-\pi)\mathbb{E}[d_{ij}(t)]$. □