Echo Chambers: Social Learning under Unobserved Heterogeneity

Cole Williams*

January 25, 2021

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.

---

*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.

1 Introduction

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 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

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

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 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

[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).

two dimensions. Firstly, holding fixed an individual's beliefs over the structure of heterogeneity in the population, social learning influences an individual'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.

Local learning further drives the evolution of an individual's choice of whom to observe. 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. Moreover, as time progresses, the degree to which an individual's behavior is influenced by social learning increases.

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

[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).

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.

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 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, 2012; 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.

Related ideas can be found in explanations for media bias (Gentzkow and Shapiro, 2006; Gentzkow, Shapiro and Stone, 2016). In a setting where the consumers of media are biased in their prior views, a firm may choose to slant reports in the direction of consumers beliefs. The idea is that bias prevents consumers from drawing negative inferences about the firm's quality, keeping its reputation intact. Che and Mierendorff (2019) study the problem of choosing a media source when the bias is publicly known. 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.

The remainder of the paper is organized as follows. Section 2 presents the model of dual learning. Section 3 analyzes the model, characterizing how individuals act after observing others (3.1), how behavior can respond non-monotonically to disagreement (3.2), and the emergence of echo chambers (3.3). Section 4 discusses how the results extend when heterogeneity comes in terms of heterogeneous priors or subjective probability distributions and Section 5 concludes.

2 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

$$\theta_{it} = \theta_t + \eta_{it} \tag{1}$$

where θ_t is unknown and common to everyone and η_{it} is specific to i and only known to him. Both components are normally distributed $\theta_t \sim \mathcal{N}(\theta_0, \sigma_0^2)$ and $(\eta_{1t}, \eta_{2t}, ..., \eta_{nt}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\tau})$ where $\mathbf{0}$ is the zero vector of length n and $\boldsymbol{\tau}$ is an $n \times n$ correlation matrix with element $\tau_{ij} = \tau_{ji} = corr(\eta_{it}, \eta_{jt}) \in [0, 1]$ called the *similarity* of i and j. Individuals do not directly observe $\boldsymbol{\tau}$, but it is commonly known to be drawn before the first period from a symmetric nondegenerate distribution $Q(\boldsymbol{\tau})$ and to remain fixed over time.

Upon entering period t, individual i receives a normally distributed private signal $s_{it} \sim \mathcal{N}(\theta_t, \sigma^2)$. With his signal in hand, individual i forms his *opinion* $x_{it} = \mathbb{E}[\theta_{it}|s_{it}, \eta_{it}]$. At this stage, with probability π picks out one other individual $j \neq i$ and observes his opinion x_{jt}. With complementary probability $1-\pi$, nature selects j according to some known stochastic process. Since their preferences are 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 beliefs over τ_{ij}. Letting h_{it}

denote the history of i's observations up to time t, refer to the marginal CDF $Q(\tau_{ij}|s_{it}, \eta_{it}, x_{jt}, h_{it})$ as i's *perceived similarity* of j. Finally, i takes an action $y_{it} \in \mathbb{R}$ which yields quadratic loss payoffs

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

Individuals do not see who others chose to observe or any other actions before acting so that the difference between y_{it} and x_{it} solely reflects what is learned from observing j.

After taking his action y_{it} individuals observe their realized payoffs and the period concludes. 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.

Remark 1. *One way to think about similarity is to imagine each individual i being endowed with a fixed underlying trait* $\xi_i \in \mathbb{R}$ *drawn according to some known distribution at the outset. In each period, private values depend on traits and are assigned according to a Gaussian process* $\eta(\xi) \sim \mathcal{GP}(0, \mathbf{k})$ *where* $\mathbf{k}(|\xi - \xi'|)$ *is the covariance function that only depends on the distance between traits. Essentially, those with closer traits experience higher correlation in their private values. For individuals with arbitrarily distant traits, private values are approximately independent.*

3 Analysis

3.1 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

[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 $y_{it}(s_{it}, x'_{jt}) \geq y_{it}(s_{it}, x_{jt})$.

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 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 then forms his opinion $x_{it} = \mathbb{E}[\theta_{it}|s_{it}, \eta_{it}]$. Supposing i is able to choose whom he observes, letting h_{it} denote the history of i's observations up to time t, i chooses an individual j, observes his opinion x_{jt}, and selects his action y_{it}, both to maximize

$$\max_{j} \mathbb{E}\left[\max_{\tilde{y}_{it}} \mathbb{E}\left[u(\tilde{y}_{it}; \theta_{it})|s_{it}, \eta_{it}, x_{jt}, h_{it}\right] \middle| s_{it}, \eta_{it}, h_{it}\right]. \tag{3}$$

For completeness, there multiple individuals are equally desirable to observe, suppose i selects one of them, each with equal probability. Since preferences may be correlated, observing how j acts allows i to update his beliefs over θ_{it} and also the perceived similarity $Q(\tau_{ij}|s_{it}, \eta_{it}, x_{jt}, h_{it}) = Q(\tau_{ij}|h_{it+1})$.

Using the distributional assumptions, we can write i's action as

$$y_{it} = \bar{\alpha} x_{jt} + \bar{\beta} x_{it} + \bar{\gamma}\theta_0 + \bar{\kappa}\eta_{it} \tag{4}$$

where $\bar{\ell} \equiv \int \left[\ell(\tau_{ij})\right] \mathrm{d}Q(\tau_{ij}|h_{it+1})$ for $\ell \in \{\alpha, \beta, \gamma, \kappa\}$ with the particular expressions for the integrands found in A.1. This representation follows simply from iterated expectations $y_{it} = \mathbb{E}\left[\mathbb{E}\left[\theta_{it}|h_{it+1}, \tau_{ij}\right] |h_{it+1}\right]$ and noting that the inner expectation takes the familiar additive form since it is a posterior expectation in the normal-normal conjugate family.

Notice that j's opinion influences i's actions in two ways, capturing the two dimensions of dual learning. First, if we fix the weights $(\bar{\alpha}, \bar{\beta}, \bar{\gamma}, \bar{\kappa})$, i's response to observing j is qualitatively the same as if they were homogeneous, y_{it} is monotonically increasing in x_{jt}. Second, x_{jt} enters the perceived similarity, and thus the weight i places on j's opinion is itself determined by the opinion.

The model holds two related properties that drive the results. First, the distribution of actions satisfies the following condition

$$\Pr\left[x_{it} < A, x_{jt} < B|\tau_{ij}\right] \leq \Pr\left[x_{it} < A, x_{jt} < B|\tau'_{ij}\right] \tag{5}$$

for all real A and B 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,[6] condition (5) 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.1).

This condition turns out to be sufficient for individuals to gauge similarity on the basis of agreement. Formally, define the *disagreement* between i and j to be difference between their actions $d_{ij}(t) = |x_{jt} - x_{it}|$.

Lemma 1. *When condition (5) is satsified, then for any δ in the support of disagreement*

$$\mathbb{E}\left[Q(\tau_{ij}|h_{it+1})\big|h_{it}, d_{ij}(t) \leq \delta\right] \leq \mathbb{E}\left[Q(\tau_{ij}|h_{it+1})\big|h_{it}, d_{ij}(t) \geq \delta\right]. \tag{6}$$

We now establish that i adjusts the weight placed on j's opinion 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. *On average, the weight placed on j's opinion is larger when there is less disagreement. That is, for $\delta > 0$*

$$\mathbb{E}\left[\bar{\alpha}|h_{it}, d_{ij}(t) \leq \delta\right] > \mathbb{E}\left[\bar{\alpha}|h_{it}, d_{ij}(t) > \delta\right].$$

In other words, if we compare two pairs of opinions (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}. The proof follows from noting that $\alpha(\tau_{ij})$ is increasing and by Lemma 1 larger disagreement yields an expected perceived

[6] 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}.

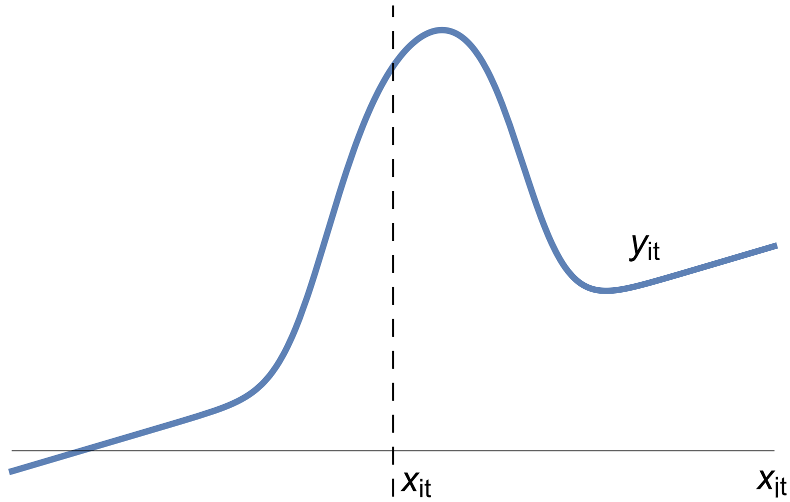

Figure 1: i's action as a function of j's opinion

similarity first order stochastically dominated by that with smaller disagreement. By the same reasoning, because $\beta'(\tau_{ij}) < 0$, $\gamma'(\tau_{ij}) < 0$, and $\kappa'(\tau_{ij}) > 0$, we can conclude that smaller disagreement leads on average to less weight on one's own opinion and the prior mean, but more weight on the personal bias.

3.2 Non-Monotonicity in Disagreement

We have seen that disagreement creates a tension between the two dimensions of learning. This section explores how this tension can produce a non-monotonic relationship between disagreement and social influence. As an example, suppose a community seeks to learn the merits of various political policies with the space of policies described by the real line. Figure 1 depicts how 1 preferred policy y_{it} changes as we alter j's opinion x_{jt} in the simple case where similarity can be either low or high $\tau_{ij} \in \{\tau^L, \tau^H\}$.[7] The direct effect of shifting j's opinion when holding the weights in (4) fixed leads y_{it} to be strictly increasing. As an indirect effect of shifting j's opinion, substantial disagreement between the two parties leads i to reduce the weight placed on j's opinion altogether. As depicted in the figure, large enough disagreement can go so far as to lead i's action to be decreasing in j's opinion; precisely the opposite of what occurs in the setup without unobserved heterogeneity.

[7]Parameter values are set at:
$(\theta_0, \sigma_0, \sigma, \tau^L, \tau^H, \eta_{it}, x_{it}, P(\tau_{ij} = \tau^H | h_{it})) = (0, 1, 5, 0.2, 0.8, 0.25, 0, 0.5)$.

We now formally characterize the portion of the domain on which i's action moves negatively with j's action. Sticking to the simple case where similarity can be either low or high let y_{it}^L and y_{it}^H be the actions taken when similarity is known to be low and high respectively, $\hat{y}_{it} \equiv y_{it}^H - y_{it}^L$ the difference between the two, and $q^H \equiv P(\tau_{ij} = \tau^H | h_{it+1})$. Define the elasticity of the perceived similarity to be

$$\varepsilon \equiv -\frac{\frac{\partial q^H}{\partial x_{jt}} \cdot \hat{y}_{it}}{\frac{\partial \hat{y}_{it}}{\partial x_{jt}} \cdot q^H}. \tag{7}$$

Proposition 2. *Suppose similarity is either low or high $\tau_{ij} \in \{\tau^L, \tau^H\}$. Individual i's action is decreasing in the observed individual j's opinion when the perceived similarity is sufficiently elastic $\varepsilon \cdot q^H > \psi$ for some constant $\psi > 0$.*

When examining how people respond to others' opinions in an environment of social learning, it would seem a natural dictate of rationality that one's response ought to move in the same direction as the observed opinion. The thrust of this section is that this not true in a world with multiple dimensions of uncertainty. In particular, when individuals attribute a difference in opinion to alternative underlying differences, stronger disagreement can yield less influence on behavior.

3.3 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.

We want to compare the level a disagreement between those who exchange opinions and those who do not in each period. 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 $D_t \equiv \frac{1}{n} \sum_{i,j} d_{ij}(t) \mathbf{1}_{j=j_{it}}$. As its complement, let *out-group* disagreement denote the average disagreement that an individual does not observe $D_t' \equiv \frac{1}{n(n-2)} \sum_{i,j} d_{ij}(t) \mathbf{1}_{j \neq i, j_{it}}$.

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 3.

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

$$\lim_{t\to+\infty} \mathbb{E}[D_t] \leq \lim_{t\to+\infty} \mathbb{E}[D'_t]. \tag{8}$$

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

$$\lim_{t\to+\infty} \mathbb{E}[D_t] \leq \mathbb{E}[D_1]. \tag{9}$$

As time progresses, if i consistently encounters strong disagreement with some individual, i will eventually find it undesirable to continue learning from him. When considering the larger network structure, those who exchange opinions are going to enjoy closer agreement than if we were to compare individuals who do exchange opinions. What is more, this difference will grow over time.

As individuals find themselves increasingly encountering opinions closer to their own, an important question is how this affects their behavior. Since x_{it} is how i acts with only private information and y_{it} is how he acts after observing j_{it}, define $(y_{it} - x_{it})^2$ to be the *social influence* of j_{it} on i in period t. For example, this captures how a conversation with j_{it} shifts i's views on a political policy.

Proposition 4. *Expected social influence is increasing over time: For $t < t'$*

$$\mathbb{E}\left[(y_{it'} - x_{it'})^2\right] > \mathbb{E}\left[(y_{it} - x_{it})^2\right].$$

Although the opinions individuals encounter are increasingly closer to their own, the influence of social learning on their behavior grows over time. The reasons are two-fold. For one, growing increasingly familiar with how similar others are makes it easier to draw inferences from their opinions. Secondly, the ability to shift who one learns from makes it possible to find others who are more similar and thus more informative for oneself.

4 Discussion

4.1 Variations

The paper focused on the case where unobserved heterogeneity comes at the level of preferences. We may also consider cases where individuals share the same preferences but differ on epistemic grounds, as in the following.

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 $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\tau})$.

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 $s_{it} \sim \mathcal{N}\left(\theta_t, \sigma^2\right)$. However, rather than an individual's prior corresponding to a fixed "perspective" as in Sethi and Yildiz (2012, 2016), priors are drawn each period $\boldsymbol{\eta}_t \sim \mathcal{N}(\boldsymbol{\eta}_0, \boldsymbol{\tau})$.

In both variations the particular expressions for the weighting functions in (4) differ, but all of the propositions hold unchanged. For illustration, we might imagine that each period the community must decide whether to adopt a proposed political policy where θ_t is the ideal policy. Individuals find it easier to learn from other with whom they tend to agree on policy views as they have a tendency to approach learning problems with similar assumptions about how to interpret data or which states are more or less likely.

5 Conclusion

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*, 175: 713–729.

Athey, Susan, and Jonathan Levin. 2018. "The value of information in monotone decision problems." *Research in Economics*, 72(1): 101–116.

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*, 99(5): 1109–1116.

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.

Boutyline, Andrei, and Robb Willer. 2017. "The Social Structure of Political Echo Chambers: Variation in Ideological Homophily in Online Networks." *Political Psychology*, 38(3): 551–569.

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*, 113(3): 554–559.

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*, 129(4): 1861–1898.

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

Gentzkow, Matthew, and Jesse M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy*, 114(2): 280–316.

Gentzkow, Matthew, Jesse M. Shapiro, and Daniel F. Stone. 2016. "Media Bias in the Marketplace." In *Handbook of Media Economics*. 623–645.

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 Matthew O. Jackson. 2012. "How homophily affects the speed of learning and best-response dynamics." *Quarterly Journal of Economics*, 127(3): 1287–1338.

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*, 143: 73–88.

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*, 5(8): 782–786.

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*, 27(1): 182–193.

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*, 108(1): 137–156.

Kunda, Ziva. 1990. "The case for motivated reasoning." *Psychological Bulletin*, 108(3): 480.

Lehmann, E. L. 1988. "Comparing Location Experiments." *The Annals of Statistics*, 16(2): 521–533.

Levin, Jonathan. 2001. "Information and the Market for Lemons." *The RAND Journal of Economics*, 32(4): 657–666.

Lobel, Ilan, and Evan Sadler. 2016. "Preferences, Homophily, and Social Learning." *Operations Research*, 64(3): 564–584.

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*, 12(2): 380–391.

Möbius, Markus, and Tanya Rosenblat. 2014. "Social Learning in Economics." *Annual Review of Economics*, 6(1): 827–847.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2014. "Asymptotic learning on Bayesian social networks." *Probability Theory and Related Fields*, 158(1-2): 127–157.

Mueller-Frank, Manuel, and Mallesh M. Pai. 2016. "Social learning with costly search." *American Economic Journal: Microeconomics*, 8(1): 83–109.

Munshi, Kaivan. 2004. "Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution." *Journal of Development Economics*, 73(1): 185–213.

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*, 66(2): 979–994.

Schelling, Thomas C. 1969. "Models of Segregation." *American Economic Review*, 59(2): 488–493.

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*, 44(3): 222–230.

Sethi, Rajiv, and Muhamet Yildiz. 2012. "Public disagreement." *American Economic Journal: Microeconomics*, 4(3): 57–95.

Sethi, Rajiv, and Muhamet Yildiz. 2016. "Communication With Unknown Perspectives." *Econometrica*, 84(6): 2029–2069.

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: Proofs

A.1 Actions

Let us compute the action i takes after observing j's opinion. For a given similarity, using the usual normal-normal-conjugate updating rule finds

$$\mathbb{E}\left[\theta_{it}|h_{it+1},\tau_{ij}\right] = \alpha(\tau_{ij})x_{jt} + \beta(\tau_{ij})x_{it} + \gamma(\tau_{ij})\theta_0 + \kappa(\tau_{ij})\eta_{it}$$

$$\alpha(\tau_{ij}) \equiv \frac{\sigma^2}{-\frac{\left(\tau_{ij}^2-1\right)\left(\sigma^2+\sigma_0^2\right)^2}{\sigma_0^4} + \frac{\sigma^2\sigma_0^2}{\sigma^2+\sigma_0^2} + \sigma^2}$$

$$\beta(\tau_{ij}) \equiv \frac{\sigma^2 - \frac{\left(\tau_{ij}^2-1\right)\left(\sigma^2+\sigma_0^2\right)^2}{\sigma_0^4}}{-\frac{\left(\tau_{ij}^2-1\right)\left(\sigma^2+\sigma_0^2\right)^2}{\sigma_0^4} + \frac{\sigma^2\sigma_0^2}{\sigma^2+\sigma_0^2} + \sigma^2}$$

$$\gamma(\tau_{ij}) \equiv -\frac{\sigma^4}{\left(\sigma^2+\sigma_0^2\right)\left(-\frac{\left(\tau_{ij}^2-1\right)\left(\sigma^2+\sigma_0^2\right)^2}{\sigma_0^4} + \frac{\sigma^2\sigma_0^2}{\sigma^2+\sigma_0^2} + \sigma^2\right)}$$

$$\kappa(\tau_{ij}) \equiv \frac{\sigma^4\tau_{ij}\sigma_0^4 + \sigma^2(\tau_{ij}-1)\sigma_0^6}{\sigma^6\left(\tau_{ij}^2-1\right) - \sigma^4\sigma_0^2\left(-3\tau_{ij}^2+\sigma_0^2+3\right) + \sigma^2\sigma_0^4\left(3\tau_{ij}^2-2\sigma_0^2-3\right) + \left(\tau_{ij}^2-1\right)\sigma_0^6}$$

with representation (4) follows from taking the expectation $y_{it} = \mathbb{E}\left[\mathbb{E}\left[\theta_{it}|h_{it+1},\tau_{ij}\right]|h_{it+1}\right]$.

A.2 Proof of Lemma 1

Lemma A.1. *For any real numbers A and B and for* $0 \leq \tau_{ij} < \tau'_{ij} \leq 1$

$$\Pr\left[x_{it} < A, x_{jt} < B|\tau_{ij}\right] \leq \Pr\left[x_{it} < A, x_{jt} < B|\tau'_{ij}\right] \tag{10}$$

Proof. As x_{it} and x_{jt} are jointly normal with correlation coefficient $Corr(x_{it}, x_{jt}|\tau_{ij}) = \frac{\nu+\tau_{ij}}{\nu+1}$ with $\nu \equiv \frac{\sigma_0^4\sigma^2}{(\sigma_0^2+\sigma^2)^2}$, Lehmann (1988) provides that $x_{jt}|\tau'_{ij}$ is Lehmann more informative for x_{it} than $x_{jt}|\tau_{ij}$. Let G denote the marginal CDF for both x_{it} and x_{jt}. Since $\Pr(x_{it} < A|x_{jt} = B, \tau_{ij})$ for all real values A and B and similarity τ_{ij}, Athey and Levin (2018) and Levin (2001) that $x_{jt}|\tau'_{ij}$ being Lehmann more informative for x_{it} than $x_{jt}|\tau_{ij}$ is equivalent to the monotone increasing order

$$\Pr\left[x_{it} < A|x_{jt} < G^{-1}(B), \tau_{ij}\right] \leq \Pr\left[x_{it} < |x_{jt} < G^{-1}(B), \tau'_{ij}\right] \tag{11}$$

which is itself equivalent to (10). □

Proof of Lemma 1. Lemma A.1 guarantees that $G(A|B,\tau_{ij}) \equiv \Pr(x_{jt} \leq A|x_{it} \leq$

$B, \tau_{ij})$ and $\tilde{G}(A|B, \tau_{ij}) \equiv \Pr(x_{jt} > A|x_{it} > B, \tau_{ij})$ are increasing in τ_{ij}.

$$
\begin{aligned}
E^C &= \{\mathbf{x} : |x_{jt} - x_{it}| \geq \delta\} = \{\mathbf{x} : x_{jt} \leq x_{it} - \delta\} \cup \{\mathbf{x} : x_{jt} \geq x_{it} + \delta\} && (12)\\
&= \{(\mathbf{x}, c) : x_{jt} \leq c - \delta, x_{it} \geq c\} \cup \{(\mathbf{x}, c) : x_{jt} \geq c + \delta, x_{it} \leq c\} && (13)\\
&= E_1^C \cup E_2^C. && (14)
\end{aligned}
$$

We can write

$$
\begin{aligned}
\Pr(E_1^C|\tau_{ij}) &= \Pr(x_{jt} \leq c - \delta|x_{it} \geq c, \tau_{ij}) \Pr(x_{it} \geq c) && (15)\\
&= 1 - \int_{-\infty}^{\infty} \tilde{G}\left(c - \delta|c, \tau_{ij}\right) \Pr(x \geq c)\mathrm{d}c && (16)
\end{aligned}
$$

and

$$
\begin{aligned}
\Pr(E_2^C|\tau_{ij}) &= \Pr(x_{jt} \geq c + \delta|x_{it} \leq c, \tau_{ij}) \Pr(x_{it} \leq c) && (17)\\
&= 1 - \int_{-\infty}^{\infty} G\left(c + \delta|c, \tau_{ij}\right) \Pr(x \leq c)\mathrm{d}c && (18)
\end{aligned}
$$

Notice that the integrands in (16) and (18) are both increasing in τ_{ij} and thus $Pr(E^C|\tau_{ij})$ is decreasing in τ_{ij}.

Notice that the ratio

$$
\frac{q(\tau_{ij}|E)}{q(\tau_{ij}|E^C)} = \frac{1 - \Pr(E^C|\tau_{ij})}{\Pr(E^C|\tau_{ij})} \frac{\Pr(E^C)}{\Pr(E)} \tag{19}
$$

is increasing in τ_{ij} and thus $Q(\cdot|E)$ has FOSD over $Q(\cdot|E^C)$. This is equivalent to (6) and directly implies Proposition 1. □

Corollary A.1. $\mathbb{E}[d_{ij}(t)|\tau_{ij}] \geq \mathbb{E}[d_{ij}(t)|\tau'_{ij}]$ *if* $\tau_{ij} < \tau_{ij}$.

Proof. Let $F(d_{ij}(t)|\tau_{ij})$ be the induced distribution of disagreement as a function of similarity. To avoid notational confusion, define $z \equiv d_{ij}(t)$. As z is a non-negative random variable, we can write its expectation as

$$
\mathbb{E}[z|\tau_{ij}] = \int_0^{\infty} (1 - F(z|\tau_{ij}))\mathrm{d}z. \tag{20}
$$

The conclusion then follows from the preceding proposition which showed that $(1 - F(z|\tau_{ij}))$ is decreasing in τ_{ij} for all $z \geq 0$. □

A.3 Proof of Proposition 2

Proof of Proposition 2. We can write i's action simply as $y_{it} = y_{it}^H q^H + y_{it}^L(1-q^H)$ where all terms are functions of $h_{it+1} = (s_{it}, \eta_{it}, x_{jt}, h_{it})$. Using the representation given in (4), differentiating with respect to x_{jt} finds $\frac{dy_{it}}{dx_{jt}} < 0$ if and only if

$$\frac{\alpha(\tau^L)}{\alpha(\tau^H) - \alpha(\tau^L)} + q^H < \varepsilon \cdot q^H. \tag{21}$$

Noting that the left side is bounded above by $\psi \equiv \frac{\alpha(\tau^H)}{\alpha(\tau^H)-\alpha(\tau^L)}$ yields the desired conclusion. □

A.4 Proof of Propositions 3 and 4

Let H_{it} denote the possible information sets of i in period t, $H_t = \times_{i=1}^n H_{it}$ be their product, $\mathcal{T}$ denote the set of similarity matrices $\boldsymbol{\tau}$. Denote $\Omega \equiv \mathcal{T} \times_{t\in\mathbb{N}} H_t$ with a typical element by $\omega \in \Omega$. Endowing $\mathcal{T}$ and each H_t with the usual topology and Ω the product topology, the equilibrium strategy induces a probability measure on the Borel σ-algebra over the paths of play Ω. For the following, probability statement are made with respect to this measure.

Lemma A.2. *Let J_i be the set of individual i chooses to observe infinitely often and $\hat{J}_i$ the set of individuals nature selects for i to observe infinitely often. Then for any $k \in J_i \cup \hat{J}_i$ such that $\tau_{ik} < \max\{\tau_{ij} : j \in J_i \cup \hat{J}_i\}$, $P(j_{it} = k | J_i, \hat{J}_i, \boldsymbol{\tau}) \to 0$ as $t \longrightarrow +\infty$.*

Proof. Let $V_j(s_{it}, \eta_{it}, h_{it})$ denote the expected payoff from observing j given the individual's private signal, bias, and personal history. For any j observed infinitely often, the random variable $V_j(s_{it}, \eta_{it}, h_{it})$ converges almost surely to the constant

$$v(\tau_{ij}) \equiv -\frac{\sigma_0^2\sigma^2\left(\frac{\sigma_0^4\sigma^2(2\sigma_0^2+\sigma^2)}{(\sigma_0^2+\sigma^2)^3} + (1-\tau_{ij})^2\right)}{\left(\sigma_0^2+\sigma^2\right)\left(\frac{\sigma_0^2\sigma^2}{\sigma_0^2+\sigma^2} + \frac{\sigma_0^4\sigma^2(2\sigma_0^2+\sigma^2)}{(\sigma_0^2+\sigma^2)^3} + (1-\tau_{ij})^2\right)}$$

which is increasing in τ_{ij}. Hence, for any j and k observed infinitely often with

$\tau_{ik} < \tau_{ij}$

$$\lim_{t\to+\infty} P(V_j(s_{it},\eta_{it},h_{it}) > V_k(s_{it},\eta_{it},h_{it})|J_i,\hat{J}_i,\boldsymbol{\tau}) = 1$$

which implies $P(j_{it} = k|J_i,\hat{J}_i,\boldsymbol{\tau}) \to 0$ in the limit. □

For the following, let J_i^* be a nonempty set of individuals such that the probability i chooses to observe $j \in J_i^*$ does not vanish. Define the random variables $\tau_i^* = \min\{\tau_{ij} : j \in J_i^*\}$ and $\hat{\tau}_i = \max\{\tau_{ij} : j \in \hat{J}_i\}$.

Lemma A.3. *$P(\tau_i^* \leq \tau) \leq Q(\tau)$ for all $\tau \in [0,1]$.*

Proof. For any observation sets J_i^* and $\hat{J}_i$ such that $P(J_i^*,\hat{J}_i) > 0$, Lemma A.2 provides $P(\tau_i^* \geq \hat{\tau}_i|J_i^*,\hat{J}_i) = 1$. Consequently

$$P(\tau_i^* \leq \tau|J_i^*,\hat{J}_i) = P(\tau_i^* \leq \tau|\hat{\tau}_i \leq \tau, J_i^*,\hat{J}_i)P(\hat{\tau}_i \leq \tau|J_i,\hat{J}_i) \leq P(\hat{\tau}_i \leq \tau|J_i,\hat{J}_i). \tag{22}$$

Multiplying both sides by $P(J_i^*|\hat{J}_i)$ and then summing over all J_i^* finds $P(\tau_i^* \leq \tau|\hat{J}_i) \leq P(\hat{\tau}_i \leq \tau|\hat{J}_i)$. Since nature's selection for whom i observes is independent of similarity $P(\hat{\tau}_i \leq \tau|\hat{J}_i) = Q(\tau)^{|\hat{J}_i|} \leq Q(\tau)$. Multiplying both sides by $P(\hat{J}_i)$ and summing over all $\hat{J}_i$ completes the proof. □

Proof of Proposition 3. Denote $\bar{D}_{it} = \sum_{j\neq i} d_{ij}(t) = d_{ij_{it}} + \sum_{j\neq i,j_{it}} d_{ij}(t)$. Since marginals of the prior are symmetric, for all t

$$E[\bar{D}_{it}] = (n-1)\mathbb{E}[d_{ij}] \tag{23}$$

but also

$$\lim_{t\to+\infty} \mathbb{E}[\bar{D}_{it}] = E[d_{ij}|\tau_{ij} = \tau^*] + \lim_{t\to+\infty} \mathbb{E}\left[\sum_{j\neq i,j_{it}} d_{ij}(t)\right]. \tag{24}$$

Taking these two equalities together, along with Lemma A.3 and Corollary A.1 finds

$$\lim_{t\to+\infty} \mathbb{E}\left[\sum_{j\neq i,j_{it}} d_{ij}(t)\right] \geq (n-2)\mathbb{E}[d_{ij}] \geq (n-2)\mathbb{E}[d_{ij}|\tau_{ij} = \tau^*]. \tag{25}$$

By symmetry of the prior, $\mathbb{E}[d_{ij}|j_{i1} = j] = \mathbb{E}[d_{ij}]$. Hence, summing the inequalities over all i and dividing both sides by $n(n-2)$ yields (a) and (b). □

Proof of Proposition 4. For any period t, we can write $\mathbb{E}\left[(y_{it} - x_{it})^2\right] = \text{Var}(y_{it}) - \text{Var}(x_{it})$. Hence, it is sufficient to show that $\text{Var}(y_{it'}) \geq \text{Var}(y_{it})$ if $t < t'$.

By the law of total variance

$$\text{Var}(\theta_{it}) = \text{Var}\left(\mathbb{E}[\theta_{it}|h_{it+1}]\right) + \mathbb{E}\left[\text{Var}(\theta_{it}|h_{it+1})\right] \tag{26}$$

$$= \text{Var}\left(y_{it}\right) - \mathbb{E}\left[\mathbb{E}\left[u(y_{it};\theta_{it})|h_{it+1}\right]\right]. \tag{27}$$

As $\text{Var}(\theta_{it}) = \text{Var}(\theta_{it'})$ we see that, for any $t' \geq t$,

$$\text{Var}(y_{it'}) - \text{Var}(y_{it}) = \mathbb{E}\left[\mathbb{E}\left[u(y_{it'};\theta_{it'})|h_{it'+1}\right]\right] - \mathbb{E}\left[\mathbb{E}\left[u(y_{it};\theta_{it})|h_{it+1}\right]\right] \tag{28}$$

Let $y^{\dagger}_{it'}$ denote the action taken when the choice of whose opinion to observe and the action ignore information not contained in the smaller information set h_{it+1}. For every $h_{it'+1}$ and h_{it+1}

$$\mathbb{E}\left[u(y_{it'};\theta_{it'})|h_{it'+1}\right] > \mathbb{E}\left[u(y^{\dagger}_{it'};\theta_{it'})|h_{it'+1}\right] = \mathbb{E}\left[u(y^{\dagger}_{it'};\theta_{it'})|h_{it+1}\right] \tag{29}$$

with

$$\mathbb{E}\left[\mathbb{E}\left[u(y^{\dagger}_{it'};\theta_{it'})|h_{it+1}\right]\right] = \mathbb{E}\left[\mathbb{E}\left[u(y_{it'};\theta_{it'})|h_{it+1}\right]\right] \tag{30}$$

where the strict inequality in (29) is due to $y_{it} \neq y^{\dagger}_{it}$ almost surely. Taking together (28), (29), and (30) completes the proof. □