Echo Chambers: Disagreement and Polarization in Bayesian Learning

Cole Randall Williams*
University of California, Irvine

December 3, 2017

Abstract

Social learning and rational disagreement have been studied in environments in which agents are either homogeneous or the distribution of types is known. We study social learning under *unobserved heterogeneity*, where the distribution of types is unknown and is itself the subject of learning. This *dual learning* process unlocks a number of new results. Rational agents display confirmation bias. Learning is local: individuals place greater weight on opinions closer to their own and rationally discount more divergent views. Not only is there asymptotic disagreement, but social learning can polarize beliefs. Dual learning also provides a basis for social identification and group formation. We explore applications to political opinion formation, extremist behavior, and choice of news media.

The diversity of our opinions does not proceed from some men being more rational than others but solely from the fact that our thoughts pass through diverse channels and the same objects are not considered by all.

Descartes, *Discourse on Method.*

*Department of Economics, University of California, Irvine, 3151 Social Science Plaza, Irvine, CA 92697, colerw@uci.edu.
I would like to thank Jean-Paul Carvalho, Stergios Skaperdas, and John Duffy for invaluable advice and support. This paper has benefited from seminar audiences at Caltech, Chapman University, and the University of California, Irvine. All errors are mine.

1 Introduction

Belief formation occurs not merely through introspection but by observing the opinions of others. Social learning governs behavior in a number of domains, including consumption decisions, occupational choice, political preferences, and scientific beliefs. Yet our understanding of how we learn from others is still incomplete.

The problem hinges on precisely how one should respond to another person's opinion. The rational response is context-dependent. It depends on the underlying structure of the population, including the information available about its structure. When agents are *homogeneous*, differing only in their private information, individuals observing each others' opinions learn to agree (Aumann, 1976; Geanakoplos and Polemarchakis, 1982). When agents are *heterogeneous*, but each agent's type is known, individuals may fail to agree. In this paper, we show that when agents are heterogeneous and individual types are *unobservable*, individuals may not only fail to agree but also display rational forms of confirmation bias and other anomalous patterns of behavior.

Social learning in a context of unobserved heterogeneity becomes a process of *dual learning*: by observing the opinions of others an agent learns both about his parameter of interest and the structure of the heterogeneity in the population. Encountering an agent with a divergent opinion could now mean that this agent is importantly different from himself. As such, learning is *local*: individuals place greater weight on opinions closer to their own and rationally discount highly divergent views. This unlocks a number of new results, including non-monotonicities in belief formation, belief polarization, and social identification through social learning.

There are many real-world examples of dual learning processes. The following presents four such examples. The first two—restaurant choice and political opinion—return several times throughout the paper to illustrate our results.

Restaurant. Restaurant choice is a canonical example of social learning.[1] Under unobserved heterogeneity, a negative review of a restaurant can either mean that the restaurant is of low quality or that the reviewer has different preferences to oneself. The more positive experiences one has of the 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

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

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 underlying differences.

Technological Innovation. Empirically consistent patterns in innovation diffusion are often explained as the result of heterogeneity among an innovation's potential adopters.[2] A prime 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.

The dual learning process that we study, arising from unobserved heterogeneity, unlocks a number of new results. The following is a non-exhaustive summary.

Our first result establishes the model's primary mechanism: learning is *local* in the sense that individuals place more weight on opinions closer to their own. By increasing the difference in opinions, this weight can be made arbitrarily close to zero. We then ask how one responds to changes in another's opinion. We find that the answer depends on which of the two countervailing forces of dual learning dominates. This leads to *non-monotonicity in disagreement*, whereby, encountering someone with a slight difference in opinion can have a larger influence on one's beliefs than if they were to hold a starkly different opinion.

In the long run, after exchanging opinions with enough other individuals, one's own opinion will converge. At this point, we observe *social identification* through social learning: learning

[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). In economics there has been disagreement over topics like the efficacy of monetary policy in stimulating the real economy and the employment effects of raising the minimum wage.

an additional individual's opinion will serve almost entirely as a means for assessing the degree of similarity between the other individual and oneself.[4] Hence dual learning provides a basis for social identification through similarity of beliefs (Akerlof and Kranton, 2000, 2010).

We then consider extensions to the basic framework, beginning with introducing our model into the observational learning environment (Bikhchandani, Hirshleifer and Welch, 1992; Banerjee, 1992) in which agents learn more coarsely by observing the actions performed by others. We show that our characterization of heterogeneity has *competing welfare effects in observational learning*: (1) agents will never converge with certainty to their optimal action in the limit of learning and (2) the process can avoid falling into an information cascade when it would have done so with certainty under homogeneity.

Another extension characterizes the behavior of media consumers.[5] Consumers choose to acquire information from sources that confirm their own beliefs. Some come to place enough trust in a media source to rely on its reports in place seeking out their own information. In existing models, the media's effect on public disagreement requires that the population itself not be aware of the disagreement. Otherwise, the agents will condition on the disagreement and the media's effect will cease. In our model, the public's awareness of disagreement can strengthen the disagreement.

Section 2 provides the background in terms of probability and decision theory for our more general approach to social learning which takes into account unobserved heterogeneity. Section 3 introduces the model of dual learning. Section 4 presents the main results of the paper, starting with the most basic setting and gradually increasing in complexity. Section 5 considers extensions of the model and Section 6 concludes.

2 Related Literature

A discussion of the rational response to someone's beliefs must begin by specifying the form that rational belief will take and how it will respond to evidence. For this, we look to seminal figures in the development of the subjectivist (or personalistic) view of probability, Ramsey (1931), de Finetti (1937), and Savage (1954), who show that, if an individual satisfies certain

[4] For example, following many conversations about climate policy with various people, hearing an additional person's view may have a negligible effect on one's own opinion, but can be quite instructive about the similarity or differences in basic values.

[5] Gentzkow, Shapiro and Stone (2016) review the related literature on media bias.

basic coherence requirements, then their beliefs can be characterized by probabilities and they will update according to conditioning in response to new evidence. At this baseline level of rationality, there is no imperative that individuals come to agreement upon discovering that they hold conflicting views—it depends exclusively on how the other's beliefs fit into their respective models of the world.

> "The criteria incorporated in the personalistic view do not guarantee agreement on all questions among all honest and freely communicating people, even in principle."
>
> Savage, (1954)

In the development of classical game theory, these minimal coherence requirements were insufficient to provide general tractability for games with *incomplete information*, that is, games in which some players are uncertain about the game being played. In such games, a player's optimal action will depend on an infinite hierarchy of beliefs: their first-order beliefs about the game, second-order beliefs about the other players' beliefs, third-order beliefs about the other players' beliefs about their beliefs, and so on ad infinitum. Harsanyi (1967) proposed the powerful simplifying assumption that it be common knowledge that players' beliefs are *mutually consistent*: any discrepancies between the various players' beliefs are driven solely by differences in private information.[6]

As shown by Aumann (1976), a strong implication of the mutual consistency assumption is that rational individuals cannot publicly disagree. More precisely, Aumann shows that with mutual consistency, if individuals' beliefs about an event are common knowledge, then they will agree.[7]

But of course, public disagreement is pervasive. Roughly 63% of Americans are "absolutely certain" of the existence of God, while 9% do not even believe in God, 48% believe that global climate change is due to human activity while 31% believe the causes to be natural, and 15% believe that the collapse of the World Trade Center resulted from controlled demolition while 75% do not.[8]

An account of the manifest public disagreement requires a weakening of the mutual consistency assumption. One weakening of the assumption that can sustain disagreement is

[6] Our discussion highlights the fact that the mutual consistency assumption is stronger than the common prior assumption, though these are often treated as equivalent in the literature. Mutual consistency entails both a common prior and common knowledge of the information structure.

[7] See Geanakoplos and Polemarchakis (1982), Bacharach (1985), and Samet (1990) for extensions and Rubinstein and Wolinsky (1990) for a discussion of the *Agreeing to Disagree* results spawned by Aumann.

[8] (Pew, Religious Landscape Study, 2014),(Pew, The Politics of Climate, 2016),(Angus Reid, Public Opinion, 2010)

to abandon the coherent belief paradigm altogether. For example, disagreement could be driven by confirmation bias (Rabin and Schrag, 1999; Fryer, Harms and Jackson, 2015), motivated reasoning (Lord, Ross and Lepper, 1979; Kunda, 1990; Bénabou and Tiróle, 2016), bounded memory (Wilson, 2014), or rule-of-thumb belief updating procedures (DeGroot, 1974; DeMarzo, Vayanos and Zwiebel, 2003).

We could alternatively deviate from mutual consistency by allowing individuals to begin with heterogeneous prior beliefs.[9] In this case, classic results in Bayesian consistency (Doob, 1949) and the merging of opinions (Blackwell and Dubins, 1962; Kalai and Lehrer, 1994) guarantees that agreement is almost surely reached over time.

As has been noted in the literature, robust disagreement can emerge when we allow for heterogeneity beyond mere differences in prior beliefs. For example, heterogeneous interpretations of public signals can explain the presence of asset trading (Harris and Raviv, 1993; Kandel and Pearson, 1995; Acemoglu, Chernozhukov and Yildiz, 2016). In contrast, the No Trade Theorems of Milgrom and Stokey (1982) and Tirole (1982) predict that risk-averse traders with mutually consistent beliefs will not engage in trade.

Unobserved heterogeneous priors have been studied in the context of information aggregation. Sethi and Yildiz (2012) find that if agents have heterogeneous prior beliefs that are unobservable but correlated, then the information is fully aggregated through successive declarations of beliefs. Sethi and Yildiz (2016) show that when agents have unobservable heterogeneous prior beliefs, agents will come to favor observing the opinions of those with whom they have become most familiar.

Smith and Sorensen (2000) study the asymptotic beliefs and actions of a population comprised of heterogeneous types in the context of observational learning. The important difference between their model and ours is that agents in our model receive a signal of their own parameter of interest (e.g. utility from performing an action) and agents in their model receive signals of the state of the world which then determines their parameter of interest. This difference is the fundamental driver of our results and leads to distinct and interesting outcomes when applied to observational learning (see section 5.1).

[9]For a discussion of the rationale for using models with heterogeneous prior beliefs see Morris (1995) and for applications see Dixit and Weibull (2007), Van den Steen (2011), Glaeser and Sunstein (2013), and Benoit and Dubra (2014).

3 Dual Learning

In our model, agents are sorted into heterogeneous and unobservable *types*. Agents of the same type seek to learn the same parameter of interest. In the examples, this translates to agents of the same type having the same tastes, values, or dispositions towards evidence. Each agent receives an informative signal of his parameter of interest and observes the opinions of the other agents from which he performs *dual learning*: he learns about his parameter as well as the likelihood that other agents are of the same type.

A way to visualize dual learning is to consider a variant of the classic ball and urn model. The following illustrates the process of dual learning and also foreshadows our result of 'non-monotonicity in disagreement.'

3.1 A Tale of Two Urns

Imagine that before you is an urn containing 100 green and red balls. You are asked to guess the number of green balls in the urn and will be paid in accordance with how close your guess $\widehat{G}$ is to the actual number of green balls G.[10] You are permitted to draw a ball from the urn 10 times (with replacement) and your draws come up with 8 green and 2 red balls. Suppose there is also another participant who makes 30 draws which you observe prior to making your guess. The left panel in Figure 1 illustrates how your guess would change by observing that the other participant, whose first 10 draws were identical to yours (8 green and 2 red), continued drawing only red balls, ending with a total of 8 green and 22 red.[11] On the y-axis is your guess $\widehat{G}$ (the posterior expected number of green balls) and on the x-axis is the additional red balls drawn by the other participant.

Now imagine there is also a second urn containing 100 green and red balls and you are uncertain which of the two urns the other participant is drawing from. Firstly, this uncertainty will lead you to place less weight on the other's draws than your own when forming your guess. Secondly, you will engage in *dual learning*—the weight you place on the other's draws will be updated based on the similarity of their draws to your own. In our example, as the other participant continues drawing red balls, it becomes increasingly likely that they are drawing from a different urn than you. Hence, you begin to place less weight on their draws in forming your guess. The right panel illustrates how dual learning leads your guess

[10]For example, your payment could be given by the quadratic loss function $1 - (\frac{G}{100} - \frac{\widehat{G}}{100})^2$.

[11]The simulation assumes a uniform prior over the number of green balls in the urns, $\pi(G) = \frac{1}{101}$ for $G = 0, 1, 2, ..., 100$.

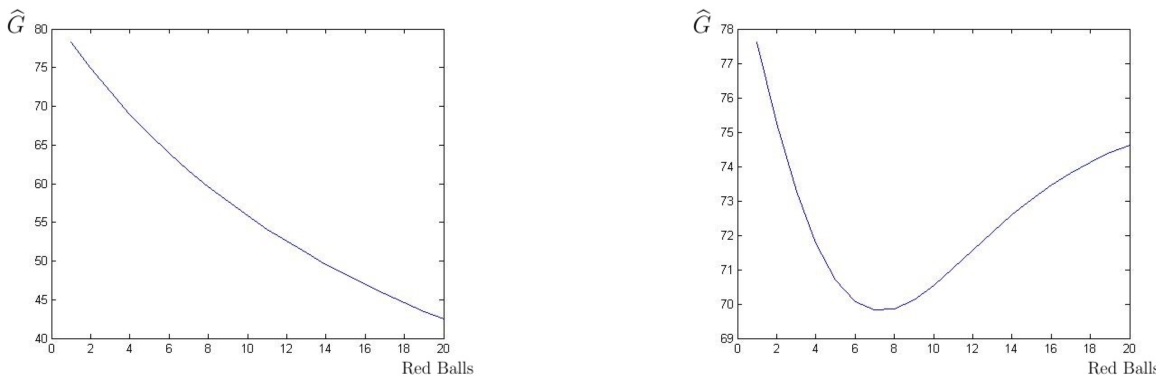

One Urn Two Urns

Figure 1: Both diagrams depict how your guess $\widehat{G}$ changes as the other participant draws additional red balls. In the left panel there is a single urn and in the right there are two urns. At the origin you and the other participant have each drawn 8 green and 2 red balls.

to move non-monotonically as we increase the disparity in the color frequency of the other participant's draws and your own.

3.2 A Model of Dual Learning

In a standard social learning model, a population of agents $i \in N$ seek to learn a single parameter $\theta^* \in \Theta$. Our model extends this to permit unobserved heterogeneity so that agents i and j seek to learn possibly distinct parameters θ_i^* and θ_j^*.

Nature first partitions the population $\bigcup_{t=1}^{T} N_t = N$ where agents i and j belonging to the same element of the partition N_t are said to be of the same *type*. The partition is formed randomly, with an agent being independently assigned to N_t with probability γ_t. We may assume the vector of assignment probabilities $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, ..., \gamma_T)$ to be either known or chosen by nature from a known distribution, but the realized partition is unknown.

Nature then independently assigns θ^t to each member of the partition N_t according to the known probability measure $\Pi(\cdot)$ with density $\pi(\theta)$. Agents of the same type seek to learn the same parameter, that is, for i and j in N_t, $\theta_i^* = \theta_j^* = \theta^t$.

Each agent receives a signal $s_i \in S$ in accordance with a conditional distribution with density $f_{\theta_i^*}(s)$. The family of conditional densities $f_\theta(s)$ are one-to-one,[12] continuous in θ and s, and the accompanying measures F_θ are mutually absolutely continuous.[13] An

[12]If $f_\theta(s) = f_{\theta'}(s)$ almost everywhere in S, then $\theta = \theta'$.

[13]Mutual absolute continuity provides that almost surely no single signal will perfectly reveal the distri-

agent then merges his signal with any of the other information he acquires (e.g. other agents' actions or opinions) to obtain his posterior probability measure $\Pi_i(\cdot|I_i)$ defined as $\Pi_i(A|I_i) \equiv Pr(\theta_i^* \in A|I_i)$ for measurable $A \subset \Theta$ with density $\pi_i(\theta_i^*|I_i)$. Observing other agents' actions or opinions can also be informative of whether they are of the same type as himself. Define $Q_{ij}(I_i) = Pr(j \text{ same type as } i|I_i)$ for $j \neq i$ to be the collection of i's posterior *perceived similarity.*

Each agent then uses his information to select an action x_i from the set X maximizing his expected payoff $E[u(x_i;\theta_i^*)|I_i]$. The particular form of the payoff function will be specified for each application.

As in Eyster and Rabin (2014), our results are most crisply articulated when agents have clarity in their inferences. For this reason, our primary analysis (section 4) allows agents to observe each others' *opinions* $\hat{\theta}(s_i)$. When parameters are real-valued $\Theta \subset \mathbb{R}$, we follow Sethi and Yildiz (2013, 2016) and specify an agent's opinion to be the expectation of his parameter given only his private information $\hat{\theta}(s_i) = E[\theta_i^*|s_i]$. More generally, we can think of an agent's opinion as some sufficient statistic of his private information. Qualitatively, our results remain under more coarse learning.

4 Interactive Belief Formation

This section analyzes the basic workings of the model. Beginning with a characterization of disagreement in the simplest setting of two agents and two actions, we move gradually toward increasing generality. Each step toward generality provides an additional insight into the role of dual learning in belief-formation and disagreement.

4.1 Two Agents & Two Actions

Our analysis begins by considering two agents 1 and 2 who are each faced with a choice between two actions $x_i \in \{0,1\}$. Agents are assigned one of two possible parameter values $\theta_i^* \in \{0,1\}$ whereby the payoff to choosing the action $x_i = \theta_i^*$ exceeds the payoff to choosing $x_i \neq \theta_i^*$. Each agent assigns a prior probability of π to $\theta_i^* = 1$ and a prior of $\tilde{\pi}$ to the other agent being of the same type. Each agent observes his signal $s_i \sim f_{\theta_i^*}$ and forms his opinion $\hat{\theta}(s_i) = E[\theta_i^*|s_i]$. In this setting, an agent's opinion coincides with what the literature calls

bution from which it is drawn.

an agent's *private belief*

$$\hat{\theta}(s_i) = p(s_i) = Pr(\theta_i^* = 1|s_i). \tag{1}$$

From this, we obtain a natural notion of disagreement whereby 1 and 2 *disagree* when $p(s_i) < \pi < p(s_j)$, that is, their opinions are pushed in different directions from the prior.

Each agent then observes the other's opinion. We take the perspective of agent 1 and simplify notation. Denote 1's posterior belief by $P(\mathbf{s}) = Pr(\theta_1^* = 1|s_1, p(s_2))$ and the posterior perceived similarity by $Q(\mathbf{s}) = Pr(\textit{2} \text{ same type as } \textit{1}|s_1, p(s_2))$ where $\mathbf{s} = (s_1, s_2)$.

As a benchmark for comparison, consider how agents update upon observing each other's opinions when they are certainly of the same type $\theta_1^* = \theta_2^* = \theta^*$. In this case, it is straightforward to see that both agents come to agreement on what we call the *shared opinion* $\hat{\theta}(s_1, s_2) = E[\theta^*|s_1, s_2]$ which simplifies

$$\hat{\theta}(s_1, s_2) = p(s_1, s_2) = Pr(\theta^* = 1|s_1, p(s_2)) \tag{2}$$

Unobserved heterogeneity adds an additional layer of complexity to the updating procedure. Fortunately, we can compute agent 1's posterior simply as the weighted average between his own opinion and the shared opinion where the weight is precisely the perceived similarity

$$P(\mathbf{s}) = p(s_1, s_2)Q(\mathbf{s}) + p(s_1)(1 - Q(\mathbf{s})). \tag{3}$$

By inspection of (3), when the perceived similarity is small, agent 1 mostly disregards 2's opinion $P(\mathbf{s}) \approx p(s_1)$. Conversely, when the perceived similarity is close to unity, 1's beliefs resemble those of the standard model $P(\mathbf{s}) \approx p(s_1, s_2)$.

The perceived similarity is itself revised upon observing opinions. It will be highest when the agents receive, in a sense, similar signals and shrink as their signals diverge. More precisely, following Bayes Rule we can write the perceived similarity

$$Q(\mathbf{s}) = \frac{f(s_1, s_2)\tilde{\pi}}{f(s_1, s_2)\tilde{\pi} + f(s_1)f(s_2)(1 - \tilde{\pi})} = \left[1 + \frac{f(s_1)}{f(s_1|s_2)}\frac{1 - \tilde{\pi}}{\tilde{\pi}}\right]^{-1} \tag{4}$$

where $f(\mathbf{s}) \equiv \int_\Theta f_\theta(\mathbf{s}) \mathrm{d}\Pi(\theta)$ is the marginal likelihood of receiving signals $\mathbf{s}$, assuming that they were drawn from the same distribution. As seen in (4), 1's perceived similarity increases upon observing the signals just in case $f(s_1|s_2) > f(s_1)$, that is, the likelihood of receiving s_1 from a distribution from which we already obtained s_2 exceeds the unconditional likelihood of having drawn s_1. With just two parameters, this inequality simplifies to $[f_1(s_1) - f_0(s_1)][f_1(s_2) - f_0(s_2)] > 0$.

The picture of interactive belief formation under dual learning is distinct from the standard model. When two agents share their opinions, their beliefs are drawn closer, but generically, we should not expect full agreement. Letting $P_i(\mathbf{s})$ represent i's posterior belief, the difference in posterior beliefs is proportional to the difference in opinions

$$P_i(\mathbf{s}) - P_j(\mathbf{s}) = \big(1 - Q(\mathbf{s})\big)\big(p(s_i) - p(s_j)\big). \tag{5}$$

Thus full agreement only occurs when agents hold equivalent private information $p(s_i) = p(s_j)$ or if they are certainly of the same type $Q(\mathbf{s}) = 1$.

The behavior predicted by the dual learning model can depart strongly from that predicted by the standard model. We shall see this difference in the following example.

4.1.1 Restaurant Example

Consider a new restaurant. If individual i chooses to dine there $x_i = 1$, he receives either high ($u_i = 1$) or low ($u_i = 0$) satisfaction. The payoff to any given visit is random and depends on an unknown parameter $\theta_i^* \in \{0, 1\}$, with $u_i(1; \theta_i^*) = \theta_i^*$ with probability 0.75. Thus i's expected payoff to dining at the new restaurant is 0.75 if $\theta_i^* = 1$ and 0.25 if $\theta_i^* = 0$. Specify i's payoff to not dining at the new restaurant $x_i = 0$ to be $u_i(0) = 0.4$ with certainty. Assume agents' prior beliefs are $\pi = \tilde{\pi} = \frac{1}{2}$.

This example resembles Bala and Goyal (1998) in that an agent's realized payoff operates as a signal. Before addressing the decision problem, let us first see how beliefs evolve if both agents repeatedly dine at the new restaurant and have opposed experiences.

Figure 2 illustrates the dynamics in agent 1's beliefs when both agents repeatedly dine at the restaurant and each time agent 1 receives high ($u_1 = 1$) satisfaction and 2 receives low ($u_2 = 0$) satisfaction. Along the horizontal axis, we increase the number of times each has dined at the restaurant. Under the standard model of learning, agents combine their opinions and form the shared opinion $p(s_1, s_2)$ which, due to their conflicting experiences, remains unchanged from the prior. Allowing for multiple types, the disparity in satisfaction provides increasingly strong evidence that agent 2 is of a distinct type to that of 1 and the perceived similarity vanishes $Q(\mathbf{s}) \to 0$. This observation, taken together with equation (3) implies that 1's posterior $P(\mathbf{s})$ quickly converges to his own opinion $p(s_1)$.

Consider how the agents' decisions are affected by their divergent experiences. For simplicity, suppose that an agent selects the dining option that maximizes his expected payoff, that

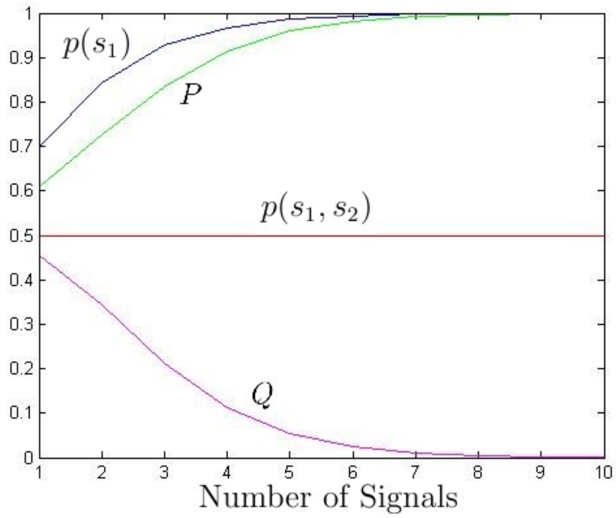

Figure 2: Belief Dynamics. The diagram depicts the change in 1's beliefs as we increase the number of high payoffs $s_1 = (1, 1, ...)$ received by 1 and low payoffs $s_2 = (0, 0, ...)$ received by 2.

is, each will dine at the new restaurant whenever it yields an expected payoff of at least 0.4.[14] The *ex ante* expected payoff to dining at the new restaurant is 0.5, and thus both agents choose this option. If the agents were homogeneous, then their persistent conflicting experiences would lead both to continue dining at the new restaurant.

After one visit to the new restaurant, before observing agent 1's opinion, agent 2's low satisfaction experience would reduce his expected payoff to dining at the new restaurant to 0.375. With no other information, agent 2 would not choose to dine there again. Upon learning that agent 1 had received high satisfaction from the new restaurant, agent 2 would revise his expected payoff to about 0.43 and would thus be willing to give the new restaurant another chance. After a second visit to the new restaurant brings agent 2 low satisfaction, he will not choose to dine there again, regardless of agent 1's satisfaction.

We can formalize our observation from the restaurant example that disparity between opinions leads one to place less weight on another's opinion in the following proposition.

Proposition 1. Say that agent 1 and 2 *agree* if $\pi < p(s_i) \leq p(s_j)$ or $p(s_i) \leq p(s_j) < \pi$ and *disagree* if $p(s_i) < \pi < p(s_j)$.

(a) If agent 1 and 2 *agree*, then the perceived similarity is strictly increasing as we increase the certainty of either of their opinions.

[14]Given our objective of illustrating opinion formation, we set aside the questions of optimal or strategic experimentation as studied in Bolton (1999) and section 5.4 of this paper.

(b) If agent 1 and 2 *disagree*, an increase in the difference of their opinions reduces the perceived similarity.

(c) Under maximal disagreement the perceived similarity vanishes and 1's posterior belief converges to his opinion: $p(s_i) \to 0$ and $p(s_j) \to 1$ imply $Q(\mathbf{s}) \to 0$ and thus $|P(\mathbf{s}) - p(s_1)| \to 0$.

Proofs of this and all further propositions can be found in the appendix.

4.2 Two Agents & Continuum of Actions

We now expand the action and parameter spaces $X = \Theta = \mathbb{R}$. In doing so, we show that increasing the disagreement between opinions can lead to more interesting, non-monotonic changes in actions.

As before, an agent observes his signal s_i and forms his opinion $\hat{\theta}(s_i) \equiv E[\theta_i^*|s_i]$. Assume θ_i^* to be normally distributed $\theta_i^* \sim \mathcal{N}(\theta_0, \sigma_0^2)$ and signals also normal distributed $s_i \sim \mathcal{N}(\theta_i^*, \sigma^2)$ so that an agent's opinion is a sufficient statistic of his information.

Each agent then observes the other's opinion, and updates his beliefs. If agents were homogeneous with certainty $\theta_1^* = \theta_2^* = \theta^*$, then they would come to agreement on the *shared opinion* $\hat{\theta}(\mathbf{s}) = E[\theta^*|\mathbf{s}]$, $\mathbf{s} = (s_1, s_2)$. For this section, payoffs are assumed to take the form

$$u(x_i; \theta_i^*) = -(x_i - \theta_i^*)^2. \tag{6}$$

Agent i's optimal action coincides with the posterior expectation of his parameter

$$x_i^* = E[\theta_i^*|s_i, \hat{\theta}(s_2)] = \hat{\theta}(\mathbf{s})Q(\mathbf{s}) + \hat{\theta}(s_i)(1 - Q(\mathbf{s})). \tag{7}$$

How does 1's action respond to a change in 2's opinion? The answer to this will depend on which of the two countervailing forces of dual learning dominates. Observe that s_2 enters (7) first through the shared opinion $\hat{\theta}(s_1, s_2)$ and second in the perceived similarity $Q(\mathbf{s})$. We can think of the movement in the shared opinion as the *direct effect* of shifting s_2. This effect captures the change in 1's beliefs if he takes 2's opinion at face value and does not consider the possible differences between them. Similarly, we can think of the adjustment of the perceived similarity as the *indirect effect* of shifting s_2. When 2's opinion $\hat{\theta}(s_2)$ is made increasingly dissimilar to 1's opinion $\hat{\theta}(s_1)$ the indirect effect counteracts the direct effect and the net result will depend on which of these two effects dominates. In the following, we continue the political opinion example from section 1 to show that the net change in 1's beliefs is not a foregone conclusion, but can vary on the domain of s_2.

4.2.1 Politics Example

Suppose the space of political policies can be described by the real line. Let us now think of s_i as i's interpretation of a piece of evidence based on his epistemic and normative values and θ_i^* as i's most preferred policy if he were to observe all the possible evidence. Specify i's interpretation of a piece of evidence as a noisy signal $s_i = \theta_i^* + \epsilon_i$ with ϵ_i distributed standard normal. Individuals who share the relevant underlying values will be receiving signals about the same preferred policy. Assume 1's prior over his preferred policy θ_1^* is standard normal and he receives evidence that suggests a policy of $0.7 = s_1$.

Figure 3 shows how 1's choice of policy $x_1^* = E[\theta_1^* | \mathbf{s}]$ changes as we alter 2's opinion. The direct effect of s_2 on the shared opinion $\hat{\theta}(s_1, s_2)$ is represented by the positive dashed line in the top sub-figure. The indirect effect of s_2 on the perceived similarity is given in the bottom sub-figure. Notice that the perceived similarity peaks near the point where 2's opinions are identical to 1's ($s_2 = 0.7$) and declines as 2's opinion moves in either direction. When 1 and 2's opinions are closest, the direct effect dominates and 1's action moves in a positive and roughly linear fashion. However, as 2's opinion moves further away from 1's in either direction, the indirect effect dominates and 1's action moves negatively with s_2. We call this pattern *non-monotonicity in disagreement*: when in close agreement 1 responds to changes in 2's opinion in a qualitatively similar way as the standard model, pushing 2's opinion too far leads 1 to respond in precisely the opposite manner of the standard model.

4.2.2 Moderating Extremists & Radicalizing Moderates

An important implication of non-monotonicity in disagreement relates to the processes of moderating extremist or radicalizing moderate behavior. In the previous example, imagine that higher actions are deemed more extreme and socially-undesirable. When confronted by a far less extreme agent 2 ($s_2 \leq -2$) figure 3 shows that 1's behavior will be almost entirely unchanged. If agent 2 were in fact more extreme ($s_2 \approx 0$), which may mean that 2 advocates performing some degree of socially-undesirable behavior, then 1 would reduce the extremity of his actions.

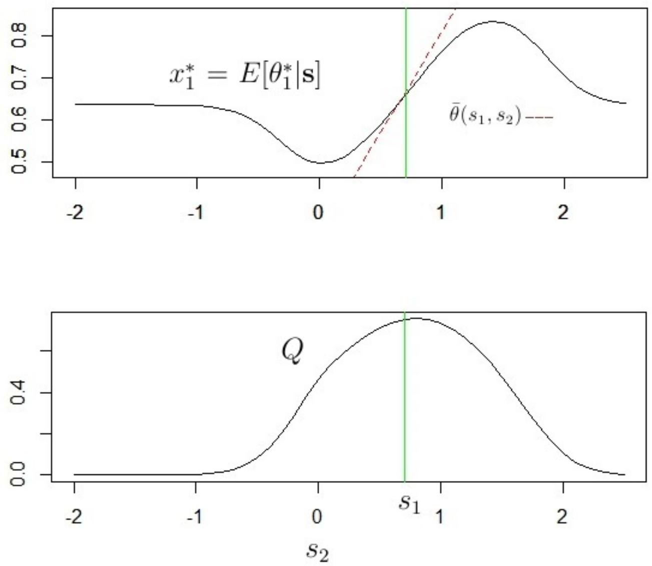

Figure 3: Non-Monotonicity in Disagreement. The diagrams illustrate the dependence of 1's beliefs and action on 2's signal.

The converse observation is made by supposing instead that agent 1 is already moderate in his behavior. There is more danger in him encountering a marginally more extreme individual ($s_2 \approx 1.5$) than someone who is vastly more extreme ($s_2 \geq 2.5$).

We now formally characterize the portions of the domain on which 1's action moves either positively or negatively to changes in 2's opinion. Consider as an analogue the identity *Revenue* = *Price* × *Quantity* from first principles. Revenue's response to a shift in the price can be positive or negative depending on the price elasticity of demand. Similarly, the response of 1's action to changes in 2's opinion will depend on the relative elasticity of the perceived similarity. To simplify notation, let $\Delta(\mathbf{s}) \equiv \hat{\theta}(s_1, s_2) - \hat{\theta}(s_1)$.

Definition 2. Define $\varepsilon \equiv -\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})} \Big/ \frac{\Delta(\mathbf{s}')-\Delta(\mathbf{s})}{\Delta(\mathbf{s}')+\Delta(\mathbf{s})}$ to be the elasticity of 1's perceived similarity $Q(\mathbf{s})$ which is said to be *relatively elastic* if $\varepsilon > 1$ and *relatively inelastic* if $\varepsilon < 1$.

The following proposition makes the regularity assumption that $\hat{\theta}(s_j) < \hat{\theta}(s_j')$ implies $\hat{\theta}(s_i, s_j) < \hat{\theta}(s_i, s_j')$. A similar (but less intuitive) statement could be made without use of this assumption.

Proposition 3 (Non-Monotonicity in Disagreement). Agent 1's action $x_1^* = E[\theta_1^*|\mathbf{s}]$ moves positively (negatively) with a change in 2's opinion $\hat{\theta}(s_2)$ if the perceived similarity $Q(\mathbf{s})$ is relatively inelastic (elastic).

4.3 Larger Finite Population ($n > 2$)

In this section, we will see that dual learning in a larger population produces different behavior than standard learning. In particular, we find *persuasion in numbers*: the opinions of the many outweigh the opinions of the few or one, even if both sets of opinions are equivalent in terms of information.

For example, consider customer product reviews. When many customers write reviews for a product, there is a good chance that some proportion of them will share the same preferences of the reader of these reviews. In contrast, when a single customer writes a review, the reader of this review cannot be sure if the customer's preferences match his own. Therefore, if many customers report satisfaction from single uses of a product it can more strongly influence the reader's purchasing decision than if a single customer were to report satisfaction from many uses of the product. In contrast, if it were known *ex ante* that everyone shared the same preferences, then both sets of reviews would influence the reader's beliefs identically. To illustrate this, we continue the restaurant example from section 4.1.1.

4.3.1 Restaurant Example Cont.

Suppose that agent 1 has not yet dined at the new restaurant and must solicit the opinions of his fellow agents prior to making his dining decision. We proceed by comparing three cases (1) agent 1 observes agent 2's repeated positive reviews (2) agent 1 observes the negative reviews from agents $3, 4, ..., n$, and (3) agent 1 observes both sets of reviews.

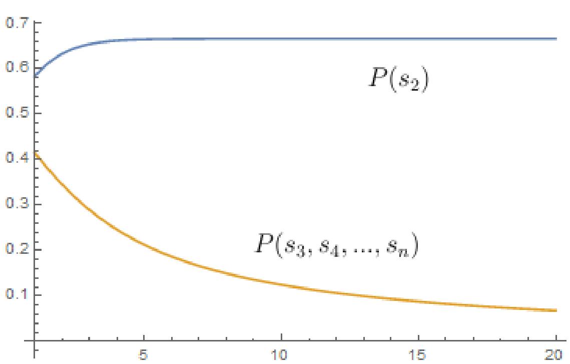

(a)

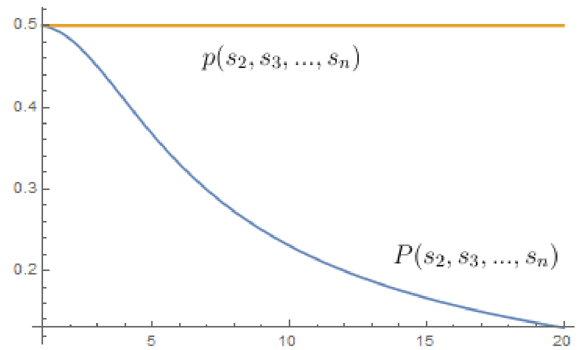

(b)

Figure 4: Belief Dynamics. The diagrams depict the change in agent 1's beliefs as we increase the number of high payoffs $s_2 = (1, 1, ...)$ received by agent 2 and low payoffs $(s_3, s_4, ...) = (0, 0, ...)$ received by agents 3, 4, ...

Suppose first that agent 1 observes that agent 2 repeatedly receives high satisfaction from dining at the restaurant. The plot of $P(s_2)$ in Figure 4.a shows how agent 1's beliefs update after each of 2's visits. Notice that, while 2's positive experiences increase 1's beliefs of his own self receiving a high expected payoff from the new restaurant, the effect tapers off. It becomes increasingly clear that 2's expected payoff from the restaurant is high, but there is no guarantee that he shares 1's tastes.

Second, suppose that agent 1 observes each agent $3, 4, ...n$ receive low satisfaction from dining at the new restaurant. The plot of $P(s_3, s_4, ..., s_n)$ in Figure 4.a reveals that these observations drive 1 to certainty that he will obtain a low expected payoff from the new restaurant. The effect is stronger in this case because there is a good chance that some fraction of these agents share the same tastes as agent 1.

Finally, we turn to Figure 4.b to see the effect of agent 1 observing agent 2 receiving ever more satisfying dining experiences and agents $3, 4, ..., n$ sequentially receiving low satisfaction experiences. The plot of the shared opinion $p(s_2, s_3, ..., s_n)$ demonstrates that, if the agents were homogeneous, then the conflicting experiences would lead 1's beliefs to remain unchanged from the prior. In contrast, the decline of $P(s_2, s_3, ..., s_n)$ reveals that the negative experiences of the many dominates the positive experiences of the one. Observing enough of these reviews will induce agent 1 to forgo dining at the new restaurant.

The following proposition describes *persuasion in numbers*. The phenomenon is most clearly identified in an environment in which X and Θ are finite, the prior over the assignment probabilities $\boldsymbol{\gamma}$ takes full support in the T-dimensional simplex, payoffs are finite, and $x(\theta) \equiv \arg\max_{x \in X} u(x; \theta)$ varies in θ.

We say that agent j is *certain* $x(\theta_j^*) = x$ if $Pr\big(x(\theta_j) = x | s_j\big) = 1$ and agent k is *boundedly certain* $x(\theta_k^*) = y$ if $\delta < Pr\big(x(\theta_k^*) = y | s_k\big) - Pr\big(x(\theta_k^*) = x | s_k\big) < 1$ for all $x \neq y$ and some $\delta > 0$.

Proposition 4 (Persuasion in Numbers). Suppose agent i's choice x_i^* is informed by the opinions of $n < +\infty$ other agents, whereby $n_x \geq 1$ of these agents are *certain* $x(\theta_j^*) = x$ and the remaining n_y agents are *boundedly certain* that $x(\theta_k^*) = y \neq x$.

(a) If the population is homogeneous, then $x_i^* = x$ for all $n_y < +\infty$.

(b) With a positive *ex ante* chance of heterogeneity and n_y sufficiently large, $x_i^* = y$.

4.4 Countably Infinite Population

What is the expected behavior of an agent's beliefs in an arbitrarily large population? Imagine that agent 1 observes his signal and the other agents' opinions in sequence. In a standard model, well-known results guarantee almost sure consistency of 1's posterior for the true θ^*. The heterogeneity in our framework precludes an immediate application of these results. To make progress on this question, we must introduce some further notation.

Let $\Omega = \Theta^T \times \Delta^T$ be the set containing the vectors of type parameters and assignment probabilities $\omega = (\theta^1, \theta^2, ..., \theta^T, \gamma_1, \gamma_2, ..., \gamma_T)$, where Θ^T is the T-fold product space ($\Theta \times \Theta \times ... \times \Theta$) and Δ^T represents the T-dimensional simplex. The population signal density belongs the finite mixture family $g_\omega(s) \equiv \sum_{t=1}^T \gamma_t f_{\theta^t}(s)$, $\omega \in \Omega$ and is said to be *identified* just in case $g_\omega(s) = g_{\omega'}(s)$ a.e. implies that both ω and ω' assign the same proportion of the population $\theta_i^* = \theta$ for all $\theta \in \Theta$ (Teicher, 1963).[15] An agent's opinion $\hat{\theta}(s_i)$ is a random quantity such that $s \mapsto \hat{\theta}(s)$ is one-to-one with realized opinions belonging to a complete separable metric space.

Define π^* to be the true distribution of thetas throughout the population as assigned by nature, i.e. $\pi^*(\theta)$ gives the proportion of the population with $\theta_i^* = \theta$. Let $\Theta^* \equiv \mathrm{supp}(\pi^*)$ be the finite support of π^*. Under the assumption of identifiability, as agent i continues observing the opinions of the other agents, his posterior almost surely converges to a function of only his own signal and the true distribution of parameters. Denote the vector containing the signals of the first n members of the population by $\mathbf{s}^n = (s_1, s_2, ..., s_n)$ and let "$\Rightarrow$" correspond to weak convergence.

Proposition 5 (Belief Convergence.). Suppose that Θ and S are complete separable metric spaces endowed with their respective Borel sigma algebras with $g_\omega(s)$, $\omega \in \Theta^T \times \Delta^T$ comprising an identified finite mixture family. Then for almost all ω^*, as $n \to +\infty$

$$\Pi_i(\cdot|\mathbf{s}^n) \Rightarrow \Pi_i(\cdot|s_i, \pi^*) \;\; a.s. \tag{8}$$

[15] Formally, $\sum_{t=1}^T \gamma_t \mathbf{1}(\theta^t = \theta) = \sum_{t=1}^T \gamma'_t \mathbf{1}(\theta^{t'} = \theta)$ for all $\theta \in \Theta$ where $\omega = (\theta^1, \theta^2, ..., \theta^T, \gamma_1, \gamma_2, ..., \gamma_T)$ and $\omega' = (\theta^{1'}, \theta^{2'}, ..., \theta^{T'}, \gamma'_1, \gamma'_2, ..., \gamma'_T)$. See Yakowitz and Spragins (1968) and Lindsay (1995) for further discussion of identified finite mixture models.

For continuity sets[16] $A \subset \Theta$, we can write i's asymptotic posterior explicitly as

$$\Pi_i(A|s_i, \pi^*) = \frac{\sum_{\theta \in \Theta^*} f_\theta(s_i)\pi^*(\theta)\delta_\theta(A)}{\sum_{\theta' \in \Theta^*} f_{\theta'}(s_i)\pi^*(\theta')} \tag{9}$$

where $\delta_\theta(A) = \mathbf{1}_\theta(A)$ is the Dirac measure assigning point mass at θ. If Θ is finite, we can write i's posterior probability mass function even more simply as

$$\pi_i(\theta|s_i, \pi^*) = \frac{f_\theta(s_i)\pi^*(\theta)}{\sum_{\theta' \in \Theta^*} f_{\theta'}(s_i)\pi^*(\theta')}. \tag{10}$$

These expressions have a nice interpretation. If agent i momentarily sets his own signal to the side and observes an infinite sequence of the other agents' opinions, he will learn the distribution of parameters throughout the population π^*. This distribution essentially becomes his new prior distribution over θ_i^* which he then updates by reintroducing his private signal.

It is worth noting that the belief convergence obtained in Proposition 5 does not guarantee that agents will likewise converge to their optimal action in the limit of learning. Section 5.1 discusses this in some detail.

In the limit of exchanging opinions, agents' beliefs converge. At this point, observing an additional agent's opinion serves almost entirely as an indicator of their underlying similarity. To examine the asymptotic behavior of i's perceived similarity, let us revisit the political opinion example where we left off in 4.2.1. We state this observation formally in the proposition that follows.

4.4.1 Politics Example Cont.

Imagine now that agent 1 continues conversing with other agents, learning their opinions about the optimal policy. After enough conversations, agent 1 will learn the distribution of opinions throughout the population. With that, the effect that each additional conversation has on his own view declines to zero. However, each conversation continues to be instructive for 1 to assess the similarity between the other agents and himself. In the long run, learning an individual's opinion functions almost entirely for social identification: serving as an indicator of the similarity in values between this individual and himself.

[16]Sets A for which the boundary has an asymptotically measure zero boundary $\Pi_i(\partial A|s_i, \pi^*) = 0$. In other words, none of the θ_i^* lie on the boundary of A.

Let ρ be the *Prokhorov metric* defined over the space of measures over Θ. For our purposes, it is sufficient to know that weak convergence corresponds to convergence in ρ. Details can be found in section 6 of Billingsley (2009).

Proposition 6 (Social Identification via Social Learning). Suppose agent i observes the other agents' opinions in sequence. As n goes to infinity, observing n's opinion has a vanishing effect on i's beliefs but a non-vanishing effect on i's perceived similarity of n,

$$\rho\big(\Pi_i(\cdot|\mathbf{s}^n), \Pi_i(\cdot|\mathbf{s}^{n-1})\big) \to 0 \tag{11}$$

$$d\big(Q_{i,n}(\mathbf{s}^n), Q_{i,n}(\mathbf{s}^{n-1})\big) \to w(s_i, s_n) \tag{12}$$

where $w(s_i, s_n)$ is not almost surely zero.

5 Extensions

Now that we have studied the basic workings of the model and identified the relationship between dual learning and disagreement, let us expand the discussion and ask the model what it has to say when other modifications and features are introduced. To summarize the findings: (1) dual learning can enhance learning while disallowing optimal action convergence in the observational learning environment, (2) agents can be more persuasive if they agree with each other on auxiliary topics, (3) over-representing those with extreme views can result in polarization when there would otherwise be none, and (4) we give a basic characterization of news media's consumer behavior.

5.1 Learning From Actions

Dual Learning delivers new insights to the observational learning literature. In the observational learning framework introduced by Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992), an individual chooses from a set of actions based on a private signal and information obtained through observing the actions taken by those who have chosen before. The principal finding has been the presence of *information cascades* whereby it becomes optimal for one (and hence all succeeding individuals) to follow the behavior of the individual who has chosen before oneself without regard for one's private information.

Let X be a finite set of actions, Θ a finite set of parameters, and $x(\theta_i^*)$ the optimal action for an agent with parameter θ_i^*. Agents choose their actions sequentially. When agent i

selects an action, he does so using the information contained in his private signal s_i and the actions of those who have selected before him. Further technical details are reserved for the appendix. Our discussion requires the following definitions:

1. Learning is *complete* if agents asymptotically assign probability one to the true distribution of parameters throughout the population $\pi^*(\theta)$.

2. A process exhibits *optimal action convergence*[17] if

$$\lim_{n\to\infty} Pr\big(x_n = x(\theta^*_n)\big) = 1. \tag{13}$$

3. An *information cascade* occurs whenever some agent's choice of action does not depend on his private signal.

4. A *confounding outcome* occurs when the population's limiting beliefs do not converge to certainty on $\pi^*(\theta)$ nor to a belief at which an information cascade would occur.

It will also be useful to define the following properties of the private beliefs that were introduced in 4.1. Letting $p_\theta(s_i) = Pr(\theta^*_i = \theta | s_i)$ be i's *private belief* in θ, we say that the signal structure has *unbounded private beliefs* if the support of p_θ contains 1 for all θ and *bounded private beliefs* if the support of p_θ does not contain 1 for any θ.

When is the process guaranteed to produce complete learning and when is there optimal action convergence? Smith and Sorensen (2000) study the case in which the population seeks to learn a single parameter θ^*. They show that if private beliefs are unbounded then either the process results in a confounding outcome or there is complete learning and actions converge to optimality. In contrast, there is never complete learning and thus no optimal action convergence when private beliefs are bounded. Acemoglu, Bimpikis and Ozdaglar (2010) extend the analysis to the case when agents only observe a subset of the history of actions. They find that with unbounded private beliefs and if agents' observations are sufficiently rich, then actions will converge in probability to the optimal action.

The introduction of unobserved heterogeneity has both a positive and a negative effect on asymptotic outcomes. Firstly, *unobserved heterogeneity completely disallows optimal action convergence.* To see why, observe that even if the population were to asymptotically learn the true distribution of thetas amongst the population $\pi^*(\theta)$, each individual i's private signal

[17] The literature often refers to almost sure convergence or convergence in probability to the optimal action as *asymptotic learning*. See Acemoglu, Bimpikis and Ozdaglar (2010), Acemoglu et al. (2011), and Mossel, Sly and Tamuz (2014).

is insufficient for him to deduce his own parameter of interest θ_i^*. For example, suppose that after observing many agents' purchasing decisions, the population learns that half of the agents in the population enjoy consuming some product while the remaining half do not. When the next agent is tasked with deciding whether or not to consume the product, he cannot be certain as to which half of the population he belongs.

Secondly, *unobserved heterogeneity can facilitate learning when homogeneity would preclude it.* In particular, when agents in the standard model have bounded private beliefs they cannot asymptotically assign certainty to the true θ^*. The reason is that as the population grows increasingly certain of the true θ^*, it will eventually be the case that an agent's action carries no information about his signal. In contrast with unobserved heterogeneity, when the population grows increasingly certain of the true distribution of parameters $\pi^*(\theta)$, there is always information about an agent's signal in his action. This finding is demonstrated by the example that follows in 5.1.1 adapted from Bikhchandani, Hirshleifer and Welch (1992).

Unobserved heterogeneity also facilitates learning when private beliefs are unbounded. When agents in the standard model have unbounded beliefs, either complete learning or a confounding outcome will obtain. In our model, each agent's action provides sufficiently rich information to prohibit the possibility of a confounding outcome and complete learning will always obtain.

5.1.1 Heterogeneity Facilitating Learning

Consider first the case which we know will result in an information cascade. A countable population of agents sequentially decide whether to *adopt* or *reject* some behavior. There is either a low or high value to adopting the behavior $\theta^* \in \{L, H\}$ and the value to rejecting it is 0. For simplicity, assume $L = -1$ and $H = 1$.

Each agent receives a privately observable signal $s_i \in \{L, H\}$, with $Pr(s_i = \theta^* | \theta^*) = r > 1/2$. When agent $n+1$ is asked to choose between adopting or rejecting the behavior, he computes his expected payoff using the information contained in his private signal as well as the actions of those who have chosen before $\mathbf{x}^n = (x_1, x_2, ..., x_n)$, $x_j \in \{\text{adopt}, \text{reject}\}$. The expected payoff to adopting is $2Pr(\theta^* = H | s_{n+1}, \mathbf{x}^n) - 1$. Thus $n+1$'s best response is to adopt just in case $Pr(\theta^* = H | s_{n+1}, \mathbf{x}^n) \geq \frac{1}{2}$. Using Bayes theorem, $n+1$ will select adopt if

$$
\begin{cases} r + Pr(\theta^* = H|\mathbf{x}^n) > 1 & s_{n+1} = H \\ \\ Pr(\theta^* = H|\mathbf{x}^n) > r & s_{n+1} = L \end{cases} \tag{14}
$$

The literature refers to $Pr(\theta^* = H|\mathbf{x}^n)$ as the *public belief*—the likelihood each agent $j > n$ assigns to the parameter being H after observing the first n actions of the other agents, but not his own private signal. Notice that the process will enter an information cascade if the public belief exceeds r as $n+1$ will choose 'adopt' regardless of his private signal. Similarly, a cascade ensues if the public belief falls below $1-r$ as $n+1$ will always choose 'reject'.

Must the process eventually enter into a cascade? There is a simple argument for why in fact it will eventually enter into a cascade with probability one. For a contradiction suppose that with positive probability the process does not at any point enter a cascade. For this to be true, the public belief could never have entered $[0, 1-r) \cup (r, 1]$. Being that the process never enters a cascade, we can infer the precise signal of each actor. Doob (1949) shows that observing an infinite sequence of IID draws will lead the public belief to almost surely converge to certainty on the true state and will thus converge to 0 or 1. This implies a contradiction as the almost sure convergence of the public belief will require it to have entered into $[0, 1-r) \cup (r, 1]$ and thus a cascade with probability one.

Now suppose that we introduce heterogeneous types into the model and define $z(\mathbf{x}^n) = Pr(\theta_j^* = H|\mathbf{x}^n)$ for $j > n$ to be the *public belief*. As before, extreme public beliefs $z(\mathbf{x}^n) \in [0, 1-r) \cup (r, 1]$ induce an information cascade. However, the process is no longer guaranteed to end up in a cascade! Notice now that if the process remains out of a cascade, the public belief will not converge to 0 or 1, but rather to the true proportion with $\theta_i^* = H$. In the appendix, we show that whenever this proportion lies within $(1-r, r)$, there is positive probability that the process never enters a cascade and complete learning will occur. The following proposition formalizes the discussion.

Proposition 7. In a population with unobserved heterogeneity:

(a) Optimal action convergence does not occur.

(b) There is generically complete learning with unbounded beliefs.

(c) Complete learning outcomes robustly exist with bounded beliefs.

5.2 Multiple Learning Problems

Up until now, we have maintained the assumption that agents form and share their opinions about a single topic. Realistically, there are many related topics we wish to learn about. The main lesson of this section is that the degree to which one can influence another's opinion is larger when they agree on auxiliary topics. Conversely, substantial disagreement on auxiliary topics can mitigate one's influence over another's opinion. Suppose that agent 1 has dined at many restaurants with agent 2 and they have shared largely the same quality of experiences each time. When 1 receives word that 2 holds a differing opinion about some new restaurant, he will be less swift in dismissing 2 as having distinct tastes. However, if 2 had a history of holding differing opinions about restaurants, he would hardly have placed any weight on 2's opinion of the new restaurant even if it had agreed with his own.

We reconsider the continuous action and parameter space of section 4.2. Agent i's payoff to an action $x_i^* \in \mathbb{R}$ takes of the form of quadratic loss from his parameter θ_i^* as in (6). He has previously engaged in L auxiliary learning problems, receiving private signals for parameters $\theta_i^\ell \in \Theta_\ell$ for $\ell = 1, 2, ..., L$. We assume that the type partition is constant between different learning problems $\theta_i^* = \theta_j^*$ and $\theta_j^\ell = \theta_i^\ell$ for all ℓ, though this could be weakened to positive correlation. We want to show how the similarity in beliefs over the L auxiliary issues θ_i^ℓ affects the susceptibility of i to be influenced in his action $x_i^* = E[\theta_i^*|\mathbf{s}]$. Assuming two agents, we can write 1's perceived similarity as

$$Q(\mathbf{s}) = \left(1 + \frac{1-\tilde{\pi}}{\tilde{\pi}} \cdot \frac{f(s_1)}{f(s_1|s_2)} \cdot \bar{R}^L\right)^{-1} \tag{15}$$

where $\bar{R} = \left(\prod_{\ell=1}^L \frac{f(s_{1\ell})}{f(s_{1\ell}|s_{2\ell})}\right)^{\frac{1}{L}}$ is the geometric mean of the likelihood ratios $\frac{f(s_{1\ell})}{f(s_{1\ell}|s_{2\ell})}$. Recall from section 4.1 that $f(s_{1\ell}|s_{2\ell}) > f(s_{1\ell})$ implies that the likelihood of observing $s_{1\ell}$ from a distribution from which we already obtained $s_{2\ell}$ exceeds the unconditional likelihood of having drawn $s_{1\ell}$. We can think of this geometric mean $\bar{R}$ as capturing the degree to which 1 and 2 agree on the L auxiliary problems. When $\bar{R} < 1$ the agents tend to agree and when $\bar{R} > 1$ the agents tend to disagree. Let "sufficient auxiliary agreement (disagreement)" denote "$\bar{R}^L$ sufficiently small (large)".

5.2.1 Politics Example (Multiple Policies)

The example illustrating non-monotonicity in disagreement in section 4.2 demonstrated that there was a necessary limit on the extent to which 2 could influence 1's view on a particular

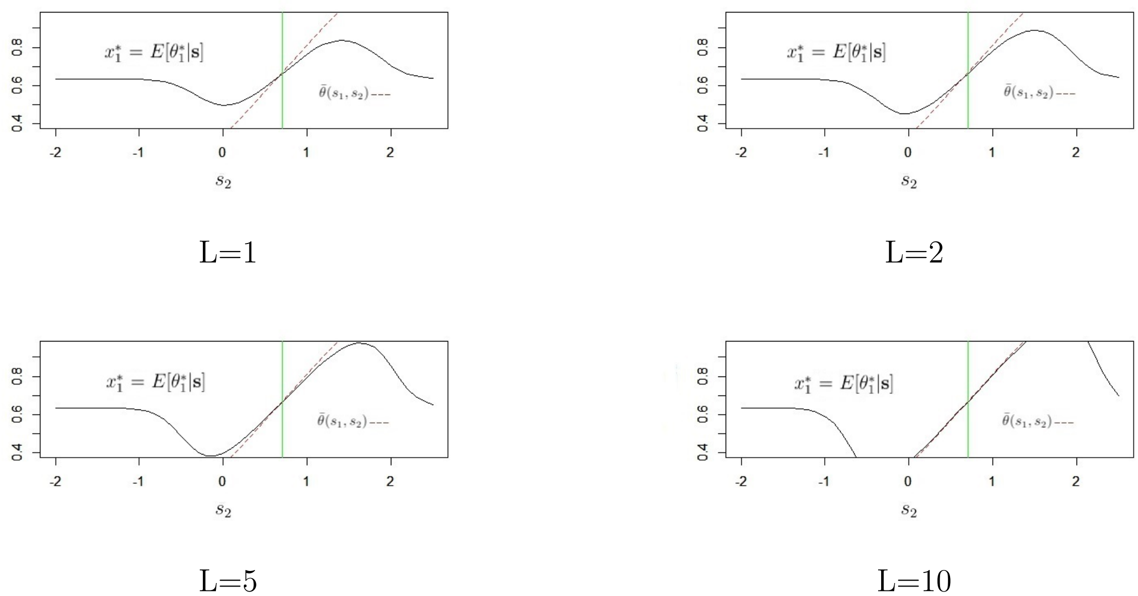

Figure 5: Multiple Learning Problems

political policy. Now suppose the two agents continue their conversation and 1 discovers that he shares much common ground with 2 on a large array of other political policies. This discovery will open up 1's belief about the original policy to being more susceptible to influence by 2. Figure 5 illustrates the persuasive power of agreeing on auxiliary issues. In each diagram, we fix $s_1 = 0.7$ and vary s_2 just as in figure 3. Between the diagrams we vary the number of other issues L on which the agents agree, where we specify the agreement as $\bar{R} = 0.75$.

There are a couple of different ways to express the idea that the degree to which 2 can influence 1's action is larger when they agree on auxiliary topics. Firstly, notice in figure 5 that by increasing L we expand the domain on which 1's action moves positively with 2's opinion. More generally, we can show that for every compact subset of the signal space $S' \subset S$, there is sufficient auxiliary agreement such that 1's action will move positively with 2's opinion for all s_1 and s_2 in S'.

Secondly, observe in figure 5 that increasing the auxiliary agreement raises the peaks and lowers the troughs of the $E[\theta_1^*|\mathbf{s}]$ curve. Generally, if the shared opinion $\hat{\theta}(s_1, s_2)$ is an unbounded function of s_2, then $\max_{s_2} E[\theta_1^*|\mathbf{s}]$ can be made arbitrarily large and $\min_{s_2} E[\theta_1^*|\mathbf{s}]$ arbitrarily small. Unbounded shared opinions can be found in the above example with normally distributed signals and a setting in which 2's signals consist of all vectors of finitely many draws from $f_{\theta_2^*}(s)$.

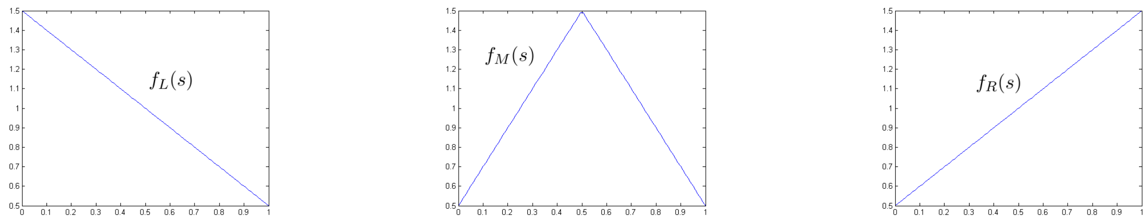

Figure 6: Conditional Densities

Consider a modification to the above example so that instead the agents disagree on the auxiliary topics $\bar{R} > 1$. This change would result in figure 5 showing the opposite qualitative effect of increasing L. In particular, an increase in L lowers the variation of x_1^* in 2's opinion and thus the influence that 2 can have on 1's beliefs.

Proposition 8.

(a) For every compact subset $S' \subset S$, there is sufficient auxiliary agreement such that 1's action x_1^* moves positively with a change in 2's opinion $\hat{\theta}(s_2)$ for all signals in S'.

(b) If the shared opinion is unbounded in s_2, then 1's action x_1^* can be made arbitrarily large or small given sufficient auxiliary agreement.

(c) The distance between x_1^* and his own opinion $\hat{\theta}(s_1)$ will be arbitrarily small under sufficient auxiliary disagreement.

5.3 (Perceived) Polarization

What happens if the media and social media skew their coverage in a way that over-represents those with more extreme views and this distortion is not accounted for by the population? We are going to look at an example of how dual learning can serve as a channel through which this type of distortion can lead to polarization where there would otherwise be none.

First, let us see why this distortion cannot drive polarization in the homogeneous case. Suppose all agents seek to learn $\theta^* \in \{L, M, R\}$ (left, moderate, right) and are faced with a set of actions $x_i \in \{L, M, R\}$ that yield a payoff of $u_i(x_i) = \mathbf{1}(x_i = \theta^*)$. Each agent receives $s_i \sim f_{\theta^*}$. The conditional densities are represented in figure 6, with $f_L(s) = \frac{3}{2} - s$ skewing signals left, $f_M(s) = \frac{1}{2} + 2s$ for $s < \frac{1}{2}$ and $\frac{5}{2} - 2s$ for $s \geq \frac{1}{2}$ giving moderate signals, and $f_R(s) = \frac{1}{2} + s$ skewing signals right.[18]

[18]These densities are used to simplify the exposition. That dual learning can serve as a conduit for generating polarization does not depend on the form of densities.

Each agent observes his own signal s_i and an infinite sequence of other agents' opinions. Consider the effect of a systematic distortion of the publicly observable opinions that over-represents extreme opinions (opinions assigning near certainty to L and R). In this case, the publicly observable opinions will outweigh each individual's private signal and the population will fully agree. Of course, the particular belief that the population settles on could be affected by the distortion, but there would nonetheless be agreement.

Now introduce the ex ante possibility of unobserved heterogeneity. Suppose that in actual fact $\theta_i^* = M$ for each and every agent i. Then observing infinitely many undistorted opinions will reveal this to be the case and the limiting behavior will be all agents selecting $x_i = M$ regardless of their signal.

As before, suppose there is a systematic distortion that over-represents extreme opinions. Then in the limit, the population might come to believe that there is true polarization. Furthermore, this very belief will drive polarized behavior.

Let $\hat{\pi} = (\hat{\pi}_L, \hat{\pi}_M, \hat{\pi}_R)$ be the limiting estimated distribution of thetas amongst the population. The expected payoff from action $x \in \{L, M, R\}$ for an agent with signal s_i is

$$U(x; s_i, \hat{\pi}) = Pr(\theta_i^* = x | s_i, \hat{\pi}) = \frac{f_x(s_i)\hat{\pi}_x}{f_L(s_i)\hat{\pi}_L + f_M(s_i)\hat{\pi}_M + f_R(s_i)\hat{\pi}_R}. \tag{16}$$

By inspection of (16), an agent's optimal action is the one maximizing the product $f_x(s_i)\hat{\pi}_x$. Figure 7 demonstrates that leading the population to believe that there are in fact fewer moderates (lowering $\hat{\pi}_M$) and more of the extremes (raising $\hat{\pi}_L$ and $\hat{\pi}_R$) will lead the population to increasingly segment between agents choosing $x_i = L$ and $x_i = R$. The top left sub-figure shows the undistorted case in which all agents optimally select $x_i = M$. With the small amount of distortion in the top right sub-figure, those on fringe with the most extreme signals are induced into choosing L or R. The bottom left shows the increase in distortion increases the proportion of the population choosing L or R. Finally, in the bottom right sub-figure, the distortion is sufficient to induce agents with any signal to choose L or R.

5.4 Application: News Media

How do individuals choose between news sources from which to acquire information? Gentzkow, Shapiro and Stone (2016) review the literature related to this topic. Mullainathan and Shleifer (2005) model consumers with a preference for reading news that confirms their own biases. We show that, with unobserved heterogeneity, consumers will rationally choose to

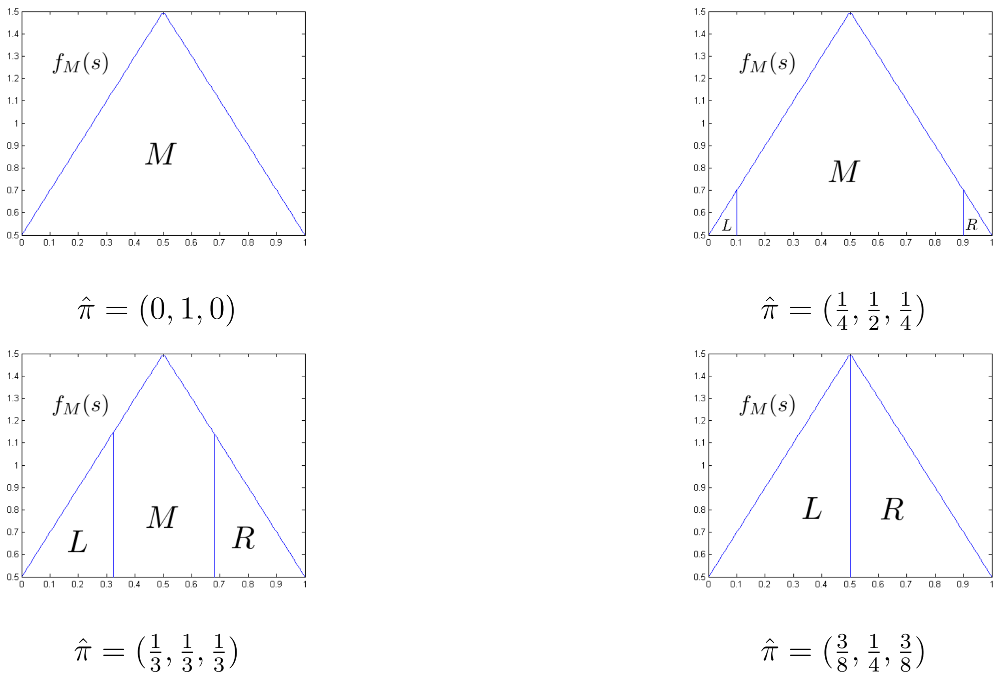

Figure 7: (Perceived) Polarization. The diagrams show that an increase in the perceived polarization $\hat{\pi}$ can drive polarized actions.

.

acquire news from sources that tend to confirm their own views—without an explicit confirmation bias. We also fill a gap in this literature by demonstrating that the media can facilitate public disagreement even when the public is aware of the disagreement.

5.4.1 Media Consumer Choice

Agents are engaged in a sequence of learning problems. In period $\ell = 1, 2, ...$ agent i selects action $x_i^\ell \in \mathbb{R}$ and receives payoff

$$u_i(x_i^\ell; \theta_i^\ell) = -\alpha_i(x_i^\ell - \theta_i^\ell)^2 \tag{17}$$

where $\theta_i^\ell \in \mathbb{R}$ and α_i measures i's idiosyncratic preference for holding accurate beliefs. Before selecting an action, i can choose to acquire a signal $s_i^\ell \sim f_{\theta_i^\ell}$ at a cost $c_i > 0$ and choose whether to observe the opinion $\hat{\theta}(s_M^\ell)$ of media firm M at opportunity cost d. We are interested in the case when acquiring direct information about an issue requires more effort than the time it takes to observe the media's report and hence we assume $c_i > d$. We also assume f_θ to be a normal density with mean θ and precision τ.

For the moment, suppose that there is a single media firm and consider the decisions faced by i in a given period ℓ. After acquiring the information that he wishes to obtain, his optimal action will be to select $x_i^\ell(\cdot) = E[\theta_i^\ell|\cdot]$ yielding an expected payoff of $-\alpha_i \text{Var}[\theta_i^\ell|\cdot]$.

When choosing the information to acquire, i must take into account both the benefits to the current period as well as the potential benefits to future periods. Assume i discounts the future at the rate $0 < \beta_i < 1$.

Proposition 9 (Media Consumer Choice). A consumer's optimal choice in period ℓ depends on α_i and $Q_{iM}^{\ell-1}$ as characterized by figure 8. Asymptotically, each α-type will select from either their left-most or right-most column.

Figure 8 charts out the optimal choice for i with cost c_i and discount rate β_i for different sensitivities to accuracy α_i and perceived similarity $Q_{iM}^{\ell-1}$ (written more clearly as Q). The symbol $\emptyset$ represents obtaining no signals, s_i obtaining i's own signal, s_M viewing the media's opinion, and s_i, s_M obtaining i's own and also view the media's opinion. In the "Experiment" region, a consumer will be willing to view both s_i^ℓ and $\hat{\theta}(s_M^\ell)$ for a period payoff that is lower than observing either only s_i or no signals $\emptyset$. The term "experiment" is drawn from the literature studying bandit processes. In the language of Gittins (1979), the consumers' decision problem is a bandit *superprocess*.

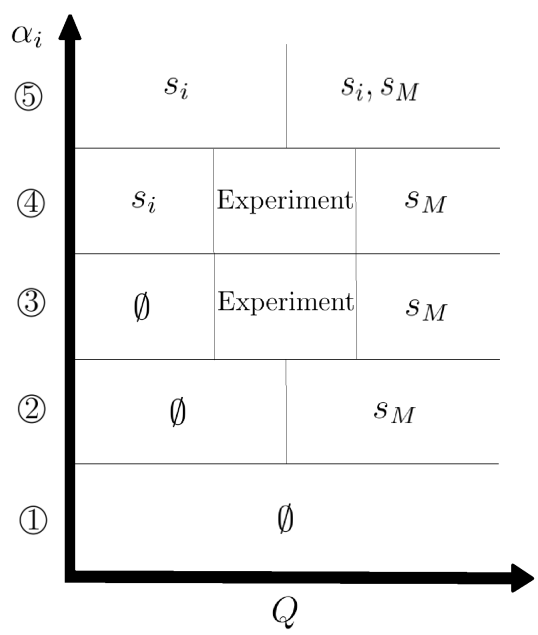

Figure 8: Optimal Choice. The diagram gives i's optimal information acquisition for various sensitivities to accuracy α_i and perceived similarity Q. Row ② vanishes if c_i is too small and there is no experimentation if d is too large.

Consider first α_i in a neighborhood of zero (row ① of figure 8). Such an agent would never deem the value high enough to purchase any information regardless of the perceived similarity. In row ②, which vanishes if the cost c_i is too small, the consumer is never willing to obtain his own signal and will only observe the media's opinion with a high enough perceived similarity.

Next, consider the other extreme of an agent with a very large α_i at level ⑤ who is quite intent on forming accurate beliefs. This agent would always find a benefit in obtaining s_i^ℓ and, so long as Q is not close to zero, will also view the media's opinion.

The behavior of agents with intermediate sensitivities to accuracy α_i is interesting. At ③, so long as the cost of observing the media's opinion d is not too large, there is a range of Q for which i will obtain a signal and observe that of the media for the sole purpose of experimentation. If the perceived similarity falls too low, the experimentation stops altogether and no information observed. Once Q is sufficiently high, the experimentation stops and the agent depends on the media's report.

At ③, middle values of the perceived similarity also involve experimentation if. If the perceived similarity falls too low, i will give up on the media and only trust his own information. As with ②, if Q becomes high, i will stop obtaining his own signal and depend only on that of the media.

Asymptotically, all α_i will at some point choose from either their respective rightmost or leftmost columns in figure 8. All types could learn to permanently ignore the media firm at some point when it has a track record of disagreeing too much with them. For α_i types ② through ④, they could find themselves in a situation where they choose to forever trust the information of this media firm.

5.4.2 Further Observations

We could conduct the same consumer analysis for the case with multiple media firms. This would show that some agents consume no media and others attend to media sources that do not have too low a perceived similarity. For those intermediate α_i types who rely on media without obtaining their own signals, there is a higher expected payoff when the consumed media sources tend to agree with each other than if they are discordant.

There are a couple more observations to make. For clarity, assume that there are two media firms A and B and the cost for a consumer to observe either of the firm's opinions is zero.

The first observation is that *a small amount of information can result in vast disagreement* *among the population.* An extreme example of this is found by supposing that the population has grown divided over the course of many periods with half the population assigning a perceived similarity of nearly one to firm A and zero to firm B and the remaining half assigning the reverse beliefs. Suppose further that no agent finds it optimal to pay c_i to purchase his own signal. Then the population will be sharply divided whenever the media firm's reports are distant from each other. This would be like Fox News and MSNBC presenting distinct opinions about some issue and viewers adopting the opinion of the news source that they have agreed with most in the past.

This example also illustrates the second observation that *public disagreement facilitated by the media does not dissipate with public awareness of the disagreement.* In existing models where diverging opinions are driven by media bias, it is important that the agents in the model are themselves unaware of this divergence. Otherwise, agents will simply condition on the opinion divergence and the media's effect vanishes. In our model, agents can grow to trust certain media sources more than others. In fact, learning that other media sources collide with one's trusted source is evidence against believing these other sources. Public disagreement is not crushed by observing the disagreement. Rather, observing public disagreement today lays the groundwork for even stronger disagreement tomorrow.

6 Conclusion

One of the unexpectedly useful insights gained from Harsanyi's mutual consistency assumption and Aumann's theorem is that we must look to differences between individuals beyond mere differences in their private information to understand the disagreement we observe in the world. This paper studies the emergence and patterns in disagreement when people take into account the unobservable differences between themselves when forming their beliefs. This more complex form of social learning, what we refer to as *dual learning*, captures many phenomena that do not fit existing models.

Disagreement is an important issue to understand. The opinions of a populace culminate in voting behavior that drives political decisions. The views of those tasked with determining research funding influence the very trajectory of science. It is our hope that this discussion and analysis will provide guidance for future empirical and theoretical work in uncovering the underlying differences that drive public disagreement.

References

Acemoglu, Daron, K. Bimpikis, and A. Ozdaglar. 2010. "Dynamics of information exchange in endogenous social networks." *Chemistry & biodiversity*, 1(11): 1829–1841.

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.

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz. 2016. "Fragility of Asymptotic Agreement Under Bayesian Learning." *Theoretical Economics*, 11: 187–225.

Akerlof, George A, and Rachel E Kranton. 2000. "Economics and Identity." *Quarterly Journal of Economics*, 115: 715–753.

Akerlof, George A, and Rachel Kranton. 2010. "Identity Economics." *The Economists' Voice*, 7(2): 11–13.

Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics*, 4(6): 1236–1239.

Bacharach, Michael. 1985. "Some extensions of a claim of Aumann in an axiomatic model of knowledge." *Journal of Economic Theory.*

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from neighbors." *Review of Economic Studies*, 65(3): 1–41.

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.*

Bénabou, Roland, and Jean Tiróle. 2016. "Mindful Economics: The Production, Consumption, and Value of Beliefs." *Journal of Economic Perspectives*, 30(3).

Benoit, Jean-Pierre, and Juan Dubra. 2014. "A Theory of Rational Attitude Polarization." *Working Paper.*

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades."

Billingsley, Patrick. 2009. *Convergence of Probability Measures.*

Blackwell, David. 1965. "Discounted Dynamic Programming." *The Annals of Mathematical Statistics*, 36(1): 226–235.

Blackwell, David, and Lester Dubins. 1962. "Merging of Opinions with Increasing Information." *The Annals of Mathematical Statistics.*

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica.*

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

de Finetti, Bruno. 1937. "La prévision: ses lois logiques, ses sources subjectives." *Annales de l'Institut H. Poincaré*, 1–68.

DeGroot, Morris H. 1974. "Reaching a Consensus." *Journal of the American Statistical Association*, 69(345): 118–121.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion bias, social influence and unidimensional opinions." *The Quarterly Journal of Economics.*

Dixit, Avinash K, and Jörgen W Weibull. 2007. "Political Polarization." *Proceedings of the National Academy of Sciences of the United States of America.*

Doob, Joseph L. 1949. "Application of the Theory of Martingales." *Le calcul des probabilites et ses applications*, 23–27.

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.*

Fryer, Roland G Jr., Philipp Harms, and Matthew O Jackson. 2015. "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization." *Working paper.*

Geanakoplos, John D., and Heraklis M. Polemarchakis. 1982. "We can't disagree forever." *Journal of Economic Theory*, 28(1): 192–200.

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

Ghosh, J.K., and R.V. Ramamoorthi. 2003. *Bayesian Nonparametrics.* New York:Springer.

Gittins, John C. 1979. “Bandit processes and dynamic allocation indices.” *Statistics*, 41(2): 148–177.

Glaeser, Edward L., and Cass R. Sunstein. 2013. “Why does balanced news produce unbalanced views?”

Harris, Milton, and Artur Raviv. 1993. “Differences of Opinion Make a Horse Race.” *Review of Financial Studies*, 6(3): 473–506.

Harsanyi, J. C. 1967. “Games with Incomplete Information Played by ”Bayesian” Players, I-III Part I. The Basic Model.” *Management Science*, 14(3): 159–182.

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.*

Kalai, Ehud, and Ehud Lehrer. 1994. “Weak and strong merging of opinions.” *Journal of Mathematical Economics*, 23(1): 73–86.

Kandel, Eugene, and Neil D. Pearson. 1995. “Differential Interpretation of Public Signals and Trade in Speculative Markets.”

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

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

Lindsay, Bruce G. 1995. “Mixture Models: Theory, Geometry, and Applications.” *NSF-CBMS Regional Conference Series in Probability and Statistics.*

Lord, Charles G, Lee Ross, and Mark R Lepper. 1979. “Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence.” *Journal of Personality and Social Psychology.*

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

Milgrom, Paul, and Nancy Stokey. 1982. "Information, trade and common knowledge." *Journal of Economic Theory*, 26(1): 17–27.

Morris, Stephen. 1995. "The Common Prior Assumption in Economic Theory." *Economics and Philosophy*, 11(02): 227.

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

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The market for news."

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

Rabin, Matthew, and Joel L Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." 1.

Ramsey, Frank P. 1931. "Truth and Probability." *The Foundations of Mathematics and Other Logical Essays*.

Rubinstein, Ariel, and Asher Wolinsky. 1990. "On the logic of "agreeing to disagree" type results." *Journal of Economic Theory*, 51(1): 184–193.

Samet, Dov. 1990. "Ignoring ignorance and agreeing to disagree." *Journal of Economic Theory*, 52(1): 190–207.

Savage, Leonard J. 1954. *The Foundations of Statistics.* John Wiley and Sons.

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. 2013. "Perspectives, Opinions, and Information Flows." *SSRN Electronic Journal*.

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.

Teicher, Henry. 1963. "Identifiability of Finite Mixtures." *The Annals of Mathematical Statistics*, 34(4): 1265–1269.

Tirole, Jean. 1982. "On the Possibility of Speculation under Rational Expectations." *Econometrica*.

Van den Steen, Eric. 2011. "Overconfidence by Bayesian-Rational Agents." *Management Science*, 57(5): 884–896.

Wilson, Andrea. 2014. "Bounded Memory and Biases in Information Processing." *Econometrica*, 82(6).

Yakowitz, S.J., and J.D. Spragins. 1968. "On the Identifiability of Finite Mixtures." *Annals of Mathematical Statistics*, 39: 209–214.

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 Mathematical Appendix

The precise statements of many of the theorems referenced in the following proofs can be found in the online appendix: `http://colerandallwilliams.com/wp-content/uploads/2017/08/Echo_Chambers_Resources.pdf`.

Proof of Proposition 1. Let $R_k \equiv \frac{f_L}{f_H}(s_k)$ be the likelihood ratio for drawing signal s_k. The assumption $p(s_i) < \pi_H < p(s_j)$ implies $R_i > 1 > R_j$. Expanding (4)

$$Q(\mathbf{s}) = \left[1 + \frac{\pi_H^2 + R_j\pi_H\pi_L + R_i\pi_H\pi_L + R_iR_j\pi_L^2}{\pi_H + R_iR_j\pi_L} \cdot \frac{1-\tilde{\pi}}{\tilde{\pi}}\right]^{-1}$$

we find Q to be differentiable in R_i and R_j. Upon differentiating, we find $\frac{dQ}{dR_i} < 0$ whenever $R_j < 1$ and $\frac{dQ}{dR_j} > 0$ for $R_i > 1$. Hence, Q is reduced by any increase in R_i or decrease in R_j, which corresponds to a decrease in $p(s_i)$ or increase in $p(s_i)$. Letting R_i^{-1} and R_j go to 0 sends Q to zero. From (3), $Q(\mathbf{s}) \to 0$ implies $P(\mathbf{s}) - p(s_1) = [p(s_1, s_2) - p(s_1)]Q(\mathbf{s}) \to 0$. ∎

The mechanism of proposition 1 can be found in more general environments. The private beliefs $p(s_i)$ of the two parameter environment provided an intuitive conceptualization of agreement and disagreement between 1 and 2. As we generalize, we now associate the agreement between 1 and 2 directly with the likelihood ratio $\frac{f(s_1)}{f(s_1|s_2)}$: lower values imply a higher degree of agreement. For measurable $A \subset \Theta$ we can write

$$\Pi_1(A|\mathbf{s}) = \Pi'(A|\mathbf{s})Q(\mathbf{s}) + \Pi'(A|s_1)\big(1 - Q(\mathbf{s})\big) \tag{18}$$

where $\Pi'(\cdot|\mathbf{s})$ is a probability measure that assumes all signals in the vector $\mathbf{s}$ come from agents of the same type. Let $\|\cdot\|$ denote the total variation metric between two probability measures, $\|\mu - \lambda\| \equiv 2\sup_B |\mu(B) - \lambda(B)|$ where the supremum is taken over measurable B.

Proposition 10. Consider a change in s_1 and/or s_2.

(a) If $\frac{f(s_1)}{f(s_1|s_2)} \to +\infty$ then $\|\Pi_i(\cdot|\mathbf{s}) - \Pi'(\cdot|s_1)\| \to 0$.

(b) If $\frac{f(s_1)}{f(s_1|s_2)} \to 0$ then $\|\Pi_i(\cdot|\mathbf{s}) - \Pi'(\cdot|\mathbf{s})\| \to 0$.

Proof.(a) For any $A \in \mathcal{B}(\Theta)$

$$|\Pi_1(A|\mathbf{s}) - \Pi'(A|s_i)| = |\Pi'(A|\mathbf{s}) - \Pi'(A|s_i)| \cdot Q(\mathbf{s}) \leq Q(\mathbf{s}) \tag{19}$$

and hence $2\sup_{A\in\mathcal{B}(\Theta)} |\Pi_1(A|\mathbf{s}) - \Pi'(A|s_i)| \leq 2 \cdot Q(\mathbf{s})$. It follows that $\frac{f(s_1)}{f(s_1|s_2)} \to +\infty$ implies $Q(\mathbf{s}) \to 0$ and thus $\|\Pi_i(\cdot|\mathbf{s}) - \Pi'(\cdot|s_1)\| \to 0$. Part (b) can be proved in a similar fashion. ■

Proof of Proposition 3. Consider a shift in s_2 to s_2' and define $\mathbf{s} = (s_1, s_2)$ and $\mathbf{s}' = (s_1, s_2')$. By definition 2, $1 - \epsilon \equiv 1 + \frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}\Big/\frac{\Delta(\mathbf{s}')-\Delta(\mathbf{s})}{\Delta(\mathbf{s}')+\Delta(\mathbf{s})}$ which rearranges to $(1-\epsilon)\cdot\big(\Delta(\mathbf{s}') - \Delta(\mathbf{s})\big)\cdot\big(Q(\mathbf{s}') + Q(\mathbf{s})\big) = 2\big(E[\theta_i^*|\mathbf{s}'] - E[\theta_i^*|\mathbf{s}]\big)$. By the regularity assumption $\operatorname{sgn}\big(\Delta(\mathbf{s}') - \Delta(\mathbf{s})\big) = \operatorname{sgn}\big(\hat\theta(s_2') - \hat\theta(s_2)\big)$ and thus

$$\operatorname{sgn}(1-\epsilon)\cdot\operatorname{sgn}\big(\hat\theta(s_2') - \hat\theta(s_2)\big) = \operatorname{sgn}\big(E[\theta_i^*|\mathbf{s}'] - E[\theta_i^*|\mathbf{s}]\big) \tag{20}$$

where $\operatorname{sgn}(\cdot)$ is the well-known signum function defined for any real number as $x = \operatorname{sgn}(x)\cdot|x|$.
■

Proof of Proposition 4. (a) Consider first the homogeneous case. If some agent j is certain that $Pr(x(\theta^*) = x|s_j) = 1$ then i too obtains this certainty and no opinions to the contrary will reduce it: $Pr(x(\theta^*) = x|s_j, \mathbf{s}) = 1$ for any $\mathbf{s}$ that is not perfectly revealing. (b) As payoffs are finite, if i assigns a high enough probability to the event $x(\theta_i^*) = y$, he will indeed choose $x_i^* = y$. Let ρ denote the probability with which nature assigns an agent to a type with parameter θ^y which is itself defined as the parameter at which $x(\theta^y) = y$. Let the probability with which nature assigns an agent to a type with parameter $\theta \neq \theta^y$ be denoted by $\eta_\theta(1-\rho)$ and $\boldsymbol{\eta}$ be the vector of length $|\Theta| - 1$ containing all such η_θ. Agent i's posterior can be written

$$Pr(x(\theta_i^*) = y|\mathbf{s}) = \int_{\mathcal{D}} Pr(x(\theta_i^*) = y|s_i, \rho, \boldsymbol{\eta})\mathrm{d}\mu(\rho, \boldsymbol{\eta}|\mathbf{s}_{-i}) \tag{21}$$

where $\mu(\cdot|\mathbf{s}_{-i})$ is the posterior probability measure over $(\rho, \boldsymbol{\eta})$ given the signals of the agents other than i and $\mathcal{D}$ is the $|\Theta|$-dimensional simplex. Expanding the integrand of (21)

$$Pr(x(\theta_i^*) = y|s_i, \rho, \boldsymbol{\eta}) = \frac{f_{\theta^y}(s_i)\rho}{f_{\theta^y}(s_i)\rho + (1-\rho)\sum_{\theta\neq\theta^y} f_\theta(s_i)\eta_\theta} \tag{22}$$

we see that there exists a cutoff such that, if $Pr(\rho > k^*|\mathbf{s}_{-i}) > 1-\epsilon$, then i will select $x_i^* = y$.

By the Lebesgue Decomposition Theorem, we can decompose $\mu = \mu_1 + \mu_2$ such that μ_1 is absolutely continuous with respect to the ($|\Theta|$-dimensional) Lebesgue measure $\mu_1 << \lambda$ and μ_2 is singular with respect to the Lebesgue measure $\mu_2 \perp \lambda$. Let $v \equiv \frac{d\mu_1}{d\lambda}$ be the Radon-Nikodym Derivative of μ_1 with respect to λ. Let $\rho < \rho' < 1$, $\mathbf{s}_j$ the signals of the n_x agents certain that $x(\theta_j^*) = x$, and $\mathbf{s}_k$ the remaining n_y signals, and write

$$\frac{v(\rho, \boldsymbol{\eta}|\mathbf{s}_j, \mathbf{s}_k)}{v(\rho', \boldsymbol{\eta}|\mathbf{s}_j, \mathbf{s}_k)} = \frac{f(\mathbf{s}_k|\rho, \boldsymbol{\eta})}{f(\mathbf{s}_k|\rho', \boldsymbol{\eta})} \cdot \frac{v(\rho, \boldsymbol{\eta}|\mathbf{s}_j)}{v(\rho', \boldsymbol{\eta}|\mathbf{s}_j)} \tag{23}$$

The density $v(\rho, \boldsymbol{\eta}|\mathbf{s}_j)$ is almost surely positive. Let us now write the ratio of likelihoods

$$\frac{f(\mathbf{s}_k|\rho, \boldsymbol{\eta})}{f(\mathbf{s}_k|\rho', \boldsymbol{\eta})} = \frac{\prod_{s\in\mathbf{s}_k}\left(f_{\theta^y}(s_k)\rho + (1-\rho)\sum_{\theta\neq\theta^y} f_\theta(s_k)\eta_\theta\right)}{\prod_{s\in\mathbf{s}_k}\left(f_{\theta^y}(s_k)\rho' + (1-\rho')\sum_{\theta\neq\theta^y} f_\theta(s_k)\eta_\theta\right)} \tag{24}$$

By assumption, $\frac{f_{\theta^y}(s_k)}{f_\theta(s_k)} > b > 1$ for all k some such b. With some algebra it can be shown that (24) is less than

$$\left(\frac{\rho b + 1 - \rho}{\rho' b + 1 - \rho'}\right)^{n_y} \tag{25}$$

which goes to zero as $n_y \to +\infty$. It follows that $v(\rho, \boldsymbol{\eta}|\mathbf{s}_{-i})$ goes to zero as $n_y \to +\infty$. Similarly, we could show that the posterior probability on any atoms goes to zero as $n_y \to$

$+\infty$. We can thus write

$$Pr(\rho < k^*) = \int_{\mathcal{D}|\rho<k^*} v(\rho, \boldsymbol{\eta}|\mathbf{s}_{-i}) \mathrm{d}\lambda(\rho, \boldsymbol{\eta}). \tag{26}$$

As the integrand of (26) is bounded and converges pointwise to 0, the Bounded Convergence Theorem provides that $Pr(\rho < k^*)$ likewise converges to 0 as $n_y \to \infty$. ∎

Remark. If the population were homogeneous $T = 1$, then Doob's Consistency Theorem (Doob, 1949) tells us that if i observes an infinite sequence of signals, with prior probability one the posterior will be consistent at the true parameter. One condition of Doob's Theorem is that the family of distributions is one-to-one. This assumption is not satisfied for the case of T component mixture models with component weights γ_t and parameterized distributions F_{θ^t}, $t = 1, 2, ..., T$. This is indeed the setting in which we are working.

In the following, we carefully define a function $h : \Omega \to \Omega$ that generates an equivalence class of ω's, in that $g_\omega(s) = g_{\tilde{\omega}}(s)$ a.e. implies that $h(\omega) = h(\tilde{\omega})$. The family $g_{\omega'}$ defined on the image $\omega' \in \Omega' \equiv h(\Omega)$ is one-to-one and hence we can apply Doob's Theorem.

By the Borel Isormorphism Theorem, there exists a Borel isormphism z between Θ and a subset of the interval $[0, 1]$ with the same cardinality as Θ. Without loss of generality assume $z^{-1}(0) \in \Theta$. Define the linear order on Θ to satisfy $\theta \leq \theta'$ iff $z(\theta) \leq z(\theta')$.

Definition. Let $h : \Omega \to \Omega$ with $h(\omega) = \omega' = (\theta^{1'}, \theta^{2'}, ..., \theta^{T'}, \gamma_{1'}, \gamma_{2'}, ..., \gamma_{T'})$ be defined by the following:

1. Combine duplicate θ's. Starting from left to right in ω, replace any of $\theta^{t+k} = \theta^t$ for some $k > 0$ with $z^{-1}(0)$ and add γ_{t+k} to γ_t while also replacing γ_{t+k} with 0.

2. Permute the indices so that the θ's are in ascending order. If there is $\theta^t = z^{-1}(0)$ with $\gamma_t > 0$ for some t, place it to the right (a higher index) to all the $\theta^\tau = z^{-1}(0)$ with $\gamma_\tau = 0$.

Lemma 11.

1. $h(\omega)$ is Borel Measurable.

2. $\Omega' \equiv h(\Omega)$ is a Borel subset of the complete separable metric space Ω.

Proof. (1) First decompose the domain $\Omega = \bigcup B_m$ whereby on each of the B_m the order of the θ^t does not change and if $\theta^t = \theta^{t'}$ or $\gamma_t = 0$ for some $\omega \in B_m$, then it does so for all $\omega' \in B_m$. Define the function $y(\omega) \equiv \big(z(\theta^1), z(\theta^2), ..., z(\theta^T), ..., \gamma_1, \gamma_2, ...\gamma_T\big)$. It can be shown that each $y(B_m)$ is a Borel subset of $[0,1]^{2T}$ and as y is Borel measurable $B_m \in \mathcal{B}_\Omega$.

By design, $h(\omega)$ is continuous on each B_m and the image of these subsets $h_m \equiv h(B_m)$ can be shown to be Borel. Take any $A \subset h(\Omega)$ such that $A \in \mathcal{B}_\Omega$ and write $A = \bigcup A_m$ where $A_m \equiv A \cap h_m$. We can write,

$$h^{-1}(A) = \bigcup (B_m \cap h^{-1}(A)) = \bigcup (B_m \cap h^{-1}(A_m)). \tag{27}$$

As h is continuous on each B_m, we know $h^{-1}(A_m)$ is contained in each sub-sigma algebra $\mathcal{B}_{B_m}$ and is thus also contained in $\mathcal{B}_\Omega$. It follows that $h^{-1}(A) \in \mathcal{B}_\Omega$.

(2) Follows immediately from $h(\Omega) = \bigcup h_m$ and the fact that each h_m is Borel. ■

Using Lemma 11, we can extend Doob's consistency theorem to the case of finite mixture models. The statement of the theorem writes "consistency*" to emphasize the use of a qualified notion of consistency. *Consistency* of the posterior at a point ω_0 entails that it will almost surely asymptotically assign probability 1 to every neighborhood of that point. For mixture models, *consistency** of the posterior at a point ω_0 entails that the posterior will almost surely assign probability 1 to every neighborhood of the *set of points equivalent to* ω_0. Here ω and ω' are *equivalent* if they assign the same weight to each θ, $\sum_{t=1}^{T} \gamma_t \mathbf{1}(\theta^t = \theta) = \sum_{t=1}^{T} \gamma'_t \mathbf{1}(\theta'^t = \theta)$.

Theorem 12 (Doob's Theorem for Finite Mixture Models.). Suppose that Θ and S are complete separable metric spaces endowed with their respective Borel sigma algebras with $g_\omega(s)$, $\omega \in \Theta^T \times \Delta^T$ comprising an identified finite mixture family. Let Π be a prior and $\{\Pi(\cdot|\mathbf{s}^n)\}$ a posterior. Then there exists $\Omega_0 \subset \Omega$ with $\Pi(\Omega_0) = 1$ such that $\{\Pi(\cdot|\mathbf{s}^n)\}_{n\geq 1}$ is consistent* at every $\omega \in \Omega_0$.

Theorem 12 is proved en route to the proof of proposition 5. Without loss of generality, the proof proceeds with i conditioning directly on the signals of agent $-i$. We could replace the s_j with j's opinion $\hat{\theta}_j = \hat{\theta}(s_j)$ and the proof would otherwise be unchanged.

Proof of Proposition 5. By lemma (11.1) $h(\omega)$ is Borel measurable so we can induce a measure λ on Ω' defined as $\lambda(A|\mathbf{s}^n) \equiv \widetilde{\Pi}(h^{-1}(A)|\mathbf{s}^n)$ where is $\widetilde{\Pi}(\cdot|\mathbf{s}^n)$ is the public belief over $\omega \in \Omega$ conditional on the vector $\mathbf{s}^n$, $\widetilde{\Pi}(B|\mathbf{s}^n) \equiv Pr(\omega \in B|\mathbf{s}^n)$ for $B \in \mathcal{B}_\Omega$.

First write

$$\Pi_i(A|\mathbf{s}^n) = \int_{\Omega'} Pr(\theta_i^* \in A|\mathbf{s}^n, \omega')\mathrm{d}\lambda(\omega'|\mathbf{s}^n) = \int_{\Omega'} \sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t} \mathbf{1}(\theta^t \in A)\mathrm{d}\lambda(\omega'|\mathbf{s}^n). \tag{28}$$

Claim 1: $\lambda(\cdot|\mathbf{s}^n) \Rightarrow \delta_{h(\omega^*)}(\cdot)$

The family $g_{\omega'}(s)$ is one-to-one on Ω' which by lemma (11.2) is a Borel subset of a complete separable metric space. Hence by Doob's Theorem (Doob, 1949) as stated in Ghosh and Ramamoorthi (2003), $\lambda(\cdot|\mathbf{s}^n)$ is almost surely consistent and by the Portmanteau Theorem paired with the fact that Ω' is a separable metric space, $\lambda(\cdot|\mathbf{s}^n)$ converges weakly to $\delta_{h(\omega^*)}(\cdot)$. Theorem 12 follows immediately by noting that if λ assigns probability 1 to every neighborhood of $h(\omega^*)$, then $\widetilde{\Pi}$ assigns probability 1 to every open set containing ω such that $h(\omega) = h(\omega^*)$.

The Portmanteu Theorem, equation (28), and claim 1 imply that whenever $\sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t}\mathbf{1}(\theta^t \in A)$ is almost surely continuous with respect to $\delta_{h(\omega^*)}(\cdot)$,

$$\Pi_i(A|\mathbf{s}^n) \to \int_{\Omega'} \sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t}\mathbf{1}(\theta^t \in A)\mathrm{d}\delta_{h(\omega^*)}(\omega). \tag{29}$$

Claim 2: For all A with $\Pi_i(\partial A|s_i, \pi^*) = 0$, $\sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t}\mathbf{1}(\theta^t \in A)$ is almost surely continuous with respect to $\delta_{h(\omega^*)}(\cdot)$.

The function $\sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t}\mathbf{1}(\theta^t \in A)$ is only discontinuous when some θ^t crosses the boundary of A, and is thus almost surely continuous with respect to $\delta_{h(\omega^*)}(\cdot)$ just in case $\delta_{h(\omega^*)}\big(D\big) = 0$ where $D \subset \Omega'$ is defined as the subset on which $\theta^t \in \partial A$ for some $\theta^t \in \omega \in D$. A set A satisfies this condition if and only if $\Pi_i(\partial A|s_i, \pi^*) = 0$ implying that such an A is a continuity set with respect to $\Pi_i(\cdot|s_i, \pi^*)$. As (29) holds for all continuity sets A, it follows by a final application of the Portmanteu Theorem that $\Pi_i(\cdot|\mathbf{s}^n)$ weakly converges to $\Pi_i(\cdot|s_i, \pi^*)$. ■

Proof of Corollary 5. Section 6 in Billingsley (2009) shows that weak convergence corresponds to convergence in ρ, hence (11) follows from proposition 5 and the triangle inequality. As $Q_{in}(\mathbf{s}^n) - Q_{in}(\mathbf{s}^{n-1}) =$

$$\int_{\Omega'} \big(Q_{in}(\mathbf{s}^n|\omega') - Q_{in}(\mathbf{s}^{n-1}|\omega')\big)\mathrm{d}\lambda(\omega'|\mathbf{s}^n) \tag{30}$$

has a continuous integrand in ω', the difference converges to $w(s_i, s_n) \equiv Q_{in}(\mathbf{s}^n|h(\omega^*)) - Q_{in}(\mathbf{s}^{n-1}|h(\omega^*))$. The second term in this difference is constant in s_n and $Q_{in}(\mathbf{s}^n|h(\omega^*))$ is not almost surely constant in s_n. ■

B Mathematical Appendix

Our discussion is made commensurate with the observational learning literature by assuming Θ to be finite, distinct θ and θ' prescribe different optimal choices from a finite set of actions X, and the vector of assignment probabilities $\boldsymbol{\gamma}$ to be known. Assume that each action is played for some open set of beliefs.

Proof of Proposition 7.

(a) Under heterogeneity, there is $\theta_i^*, \theta_j^* \in supp(\pi^*)$ with $x(\theta_i^*) \neq x(\theta_j^*)$. Optimal action convergence implies that with probability one such an agent i receives signal s_i such that they choose $x_i = x(\theta_i^*)$. Mutual absolute continuity of the signaling distributions would also imply j receives s_j inducing $x_j = x(\theta_i^*)$ with probability one. Hence, there is no optimal action convergence.

(b) This proof draws strongly from Smith and Sorensen (2000) (S&S). In this environment, nature chooses between only finitely many parameter vectors $\boldsymbol{\theta} = (\theta^1, \theta^2, ..., \theta^T)$ and hence only finitely many population distributions of thetas. Denote by $\tilde{\pi}(\theta) = \sum_{t=1}^{T} \gamma_t \mathbf{1}(\theta^t = \theta)$ a generic distribution of thetas and π^* the true distribution as chosen by nature. The likelihood ratios $\ell_{\tilde{\pi}}(\mathbf{x}^n) = \frac{Pr(\tilde{\pi}|\mathbf{x}^n)}{Pr(\pi^*|\mathbf{x}^n)}$ for $\tilde{\pi} \neq \pi^*$ and $\mathbf{x}^n = (x_1, x_2, ..., x_n)$ the first n actions chosen form a Martingale conditional on π^*. Define $\psi(x|\tilde{\pi}, \boldsymbol{\ell})$ to be the ex ante probability of an agent performing action x conditional on $\tilde{\pi}$ being the true distribution of parameters and $\boldsymbol{\ell}$ being their prior vector of likelihood ratios $\ell_{\tilde{\pi}}$.

By the Martingale Convergence Theorem, there exists a real, nonnegative stochastic variable $\ell_{\tilde{\pi}}^{\infty}$ such that $\ell_{\tilde{\pi}}(\mathbf{x}^n) \to \ell_{\tilde{\pi}}^{\infty}$ almost surely. This implies that asymptotically for all $\tilde{\pi}$ and all actions x played with positive probability

$$\ell_{\tilde{\pi}}^{\infty} = \frac{\psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty})}{\psi(x|\pi^*, \boldsymbol{\ell}^{\infty})} \ell_{\tilde{\pi}}^{\infty} \tag{31}$$

If only one action is taken with positive probability at $\boldsymbol{\ell}^{\infty}$, then because private beliefs are unbounded, the public belief must assign certainty to the population being homogeneous. As the likelihood ratios almost surely will not converge to certainty on the false $\tilde{\pi}$, then the population truly is homogeneous π^* and complete learning has occurred.

Consider the case where two actions x and x' are active in the limit. This would imply that at $\boldsymbol{\ell}^{\infty}$ agents assign positive prior probability to exactly two parameters θ and θ'. If only one parameter θ were assigned positive probability only the action $x(\theta)$ would be active. If more than two parameters were assigned positive probability, then unbounded private beliefs would entail that more than two actions would be active. For equation (31) to be satisfied, either $\ell_{\tilde{\pi}}^{\infty} = 0$ or $\psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty}) = \psi(x|\pi^*, \boldsymbol{\ell}^{\infty})$ and we shall proceed to show that the latter equality cannot hold for $\tilde{\pi} \neq \pi^*$.

For a given $\boldsymbol{\ell}$, i's best response will be x just in case her private belief $p_\theta(s_i)$ exceeds some threshold $K(\boldsymbol{\ell})$. Define $\widetilde{F}_\theta(p_\theta)$ and $\widetilde{F}_{\theta'}(p_\theta)$ to be the conditional distributions of the perceived similarity p_θ. The ex ante probability that i chooses x for a given $\tilde{\pi}$ is

$$\psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty}) = \tilde{\pi}(\theta)\big(1 - \widetilde{F}_\theta(K(\boldsymbol{\ell}))\big) + \tilde{\pi}(\theta')\big(1 - \widetilde{F}_{\theta'}(K(\boldsymbol{\ell}))\big). \tag{32}$$

From lemma A.1 in S&S $F_\theta(p_\theta) > F_{\theta'}(p_\theta)$ whenever both are not zero or one. Thus for distinct $\tilde{\pi}$ and π^*, almost surely $\psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty}) \neq \psi(x|\pi^*, \boldsymbol{\ell}^{\infty})$. It follows that $\ell_{\tilde{\pi}}^{\infty} = 0$ and complete learning has occurred.

Generically, more than two actions cannot be active at $\boldsymbol{\ell}^{\infty}$. To see this, let J be the number of $\tilde{\pi} \neq \pi^*$ with $\ell_{\tilde{\pi}}^{\infty} > 0$ and M the number of actions active at $\boldsymbol{\ell}^{\infty}$. Because of the identity $\sum_{x \in X} \psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty}) = 1$, if the equality $\psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty}) = \psi(x|\pi^*, \boldsymbol{\ell}^{\infty})$ holds for $M-1$ of the active actions, it must also hold for the remaining active action. Hence satisfying equation (31) for all $\tilde{\pi}$ and active actions generates a system of $J(M-1)$ equations in J unknowns $\ell_{\tilde{\pi}}^{\infty}$. nown 1. As the equations generically differ, they can only be solved when $M = 2$.

(c) Assume $\pi^*(H) \in (1-r, r)$. The idea is to first isolate a positive measure of trajectories that after M_ϵ steps will never leave a small radius around $\pi^*(H)$. Then we show that, of these trajectories, a positive proportion will have never left $(1-r, r)$ in the first M_ϵ steps.

Let $z(\mathbf{s}^n) = Pr(\theta_j^* = H|\mathbf{s}^n)$ for $j > n$ be a variant of the public belief $z(\mathbf{x}^n)$ defined in 5.1.1. Prior to a cascade $z(\mathbf{s}^n) = z(\mathbf{x}^n)$. As per theorem 12, $z(\mathbf{s}^n) \to \pi^*(H)$ almost surely as $n \to +\infty$. Let Z be the set of the trajectories of the public belief with Borel sigma algebra $\mathcal{B}_Z$ and probability measure μ induced from the signaling distribution G_{ω^*}. Define $Z' \subset Z$ to be the subset of trajectories such that $z_m \in (1-r, r)$ and at all points in the sequence. We want to show that $\mu(Z') > 0$.

Let $\delta \equiv \frac{1}{2}\min(|r - \pi^*(H)|, |\pi^*(H) + 1 - r|)$. For almost all trajectories, there exists an $M < +\infty$ such that, for all $m > M$, $|z_m - \pi^*(H)| < \delta$. Let M_ϵ be the smallest integer such that $\mu(Z_\epsilon) > \epsilon$ where $Z_\epsilon = \{z \in Z : \forall m \geq M_\epsilon, |z_m - \pi^*(H)| < \delta\}$.

As each trajectory has only received finitely many signals there are only finitely many unique of signal frequencies observed in the first M_ϵ steps of each trajectory. Partition $Z_\epsilon = \bigcup Z_\epsilon^l$

where each Z_ϵ^l comprises of trajectories which received the same signal frequencies in the first M_ϵ steps. The positive measure for Z_ϵ requires that at least one member of its partition has positive measure and thus suppose $\mu(Z_\epsilon^l) = \epsilon_1 > 0$. Denote by k the number of H signals in the first M_ϵ steps for trajectories in Z_ϵ^l.

Every permutation of k "H" signals and $M_\epsilon - k$ "L" signals has the same positive probability denoted by $\epsilon_2 > 0$. Thus we can find a subset of trajectories $\hat{Z}_\epsilon^l \subset Z_\epsilon^l$ that never leave $(1-r, r)$ defined by the permutation in which the first $2 * \min\{k, M_\epsilon - k\}$ steps form an oscillating sequence of $L, H, L, H, ...$ and then including whatever signals remain at the end. At no point during the oscillation does $z \in \hat{Z}_\epsilon^l$ leave $(1-r, r)$ and if it leaves after then necessarily it's M_ϵth entry z_{M_ϵ} would too, contradicting $\hat{Z}_\epsilon^l \subset Z_\epsilon$. It follows that $\mu(\hat{Z}_\epsilon^l) = \epsilon_1\epsilon_2 > 0$. Notice that $\hat{Z}_\epsilon^l \subset Z'$ as all $z \in \hat{Z}_\epsilon^l$ never leave $(1-r, r)$.

Thus we know $\mu(Z') \geq \mu(\hat{Z}_\epsilon^l) \geq \epsilon_1\epsilon_2 > 0$ ■

In the low probability event of avoiding an information cascade each agent follows the action dictated by his private signal. In the limit, the proportion choosing the correct action is r.

Proof of Proposition 8. (a) By proposition 3, it will be enough to show that for every compact $S' \subset S$, $\bar{R}^L$ can be made sufficiently small so that 1's perceived similarity is relatively inelastic $\epsilon < 1$ for all signals in S'. Because S' is compact and the conditional densities $f_\theta(s)$ continuous in s, both $\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ and $\max_{s_1,s_2,s_2' \in S'} \frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ are well defined on S'. It can be shown that for any given signals, $\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ goes to zero as $\bar{R}^L$ goes to zero. It follows that $\max_{s_1,s_2,s_2' \in S'} \frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ also goes to zero as $\bar{R}^L$ goes to zero. From definition 2, it follows that for $\bar{R}^L$ sufficiently small, $\epsilon < 1$ on S'.

(b) We will show that for sufficient auxiliary agreement $E[\theta_1^*|\mathbf{s}]$ can be made larger than any $M < +\infty$. Having assumed the shared estimate to be an unbounded function of s_2, we can find an s_2 such that $\hat{\theta}(s_1, s_2) > M$ for any M and s_1. From

$$E[\theta_1^*|\mathbf{s}] - \hat{\theta}(s_1, s_2) = \big(1 - Q(\mathbf{s})\big)\big(\hat{\theta}(s_1) - \hat{\theta}(\mathbf{s})\big) \tag{33}$$

and equation 15, $E[\theta_1^*|\mathbf{s}] - \hat{\theta}(s_1, s_2) \to 0$ as $\bar{R}^L \to 0$, completing the proof. The proof that $E[\theta_1^*|\mathbf{s}]$ can be made smaller than any $m > -\infty$ and for part (c) follow by similar arguments. ■

C Mathematical Appendix

Blackwell (1965) has shown that with a fixed and finite set of choices in each period, there is a deterministic stationary Markov policy for which, for any initial state, the total expected reward is the supremum of the total expected rewards for the class of all policies. The optimal policy satisfies the functional equation

$$V(Q) = \max_{\hat{s} \in \hat{S}} U(\hat{s}; Q) + \beta E\big[V(Q')|\hat{s}, Q\big] \tag{34}$$

with $\hat{S} = \{\emptyset, \hat{s}_i, \hat{s}_M, (\hat{s}_i, \hat{s}_M)\}$ the set of signal combinations the consumer can choose to observe. Define $A(Q)$, $B(Q)$, C, and D to be the expected period payoffs to observing (s_i, s_M), s_M, s_i, and $\emptyset$ (no signals) respectively for $Q^{\ell}_{im} = Q$ and a given α_i. Let $\bar{A} \equiv A(1)$ and $\bar{B} \equiv B(1)$.

Lemma 13.

1. The expected period payoff from observing s^{ℓ}_M is increasing in Q.
2. The value function $V(Q)$ is non-decreasing in Q. If at Q, there is positive probability of ever observing s^{ℓ}_M, then $V(Q)$ is increasing in Q.
3. There are diminishing expected period returns to information at $Q = 1$.
4. $A(Q)$ and $B(Q)$ are continuous in Q.
5. For a given c_i and β_i, we can partition the domain for $\alpha_i \in \mathbb{R}_+ = [0, b_0] \cup (a_1, b_1] \cup (a_2, b_2] \cup (a_3, +\infty)$ where $a_m < a_{m'}$ and $b_m < b_{m'}$ whenever $m < m'$. The following inequalities hold for the various α_i-types.

$$\begin{cases} \bar{B} - d \leq D,\ C - c_i < D,\ \bar{A} - c_i - d < D, & \alpha_i \in [0, b_0] \\ \bar{B} - d > D,\ C - c_i \leq D,\ \bar{A} - c_i - d < D, & \alpha_i \in (a_1, b_1] \\ \bar{B} - d > D,\ C - c_i > D,\ \bar{A} - c_i - d \leq D, & \alpha_i \in (a_2, b_2] \\ \bar{B} - d > D,\ C - c_i > D,\ \bar{A} - c_i - d > D, & \alpha_i \in (a_3, +\infty) \end{cases}$$

Proof of Lemma 13.1. The proof follows the same form as the proof for 13.2. ∎

Proof of Lemma 13.2. Expand the value function

$$V(Q) = Q\, V^1(Q) + (1 - Q)\, V^2(Q) \tag{35}$$

with $V^1(Q)$ to be value at Q if in fact i and M are of the same type and $V^2(Q)$ if i and M are of different types. Let $\widetilde{V}(Q)$ to be the value function if we remove the option of the agent observing s_M^ℓ. We want to show

$$V^1(Q) \geq \widetilde{V}(Q) \geq V^2(Q) \tag{36}$$

with the last inequality holding strictly whenever the optimal policy assigns positive probability to the consumer ever observing s_M^ℓ.

The second inequality in (36) comes from the fact that $\widetilde{V}(Q)$ follows the policy that maximizes the flow of utility when observing s_M^ℓ is not an option. If i and M are not of the same type, then they can do not better than by following the policy of $\widetilde{V}(Q)$, hence, $\widetilde{V}(Q) \geq V^2(Q)$. If the consumer ever observes s_M^ℓ, then they both pay the cost d and also receive information that will almost surely lead them to choose a suboptimal action. In this case, $\widetilde{V}(Q) > V^2(Q)$.

To demonstrate that the first inequality in (36) holds, assume for a contradiction that it does not hold $V^1(Q) < \widetilde{V}(Q)$. From what was previously shown, this would imply $V(Q) < \widetilde{V}(Q)$. This implies a contradiction as the policy for V could be modified to never acquire s_M^ℓ guaranteeing $V(Q) = \widetilde{V}(Q)$.

Finally, it follows from the inequalities in (36) that if $Q' > Q$

$$V(Q') \geq Q'V^1(Q) + (1 - Q')V^2(Q) \geq V(Q) \tag{37}$$

with the last inequality holding strictly if there is a positive probability of the consumer ever observing s_M^ℓ. ■

Proof of Lemma 13.3. This is where the assumed normal-normal conjugate environment comes into play. Recall that a consumer's expected payoff to $\hat{\mathbf{s}}$ (ignoring the costs) is $-\alpha_i Var(\theta_i^*|\hat{\mathbf{s}})$. Without loss of generality, set $\alpha_i = 1$.

We want to show $\bar{A} - \bar{B} < C - D$. The payoff to $\hat{s}_i$ with precision τ is $-(\tau + \tau_0)^{-1}$ and the payoff to not observing s_i is $-\tau_0^{-1}$. The effect of observing s_j is to increase τ_0. The desired inequality is obtained by differentiating

$$\frac{d}{d\tau_0}\big[-(\tau + \tau_0)^{-1} + \tau_0^{-1}\big] = (\tau + \tau_0)^{-2} - \tau_0^{-2} \tag{38}$$

which is less than zero whenever $\tau > 0$. ■

Proof of Lemma 13.4. As $A(Q)$ and $B(Q)$ give the maximized values for objective functions that are differentiable in x and continuous in $Q \in [0,1]$, the *Theorem of the Maximum* implies that $A(Q)$ and $B(Q)$ are continuous in Q. ■

Proof of Lemma 13.5. $\alpha_i \in [0, b_0]$. All inequalities hold strictly at $\alpha_i = 0$. Both sides of each inequality decrease linearly in α_i, with the right side decreasing at a faster rate. Hence, there exists a unique α_i at which point the each inequality becomes an equality. That $\bar{B} - d = D$ implies $C - c_i < D$ follows from $\bar{B} = C$ and $d < c_i$. That $\bar{B} - d = D$ implies $\bar{A} - c_i - d < D$ follows from part 3 of the lemma and some algebra.

$\alpha_i \in (a_1, b_1]$. That $C - c_i = D$ implies $\bar{A} - c_i - d < D$ follows again from part 3 of the lemma. ■

Proof of Proposition 9. For $\alpha_i \in [0, b_0]$, any policy that acquires a signal yields an expected payoff that is less than (strictly so if $Q^0 < 1$) the one that does not select any in each period.

For $\alpha_i \in (a_3, +\infty)$, the inequalities in Lemma 13 entail that no policy that ever selects $\emptyset$ or $\hat{s}_m$ is optimal. Hence, we can consider only plans that select $\hat{s}_i$ or $(\hat{s}_i, \hat{s}_M)$ in each period.

Stationarity of the optimal policy implies $V(Q|$ if $\hat{s}_i$ is chosen$) = \frac{C - c_i}{1+\beta_i}$ and is constant in Q. $V(0|$ if $\hat{s}_i$ is chosen$) > V(0|$ if $(\hat{s}_i \hat{s}_M)$ is chosen$)$, $V(1|$ if $\hat{s}_i$ is chosen$) < V(1|$ if $(\hat{s}_i \hat{s}_M)$ is chosen$)$, and by Lemma 13.2, $V(Q|$ if $(\hat{s}_i \hat{s}_M)$ is chosen$)$ is increasing in Q. Hence, there exists a $\bar{Q}$ such that $\hat{s}_i$ is the optimal control for $Q \leq \bar{Q}$ and $(\hat{s}_i, \hat{s}_M)$ the optimal control for $Q \geq \bar{Q}$.

For $\alpha_i \in (a_1, b_1)$, no policy that ever selects only $\hat{s}_i$ in a period is optimal. We can thus restrict our attention to policies that only ever select from $\{\emptyset, \hat{s}_M, (\hat{s}_i, \hat{s}_M)\}$. As before, there exists cutoff $\bar{Q}$ such that $\emptyset$ is optimal at $Q \leq \bar{Q}$ and either $\hat{s}_M$ or $(\hat{s}_i, \hat{s}_M)$ are optimal for $Q \geq \bar{Q}$

There exists another interior cutoff $\widehat{Q}$ such that if $Q > \widehat{Q}$ then $\hat{s}_M$ is optimal. This follows by noting that the value function, when $(\hat{s}_i, \hat{s}_M)$ is optimal, is bounded from above

$$V(Q| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) < A(Q) - c_i - d + \beta_i \frac{\bar{B} - d}{1 - \beta_i}. \tag{39}$$

As the optimal policy is stationary, if only $\hat{s}_M$ is ever chosen, it will always be chosen thereafter, yielding a value

$$
\begin{aligned}
V(Q| \text{ if } \hat{s}_M \text{ is chosen}) &= B(Q) - d + \beta_i V(Q| \text{ if } \hat{s}_M \text{ is chosen}) \\
\iff V(Q| \text{ if } \hat{s}_M \text{ is chosen}) &= \frac{B(Q) - d}{1 - \beta_i}
\end{aligned} \tag{40}
$$

The right side of (40) exceeds the right side of (39) when

$$
B(Q) - A(Q) > -c_i + \frac{\beta_i}{1 - \beta_i}(\bar{B} - B(Q)). \tag{41}
$$

As the above inequality holds at $Q = 1$ for this domain of α_i and $A(Q)$ and $B(Q)$ are continuous in Q, the aforementioned cutoff $\widehat{Q}$ exists.

Let us show that for c_i sufficiently high $\widehat{Q} < \bar{Q}$, i.e. the consumer will never experiment, either always selecting $\emptyset$ or $\hat{s}_M$ depending on the relation of Q^0 and $\bar{Q}$. If this is true, then the stationarity of the policy implies

$$
V(\bar{Q}) = \frac{D}{1 - \beta_i} = \frac{B(\bar{Q}) - d}{1 - \beta_i} \tag{42}
$$

Appealing to the upper bound on $V(\bar{Q}| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen}))$, we can select c_i sufficiently large such that the right side of (39) is less both terms in the equality of (42).

Next, we show that if d is not too large, there will be an open interval under which the consumer experiments, i.e. $\bar{Q} < \hat{Q}$. Suppose for a contradiction that for all $\alpha_i \in (a_1, b_1)$ and all $d > 0$, $\hat{Q} \leq \bar{Q}$. This implies again that equation (42) holds.

For all $\epsilon_1 > 0$, we can choose α_i arbitrarily close to b_1, such that for

$$
C - c_i < D < C - c_i + \epsilon_1. \tag{43}
$$

Suppose we modify the policy to select $(\hat{s}_i, \hat{s}_M)$ at $\bar{Q}$. Then the value at $\bar{Q}$ is given by

$$
V(\bar{Q}| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) = A(\bar{Q}) - c_i - d + \beta_i E[V(Q')|\bar{Q}, \hat{s}_i, \hat{s}_M] \tag{44}
$$

We can write

$$E[V(Q')|\bar{Q}, \hat{s}_i, \hat{s}_M] =$$

$$\Pr(Q' > \bar{Q}) \cdot E[V(Q')|Q' > \bar{Q}] + Pr(Q' \leq \bar{Q}) \cdot E[V(Q')|Q' \leq \bar{Q}] \tag{45}$$

where all terms in the above equality condition on $\bar{Q}$ and $\hat{s}_i, \hat{s}_M$. As $E[V(Q')|Q' > \bar{Q}] = E[B(Q') - d|Q' > Q] > \frac{B(\bar{Q})-d}{1-\beta_i}$ and $E[V(Q')|Q' \leq \bar{Q}] = \frac{D}{1-\beta_i}$,

$$V(\bar{Q}| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) > A(\bar{Q}) - c_i - d + \beta_i \frac{D}{1-\beta_i} > A(\bar{Q}) - C + D - \epsilon_1 - d + \beta_i \frac{D}{1-\beta_i} \tag{46}$$

where the inequality $A(\bar{Q}) - C + C - c_i > A(\bar{Q}) - C + D - \epsilon_1$ follows from (43). For d and ϵ_1 sufficiently small, $V(\bar{Q}|$ if $(\hat{s}_i, \hat{s}_M)$ is chosen$)) > V(\bar{Q}|$ if $\emptyset$ is chosen$))$. Hence, the assumed optimal policy is in fact suboptimal, implying a contradiction.

The proof for "experimentation" at $\alpha_i = b_1$ and $\alpha_i \in (a_2, b_2]$, is nearly identical to the proof above. ■