Ideological Bias and Trust in Information Sources

Matthew Gentzkow* — *Stanford and NBER*
Michael B. Wong — *MIT*
Allen T. Zhang — *Harvard*

September 2020

Abstract

We study the role of endogenous trust in amplifying ideological bias. Agents in our model seek to learn a sequence of states using information from sources whose accuracy is *ex ante* uncertain. Agents also receive noisy feedback about the states from direct observation. Arbitrarily small biases in the processing of this feedback can cause large ideological differences in the agents' trust in information sources and their beliefs about the states, and may lead agents to become overconfident in their own judgment. These patterns can be similar whether agents see only ideologically aligned sources or see a diverse range of sources.

Keywords: Bias, trust, polarization, media bias
JEL: D83, D85, L82

*Email: gentzkow@stanford.edu, mbwong@mit.edu, allenzhang@fas.harvard.edu. We thank Daron Acemoglu, Isaiah Andrews, David Autor, Abhijit Banerjee, Ben Golub, Jesse Shapiro, Francesco Trebbi, David Ritzwoller, Ashesh Rambachan and seminar participants at UC Berkeley, the University of Michigan, MIT, and Stanford for helpful comments and suggestions. We acknowledge funding from the Stanford Institute for Economic Policy Research (SIEPR), the Toulouse Network for Information Technology, the Knight Foundation, the National Science Foundation and the Stanford Cyber Initiative.

1 Introduction

Ideological divisions in society often seem intractable, with those on either side persistently disagreeing about objective facts. In recent years, for example, fervent debates over the validity of global warming, evolution, and vaccination have persisted long after the establishment of a scientific consensus. Partisans also disagree about which sources can be trusted to provide reliable information about such facts. In the United States, for instance, 75 percent of conservative Republicans say they trust news and information from Fox News, while 77 percent of liberal Democrats say they distrust it (Pew 2020). Such divisions have deepened even as new media technologies have made information more widely and cheaply available than ever before. The information age has, paradoxically, produced what has been dubbed a “post-truth” era (Keyes 2004).

Such patterns seem at odds with the prediction of many Bayesian (e.g., Blackwell and Dubins 1962) and non-Bayesian (e.g., DeGroot 1974) learning models in which widespread availability and distribution of information leads all agents’ beliefs to converge to the truth. Many possible alternatives have been proposed, including models with psychological biases that give individuals a taste for cognitive consistency or confirmation (e.g., Lord Ross and Lepper 1979; Cotton 1985). To explain the patterns we observe, such accounts generally require the magnitude of biases to be large, so that individuals effectively place limited weight on learning the truth relative to other explicit or implicit objectives.

In this paper, we explore a different possibility, which is that rational Bayesian inference may magnify the influence of even small biases when agents are uncertain which sources they can trust. Building on recent insights by Acemoglu, Chernozhukov, and Yildiz (2016) and Sethi and Yildiz (2016), among others, we show that arbitrarily small biases in information processing may lead to substantial and persistent divergence in both trust in information sources and beliefs about facts, with partisans on each side trusting unreliable ideologically aligned sources more than accurate neutral sources, and also becoming overconfident in their own judgment. Consistent with recent evidence suggesting that the magnitude of selective exposure has generally been limited (Gentzkow and Shapiro 2011; Flaxman et al. 2016), we show that these patterns arise whether agents selectively view only ideologically aligned sources or are exposed to a diverse range of sources. Increasing the number of available information sources in such a setting may deepen rather than mitigate ideological differences.

An agent in our model wishes to learn about a sequence of unobserved states $\omega_t \sim N(0,1)$, which are drawn independently in each period t. We think of each period's ω_t as capturing a distinct item discussed in the news. In one period this might be the magnitude of global warming due to human activity, in the next period the truth about a political scandal, in a third period the risk of vaccines, and so on. Each period the agent observes a normally distributed signal s_{jt} correlated with ω_t from one or more information sources j. We refer to the correlation between s_{jt} and ω_t as the *accuracy* of source j. We analyze two scenarios, one in which the agent observes exactly one source j in each period (she "single-homes" in the language of Rochet and Tirole 2003), and another in which she observes all j in each period (she "multi-homes").

To introduce a political dimension to the model, we assume that the issue in each period t is associated with an *ideological valence* $r_t \sim N(0,1)$. We think of r_t as the belief about ω_t that would be most favorable to conservatives and $-r_t$ as the belief that would be most favorable to liberals, normalizing the state so that $r_t = 0$ is the politically neutral position in each period. We allow the information sources j to have ideological *biases* in the sense that their errors may be correlated with r_t. The accuracies and biases of the sources are the main persistent state variables that the agent seeks to learn over time.

The final essential ingredient of our model is feedback x_t about the true state that the agent observes directly. In the case of global warming, this might be weather events she experiences. In the case of vaccine risk, this might be observations of children in her social network who have developed autism. While such feedback will typically be noisy and so not be very informative about ω_t on its own, its key features are that the agent knows how to interpret it (in the sense that she has a strong prior on its correlation with ω_t) and that she believes it to be unbiased (in the sense that she believes x_t is conditionally independent of r_t given ω_t). As a result, comparing the feedback x_t to the reports of the information sources s_{jt} allows her to learn which sources she can trust.

We model ideological bias on the part of the agent by assuming the x_t she observes is in fact distorted in the direction of r_t or $-r_t$. Such bias could arise due to motivated reasoning, selective memory, or availability, for example. Because the agent rules out the possibility of such biases by assumption, they may over time lead to distorted learning about the accuracy of information sources and, consequently, the states ω_t. We characterize the form such distortions take and ask whether they can be large even when the magnitude of the agent's bias is small.

Our formal results characterize the limiting distribution of the agent's beliefs about accuracies and states as $t \to \infty$, focusing on the special case in which the agent's priors on the accuracy of the feedback x_t are concentrated at the true value. We show that in this case her beliefs about the accuracies of information sources eventually place probability one on a single vector, which we define as the agent's asymptotic *trust* in the respective sources.

In the benchmark case in which the agent has no ideological bias, her trust α^* is equal to the true accuracies α_0 of the information sources, and her asymptotic beliefs about ω_t are the same as if she knew the true data generating process. If the accuracy of the information sources is sufficiently high, her beliefs about ω_t are close to correct.

Introducing arbitrarily small biases in the agent's feedback x_t changes the results of the benchmark case dramatically. An agent with a small conservative bias may come to trust right-leaning sources more than is warranted by their true accuracy, trust unbiased sources less than is warranted, and believe that left-leaning sources are perverse, in the sense that their signals are negatively correlated with the true state. She may become overconfident in the sense that she believes the accuracy of her own feedback x_t is greater than it really is. She will generally come to believe that the state ω_t is positively correlated with the ideological state r_t, and thus begin any period in which she knows r_t with a conservatively biased prior. All of these effects may be large if the accuracy of x_t is sufficiently low.

To see the intuition for the way small biases are amplified in our model, note that an agent will come to see source j as more accurate the greater the observed correlation between its report s_{jt} and her feedback x_t. When the feedback x_t is noisy, this correlation will be small even when s_{jt} is perfectly accurate. Small differences in observed correlation—such as those that might be induced by a small ideological bias—thus imply large differences in accuracy.

Distortions in the way agents learn about the informational environment can translate into disagreements about the states ω_t. We first show how biases affect the accuracy of agents' posterior beliefs about ω_t. We then show how the magnitude of disagreement between different agents depends on the accuracy and bias of their prior information, as well as on those of the observed sources. When agents all observe a common unbiased source, disagreement is generally small.

When biased sources are introduced to the market, arbitrarily small biases can lead to arbitrarily large disagreement.

The final section of our main results considers how these results differ under single and multi-homing. A common intuition is that divergent trust and polarization might be mitigated if agents were exposed to an ideologically diverse set of information sources. We show that it is possible for multi-homing to have beneficial effects consistent with this intuition, but also that this need not be the case. Multi-homing may leave trust and polarization unchanged, or even exacerbate them.

Two extensions explore the implications of ideological bias for media competition and political behavior. First, we endogenize the choice of slant by media outlets in a sequential positioning game. We find that media competition can lead to greater media slant as well as intensified disagreements among viewers. Second, we show mistrust of motives across ideological divides can arise when agents underestimate both their own and others' biases. Ideological bias in this case can intensify political conflict, leading to costlier battles for power.

Our paper is closely related to the literature on Bayesian asymptotics. In a seminal contribution, Blackwell and Dubins (1962) show that Bayesian agents observing signals with increasing information eventually agree on the distribution of future signals. However, beliefs about payoff-relevant states need not converge. Berk (1966) shows that limiting posteriors have nonzero support on the set of all identifiable values. Acemoglu, Chernozhukov, and Yildiz (2016) show that asymptotic agreement is fragile in that arbitrarily small differences in beliefs about the interpretation of signals can generate large disagreements about an underlying state. Our contribution follows a similar logic. We study a setting with multiple states and multiple information sources, and show that an arbitrarily small misspecification in the agent's model of the signal generation process generates significantly divergent trust of information sources.

Our work relates to a growing literature on models of opinion polarization and ideology. Baliga, Hanany, and Klibanoff (2013) provide a simple statement that beliefs cannot polarize under Bayesian updating when agents share the same theory connecting parameters to signals.[1] However, belief polarization is shown to arise when Bayesian agents have confirmatory bias (Rabin

[1]Their stark result contrasts with a large number of experimental studies (e.g., Lord, Ross, and Lepper 1979) wherein the beliefs of subjects polarized after the presentation of new evidence. Dixit and Weibull (2007) is an early paper that considers the possibility of belief polarization under Bayesian updating. Kartik, Lee, and Suen (2020) show that Bayesians expect disagreements to reduce upon the arrival of any new information.

and Shrag 1999), face costly information (Suen 2004), are averse to ambiguity (Baliga, Hanany, and Klibanoff 2013), have limited memory or attention (Fryer, Harms, and Jackson 2015; Che and Mierendorff 2019), make inferential mistakes about the credibility of sources (Cheng and Hsiaw 2019), or disagree about the interpretation of signals (Andreoni and Mylovanov 2012; Kondor 2012; Glaeser and Sunstein 2014; Benoit and Dubra 2019). Ortoleva and Snowberg (2015) and Levy and Razin (2015) explore how correlation neglect and resulting overconfidence impact polarization and political behavior.

Our theory also relates to the literature on observational learning in social networks. One strand of this literature considers opinion dynamics in social networks with non-Bayesian learning rules, beginning with DeGroot (1974). DeMarzo, Vayanos, and Zwiebel (2003) and Golub and Jackson (2010, 2012) are state of the art models that characterize conditions under which disagreements may persist in groups. Another strand of this literature considers Bayesian learning, and begins with Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). In these models, Bayesian individuals only observe posterior beliefs or actions of other individuals, and may fail to learn underlying states because they are able only to recall or communicate coarse information. For instance, Sethi and Yildiz (2012; 2016) examine disagreements between individuals with heterogeneous priors who communicate posterior beliefs. By contrast, our model assumes no coarseness of information. Disagreements arise due to small deviations from Bayesian information processing, and persist even when individuals have access to the same set of signals.

Finally, our work relates to the literature on media bias (Mullainathan and Shleifer 2005; Gentzkow, Shapiro and Stone 2016). The mechanism by which agents in our model come to trust like-minded sources is closely related to the one explored by Gentzkow and Shapiro (2006). That model is essentially static, however, and does not provide a mechanism by which diverging beliefs or trust can persist over time. A large empirical literature studies the link between media markets and political polarization (Glaeser and Ward 2006; McCarty, Poole, and Rosenthal 2006; Campante and Hojman 2013; Prior 2013). A growing experimental literature examines the link between trust in information sources and political beliefs (Levendusky 2013; Nisbet, Cooper, and Garrett 2015; Benedictis-Kessner et al 2019; Thaler 2020; Jo 2020).

The paper proceeds as follows. Section 2 describes the model. Section 3 characterizes the asymptotic distribution of beliefs. Section 4 presents our results on overconfidence, trust, and disagreement. Section 5 compares the single and multi-homing cases. Section 6 presents extensions. Section 7 concludes.

2 Model

2.1 Setup

An agent wishes to learn about a sequence of states $\omega_t \sim N(0,1)$ over time periods $t = 1,2,\ldots\infty$. In each period she observes a signal $s_{jt} \sim N(0,1)$ informative about ω_t from one of the available information sources $j = 1,...,J$. Each source is observed infinitely many times in the limit as $t \to \infty$. We refer to this case where the agent observes a single source in each period as *single-homing.* We consider the *multi-homing* case where the agent observes multiple sources in each period in Section 5 below.

While we describe the model from the perspective of a representative agent, the model can also be interpreted as describing a population of heterogeneous agents each of which solves the problem we describe. We revisit this interpretation when we consider polarization of beliefs.

The agent may also observe information $x_t \sim N(0,1)$ about ω_t based on direct observation. We describe this as "feedback" about the state realized after signals are observed, but, as discussed below, it can also represent prior information.

In each period, the ideological valence of ω_t is represented by an ideological state $r_t \sim N(0,1)$. We think of r_t as representing the conservative position on issue ω_t (i.e., the realization of ω_t most consistent with a conservative platform) and $-r_t$ as representing the analogous liberal position.

With some probability strictly between zero and one both x_t and r_t are observed by the agent. Otherwise, neither is observed. This is meant to capture the idea that there is a subset of issues about which the agent has some independent knowledge and other issues about which her only information comes through s_{jt}. Whether or not x_t and r_t are observed is independent across periods.

Together, ω_t, r_t, x_t, and the J-vector s_t of s_{jt} are jointly standard normal and are drawn independently over time. The ideological state r_t is given by

$$r_t = \gamma\omega_t + \sqrt{1-\gamma^2}\tilde{r}_t \tag{1}$$

where $\gamma \in [-1,1]$ is the correlation of r_t with ω_t and $\tilde{r}_t \sim N(0,1)$ is a residual component independent of ω_t. In our benchmark case we will assume the true value of γ is zero, so that ω_t and r_t are independent, but biased agents may come to mistakenly believe that $\gamma \neq 0$.

Each signal s_{jt} is given by

$$s_{jt} = \alpha_j \omega_t + \beta_j \tilde{r}_t + \varepsilon_{jt} \tag{2}$$

where $\alpha_j \in [-1,1]$ is the correlation of s_{jt} with ω_t, $\beta_j \in [-1,1]$ is the correlation of s_{jt} with $\tilde{r}_t$, these satisfy $\alpha_j^2 + \beta_j^2 \leq 1$, and $\varepsilon_{jt} \sim N\left(0, 1-\alpha_j^2-\beta_j^2\right)$ is independent of ω_t and $\tilde{r}_t$. We refer to α_j and β_j as the *accuracy* and *bias* of signal j respectively. We let s_t, ε_t, α, and β denote the J-vectors of s_{jt}, ε_{jt}, α_j, and β_j respectively.

We refer to sources that satisfy the condition $\alpha_j^2 + \beta_j^2 \leq 1$ as *feasible*. We refer to sources that satisfy this condition with equality as *frontier source*s. Note that the feasibility condition—which follows from our assumption that all signals are marginally distributed $N(0,1)$—builds into the model an inherent trade-off between accuracy and bias.

Feedback x_t is given by

$$x_t = a\omega_t + b\tilde{r}_t + \eta_t, \tag{3}$$

where $a \in (0,1]$ is the correlation of x_t with ω_t, $b \in [-1,1]$ is the correlation of x_t with $\tilde{r}_t$, these satisfy $a^2 + b^2 \leq 1$, and $\eta_t \sim N\left(0, 1-a^2-b^2\right)$ is independent of ω_t, $\tilde{r}_t$, and all ε_{jt}. We refer to a and b as the agent's *accuracy* and *bias.* As discussed in more detail below, we are mainly interested in the case where x_t provides noisy information about ω_t and so a is small. When b and β_j have the same sign we say that the agent and source j are *like-minded* and when they have the opposite sign we say the agent and source j are *opposite-minded.*

The parameters of this model are $\theta = (a, b, \alpha, \beta, \gamma)$. We denote the set of all such parameters by Θ.[2] We let θ_0 denote the true value of θ, and where it adds clarity we will use a_0, α_{0j}, b_0, β_{0j}, and γ_0 to refer to the true values of the individual components. While the agent entertains the possibility of any $\theta_0 \in \Theta$, we will focus throughout on the case where $\gamma_0 = 0$, so the truth is uncorrelated with the ideological state, where $a_0^2 + b_0^2 < 1$, so the feedback x_t is not perfectly correlated with any other elements of the model, and where $\left|\beta_{0j}\right| < 1$, so signals are not perfectly correlated with r_t.

[2] Θ is the set of all $\theta \in (0,1] \times [-1,1]^{2J+2}$ such that $a^2 + b^2 \leq 1$ and $\alpha_j^2 + \beta_j^2 \leq 1$.

2.2 Prior Beliefs

Agents in the model are initially uncertain not only about the states ω_t, but also about the parameters θ that govern the distribution of prior information and signals. Let $(\Theta, \mathscr{L}_\Theta, \nu)$ denote the Lebesgue space on Θ.[3] At the beginning of the first period, the agent has an absolutely continuous prior belief μ_0 on Θ with a continuous density with respect to ν.

The central assumption of our model is that the agent believes that her own independent information x_t is unbiased. We assume that μ_0 has full support on the subset of Θ consistent with this restriction.

Assumption 1. *The support of μ_0 is the set $\Theta^{prior} \subset \Theta$ for which $b = 0$.*

We will be mainly interested in the case where the agent's beliefs about the accuracy a of her own signal are concentrated at the true value a_0. However, we do not assume that the agent places prior probability one on this value because, as will become clear below, the data the agent observes may not be consistent with both $b = 0$ and $a = a_0$. In these cases, we want to allow the agent to adjust her belief about a to rationalize the data. We will thus focus on the limit of a sequence of priors that become increasingly concentrated at values close to a_0 while maintaining the property of full support on Θ^{prior}.

Definition. Let $\{\mu_{0,n}\}_{n=1}^{\infty}$ be a sequence of prior distributions with full support on Θ^{prior} and continuous densities that converges weakly to a limit μ_0^*. For a given measurable set $\vartheta \in \mathscr{L}_\Theta$, define $d(\vartheta, a_0) = \inf_{\theta \in \vartheta} |a(\theta) - a_0|$. We say that such a sequence *becomes concentrated at* a_0 if for any $\vartheta', \vartheta'' \in \mathscr{L}_\Theta$ such that $d(\vartheta', a_0) < d(\vartheta'', a_0)$, we have that $\mu_{0,n}(\vartheta'')/\mu_{0,n}(\vartheta') \to 0$ as $n \to \infty$.

This definition implies that $\mu_{0,n}$ converges to a distribution degenerate at $a = a_0$. It is stronger than this, however, because it also requires that for any values a and a' away from a_0 such that a is closer to a_0 than a', the prior eventually places arbitrarily more weight on a than on a'.

Our characterization of asymptotic beliefs below will focus on the limiting posterior as priors become concentrated at a_0 in this sense. If the distribution of the data is in fact inconsistent with a_0, this posterior will place arbitrarily high weight on the value of a closest to a_0 that is consistent with the data.

[3] $\mathscr{L}_\Theta$ is the σ-algebra of Lebesgue measurable sets, and ν is the corresponding Lebesgue measure.

2.3 Discussion

The feedback x_t plays a central role in our model. As a baseline case, we think of x_t as information about the true value of the state ω_t that the agent observes directly. This could be weather events in the agent's locality (when ω_t relates to global warming), outcomes of vaccinated and unvaccinated children in the agent's social network (when ω_t relates to the safety of vaccines), the agent's own experience with public schools (when ω_t relates to education policy), or the agent's personal economic situation (when ω_t relates to economic policy). We describe x_t for simplicity as feedback observed *ex post,* after the agent has seen the reports s_{jt} of the information sources. However, it could equally well be interpreted as prior information that the agent obtained before seeing s_{jt}.

What is crucial for our model is that x_t satisfies three conditions. First, the agent has a strong prior on its accuracy a which is concentrated at the truth. Second, the agent believes that it is free from ideological bias ($b = 0$) with probability one. Together, these two properties define the difference between x_t and the information sources s_{jt}. Whereas the accuracy and bias of the latter are uncertain from the agent's perspective, the agent is confident that she understands and can trust x_t. Thus, x_t provides the key reference point that the agent uses to learn which other sources can be trusted.

The third condition is that, contrary to the agent's belief, x_t may in fact be subject to ideological bias. There are a large number of well-studied psychological phenomena that could provide a micro-foundation for this bias. One set of these falls under the heading of motivated reasoning (Kunda 1990). Consider an agent who has grown up in a liberal family, benefitted from liberal policies, and taken actions (like voting) consistent with liberal ideology. She may underweight feedback pointing in the conservative direction in order to reduce cognitive dissonance (Festinger 1957). She may be more likely to remember evidence consistent with a liberal view — for example, remembering unusually hot days or unusually severe storms that suggest global warming is severe. She may tilt her assessment of the credibility of evidence due to confirmation bias (Lord Ross and Lepper 1979). She may also live in an environment in which information that supports her position is more "available" in the sense of Tversky and Kahneman (1973) – for example, if she has gone to high quality public schools she may find it easier to think of the benefits of teachers' unions than their costs. Finally, it may be that evidence that supports the liberal position is more salient in the sense of Bordalo, Gennaioli, and Shleifer (2012).

While we describe x_t as directly observed information, several other interpretations are possible. One is that x_t is the result of the agent's reasoning and introspection about the likely value of ω_t. If ω_t relates to economic stimulus policy, the agent might reason from first principles about how large the plausible costs and benefits could be. (She might even write down and solve a model!) The outcome of such reasoning could be captured in x_t. Here, the key assumptions are that she has a confident assessment of her own mental capacities and does not think her own reasoning is biased. Another possibility is that x_t is the signal of a particular information source that the agent believes *a priori* to be unbiased. This might be what her mother says, or what the Bible says, or what scientists say, or even the report of a particular news source that she begins with extraordinary faith in. What is key is that she takes the reliability of this source as axiomatic and does not entertain the possibility that it could be biased.

It may seem strong to assume that the agent puts such dogmatic *ex ante* faith in any particular source of information. However, if there is *no* information in which she would place such faith, the agent's problem is fundamentally unidentified. She would never be able reject either (i) all sources are unbiased and at least one of them is perfectly accurate; (ii) all sources have accuracy zero and at least one of them is maximally biased. We make this precise in Proposition 5 below.

3 Asymptotic Learning

Because each of the observed variables (x_t, r_t, s_t) is distributed marginally $N(0,1)$, the correlations among these variables exhaust the information that the agent can learn from the data. Moreover, since (in the single-homing case) only one s_{jt} is ever observed in a particular period and all variables are independent across t, the agent learns no information about the correlation among the elements of s_t. Thus, the distribution of observable data is fully parameterized by the vector of correlations $R = (\rho_{xs}, \rho_{xr}, \rho_{rs})$, where the J-vector ρ_{xs} is the correlation of x_t and s_t, the scalar ρ_{xr} is the correlation of x_t and r_t, and the J-vector ρ_{rs} is the correlation of r_t and s_t. We let $R(\theta)$ denote the correlations implied by θ and we let $\mathscr{R} = \{R : R = R(\theta), \theta \in \Theta\}$.

Remark 1. Given parameters $\theta = (a, b, \alpha, \beta, \gamma)$, the elements of $R(\theta)$ are

$$\rho_{xs} = a\alpha + b\beta$$

$$\rho_{xr} = a\gamma + b\sqrt{1-\gamma^2}$$

$$\rho_{rs} = \alpha\gamma + \beta\sqrt{1-\gamma^2}.$$

As discussed in Section 2.1, we focus on the case where none of x_t, s_t, and r_t are perfectly correlated with each other, so the vector of true correlations $R_0 \in \text{int}(\mathscr{R})$.

Given that there are $2J+3$ parameters and $2J+1$ observed moments, θ will be partially identified by the data. Let $I(R) = \{\theta : R(\theta) = R\}$ denote the *identified set* of parameters consistent with correlations R and let $I(R; b=0)$ denote the subset of these parameters consistent with the agent's assumption that $b = 0$. Both of these sets are always non-empty, since for any $R \in \mathscr{R}$ we can choose parameters $a = 1$, $b = 0$, $\alpha = \rho_{xs}$, $\gamma = \rho_{xr}$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \alpha\gamma)$. Both of these sets will also typically include elements with $b = 0$ and $a < 1$, as we will show in proposition 1.

A key feature of inference under the agent's model of the world is that the magnitude of the observed correlations ρ_{xs} and ρ_{xr} place a lower bound on the accuracy a of the agent's own feedback x_t. Since the agent believes $b = 0$, correlation with the state ω_t provides the only mechanism by which x_t can be correlated with s_t and r_t. If her feedback was completely uninformative ($a = 0$), there should be no such correlation and the agent should see $\rho_{xs} = \rho_{xr} = 0$. The larger are $|\rho_{xs}|$ and $|\rho_{xr}|$, the larger is the value of a needed to rationalize the observed data.

The precise bound on a depends not only on $|\rho_{xs}|$ and $|\rho_{xr}|$ but on the covariances between r and s as well. If the agent rationalizes the correlations $|\rho_{xs}|$ and $|\rho_{xr}|$ by assuming a low value of a and relatively high values of α and γ, she expects the correlation between r and s to be relatively high. Observing that the correlation is in fact lower means that the data can only be rationalized by a higher value of a and a lower value of α and γ. The exact condition is that a cannot be smaller than $\underline{a}^R = \max_j \sqrt{\zeta_j}$, where ζ_j is the population R^2 from a regression of x_t on r_t and s_{jt}.

Proposition 1. *The identified set $I(R; b=0)$ consistent with the observed data $R = (\rho_{xs}, \rho_{xr}, \rho_{rs})$ is non-empty and consists of all $\theta \in \Theta$ with $a \in [\underline{a}^R, 1]$, $b = 0$, $\alpha = \frac{\rho_{xs}}{a}$, $\gamma = \frac{\rho_{xr}}{a}$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$.*

Proof. See appendix. □

When the agent is biased ($b_0 > 0$), the true parameters of the model lie outside the support of the agent's prior, and our model is an example of Bayesian learning under misspecification (Lian 2009). Characterizing the evolution of beliefs in such cases is complicated in general, and can lead to instability or lack of convergence in the limit (Berk 1966). However, Proposition 1 shows that the data the agent observes can always be rationalized by some $\theta \in \Theta^{prior}$ that *does* fall within the support of her prior. Thus, the data will never violate her model of the world, and we show that her beliefs will be well behaved asymptotically as a result.

Proposition 2. *Suppose the true correlations of the observed data are $R_0 \in \mathscr{R}$. Then as $t \to \infty$ under single homing, the agent's posterior distribution converges to a limit μ_∞ such that for all measurable $\vartheta \subseteq \mathscr{L}_\Theta$,*

$$\mu_\infty(\vartheta) = \frac{\mu_0\left(\vartheta \cap I\left(R_0 : b = 0\right)\right)}{\mu_0\left(I\left(R_0 : b = 0\right)\right)}.$$

Proof. See appendix. □

Thus, the agent's beliefs converge asymptotically to an identified set consistent with the observed correlations R. They place probability zero on parameter values outside this set. Because all parameters in the identified set imply the same distribution of observed data, beliefs within the set remain proportional to the prior.

The structure of the identified set makes clear that lack of identification of the accuracy a of the agent's own signal and of the accuracy α of the information sources go hand-in-hand. All a single-homing agent can ever learn is the correlations among x_t, r_t, and the elements of s_t. A given R could result from a high value of a and relatively low values of the α_j, or a low value of a and relatively high values of the α_j; these will never be distinguished by the observed data. There is a one-to-one correspondence between values of a and values of α within the identified set.

As discussed above, we are mainly interested in the case where the agent's prior becomes concentrated at the true value a_0. We show that in this case the agent's posterior on a converges to a unique value—either a_0 if this is included in the identified set, or the closest value that is included—which turns out to be $\underline{a}^R$—if not. The agent's posterior on α converges to the unique value associated with that a.

Definition. Let $\{\mu_{0,n}\}_{n=1}^{\infty}$ be a sequence of prior beliefs that becomes concentrated at a_0, and let $\mu_{\infty,n}$ be the asymptotic posterior distribution defined by Proposition 2 when the prior is $\mu_{0,n}$. Then the agent's *limiting posterior* is μ_∞^* if $\mu_{\infty,n} \xrightarrow{d} \mu_\infty^*$.

Proposition 3. *Suppose the true correlations of the observed data are $R_0 \in \mathscr{R}$. Then the limiting posterior μ_∞^* exists and is well-defined. Its support contains a unique value θ^* whose elements are $a^* = \max\{a_0, \underline{a}^R\}$, $\gamma^* = \frac{\rho_{xr}}{a^*}$, $\alpha_j^* = \frac{\rho_{xsj}}{a^*}\,\forall j$, and $\beta_j^* = \frac{1}{\sqrt{1-\gamma^{*2}}}\left(\rho_{rsj} - \gamma^* \alpha_j^*\right)\forall j$.*

Proof. See appendix. □

The limiting belief θ^* will be one of the main focuses of our analysis. We refer to a^* as the agent's *confidence* and if $a^* > a_0$ we say she is *overconfident*. We refer to α_j^* as the agent's *trust* in source j and to β_j^* as the *perceived bias* of source j. Finally, we refer to γ^* as the agent's *ideology*. An agent with $\gamma^* > 0$ comes to believe that conservative views are on average closer to the truth than liberal views and she starts each period when r_t is observed with a prior belief about ω_t biased toward r_t. An agent with $\gamma^* < 0$ comes to believe the opposite.

Note that when the agent is unbiased ($b_0 = 0$), we recover the standard result that her beliefs converge to the truth (Blackwell and Dubins 1962). In this case $\underline{a}^R = a_0$ and so the agent is not overconfident ($a^* = a_0$). This in turn implies $\gamma^* = \gamma_0$, $\alpha_j^* = \alpha_{0j} \forall j$, and $\beta_j^* = \beta_{0j} \forall j$.

We define the agent's asymptotic *distribution of beliefs about ω_t* to be the distribution of her posterior mean $\overline{\omega}_t$ about ω_t given that her beliefs about θ are given by the limiting posterior μ_∞^*. For simplicity, we focus on the distribution in periods when the agent observes neither prior information x_t nor ideological state r_t, so we capture only the influence of the information source(s) she observes. The following characterization follows from standard conjugate prior results for the normal distribution.

Proposition 4. *Suppose the agent's beliefs about θ are the limiting posterior μ_∞^* associated with $R \in \mathscr{R}$. In periods where the agent does not observe (x_t, r_t), and observes source j, her posterior belief about ω_t will be a normal distribution with mean $\overline{\omega}_t = \alpha_j^* s_{jt}$ and variance $v_t = 1 - \alpha_j^{*2}$.*

Proof. See appendix. □

The average belief $\overline{\omega}_t$ is a linear function of the observed signal and so the asymptotic distribution of beliefs about ω_t follows immediately from the known distribution of the signal. Note that because a biased agent's limiting posterior may be incorrect, her beliefs need not satisfy the martingale property and so her expected posterior mean on ω_t given the ideological state r_t may be correlated with r_t and significantly different from zero even though the true distribution of ω_t is independent of r_t. When we consider multiple agents, this will also lead to systematic disagreement, and even the possibility that two agents may both be certain about the value of ω_t but differ in what they think is the true value.

As a final result in this section, we consider how asymptotic learning would change if feedback x_t were not available, so the agent did not observe any information source that she *ex ante* believed to be unbiased. In this case, the distribution of observable data is given by $R^{\emptyset} = \rho_{rs}$, since ρ_{xs} and ρ_{xr} are not observed. As discussed in Section 2.3, the agent's problem is fundamentally unidentified in this case. The identified set $I\left(R^{\emptyset}; b = 0\right)$ consistent with observed data $R^{\emptyset}$ contains a wide range of parameter values, including $\alpha_j = 1$ for some source j, or $\alpha_j = 0$ for all sources j. The agent thus cannot rule out the extreme possibilities that either one of the sources always reports the true states or none of them contain any information about the true states. In this sense, asymptotic learning requires that there is some source she is willing to put dogmatic faith in *ex ante.*

Proposition 5. *Suppose an agent does not observe x_t in any period, but still observes r_t. The true correlations of the observed data are $R^{\emptyset} = \rho_{rs}$. The identified set $I\left(R^{\emptyset}; b = 0\right)$ consistent with the observed data $R^{\emptyset}$ includes (i) all $\beta_j = 0$ and at least one $\alpha_j = 1$; (ii) all $\alpha_j = 0$ and at least one $\beta_j \in \{-1, 1\}$.*

Proof. See appendix. □

4 The Implications of Ideological Bias

In this section, we turn to our main results on the evolution of the agent's beliefs about θ and ω_t. The main results in this section are straightforward implications of Proposition 3, and so we state them as corollaries. To simplify the exposition, we focus without loss of generality on the case of $b_0 \geq 0$ and assume throughout that there is at least one information source with $\alpha_{0j} \neq 0$. We will also be particularly interested in cases where at least one informative information source locates on the frontier (i.e., there exists a j with $\alpha_{0j} > 0$ and $\alpha_{0j}^2 + \beta_{0j}^2 = 1$). The model of endogenous entry in Section 6.1 below shows conditions under which sources will choose such locations endogenously.

4.1 Confidence

We begin by considering the determinants of the agent's confidence a^*. Ortoleva and Snowberg (2015) explore in detail the implications of overconfidence for political behavior.[4] In our model, overconfidence arises endogenously as a result of small biases in x_t.

From Proposition 3, the agent will be overconfident if and only if $\underline{a}^R > a_0$. This will occur if at least one ζ_j (the population R^2 from a regression of x_t on r_t and s_{jt}) is greater than a_0^2, and in particular if the absolute correlation of x_t with either r_t or some s_{jt} is greater than a_0.

The following lemma derives the conditions under which $\underline{a}^R > a_0$ in our main case of interest, where the true correlation γ_0 between r_t and ω_t is zero.

Lemma 1. *When $\gamma_0 = 0$, the R^2 from a regression of x_t on r_t and s_{jt} is*

$$\zeta_j = b_0^2 + a_0^2 \left(\frac{\alpha_{0j}^2}{1-\beta_{0j}^2} \right).$$

Therefore, the agent is overconfident if and only if

$$\frac{b_0^2}{a_0^2} > 1 - \max_j \left\{ \frac{\alpha_{0j}^2}{1-\beta_{0j}^2} \right\}.$$

Proof. See appendix. □

Biased agents need not be overconfident. For example, if no sources in the market are on the frontier ($\alpha_{0j}^2 + \beta_{0j}^2 < 1 \; \forall j$), an agent with b_0 is not overconfident provided that b_0 is sufficiently small relative to a_0.

Overconfidence will arise when bias is large. However, even when bias is arbitrarily small, the agent will be overconfident provided the accuracy of her feedback a_0 is sufficiently low. She will also be overconfident if at least one source in the market is sufficiently accurate and sufficiently close to the frontier.

[4] In Ortoleva and Snowberg (2015), agents over-estimate the precision of their information because they ignore correlation in the underlying signals they see. This leads overconfident citizens to have excess variance in their posterior beliefs. Overconfidence in our model has the same excess variance implication (in periods when the agent observes x_t), but also has a further effect on polarization via endogenous trust.

Corollary 1. *Suppose the agent has bias $b_0 > 0$. The agent will be overconfident provided $a_0 < \hat{a}$, where $\hat{a} > b_0$ is a bound that depends on the values of α_0 and β_0. The agent will also be overconfident regardless of the value of a_0 provided that $\alpha_{0j}^2 / \left(1 - \beta_{0j}^2\right)$ is sufficiently close to one for some j. If there is at least one frontier source with $\alpha_j > 0$, the agent is overconfident regardless of the value of a_0, with $a^* = \sqrt{a_0^2 + b_0^2}$.*

4.2 Trust

We can combine Proposition 3 and Remark 1 to derive the agent's trust as a function of her confidence a^*.

Remark 2. The agent's trust in information source j is

$$\alpha_j^* = \frac{a_0}{a^*}\alpha_{0j} + \frac{b_0}{a^*}\beta_{0j}.$$

For a given value of a^*, an agent will thus come to trust source j more the greater the alignment of their biases (i.e., the greater that $b_0\beta_{0j}$ is). The model thus predicts an endogenous preference for like-minded information sources. Trust is always increasing in true accuracy α_{0j} holding both a^* and β_{0j} constant. However, it is possible that increasing α_{0j} actually reduces trust because it causes confidence a^* to rise.

Remark 2 shows the key amplification mechanism that will drive a number of our main results: when the agent believes her own information is noisy ($a^* \approx 0$), small biases b_0 and β_{0j} can translate into large differences in trust. When x_t is noisy its correlation with even a perfectly accurate signal s_{jt} will be small. Its correlation with a completely inaccurate signal will be zero. Thus, small differences in observed correlation lead a rational agent to infer large differences in the accuracy of s_{jt}.

It will be useful to define the source that maximizes the agent's trust given the values of a_0 and b_0 and holding the value of a^* fixed. Recall that a combination $\left(\alpha_j, \beta_j\right)$ is feasible if $\alpha_j^2 + \beta_j^2 \leq 1$. It is straightforward to show that the unique trust-maximizing position is the one given in the following definition.[5]

[5] Note also that the trust-maximizing position does not depend on the value of a^*, and it can be equivalently defined as the source for which the agent's trust is weakly higher than her trust for any other feasible source whenever both are present.

Figure 1: Iso-trust curves

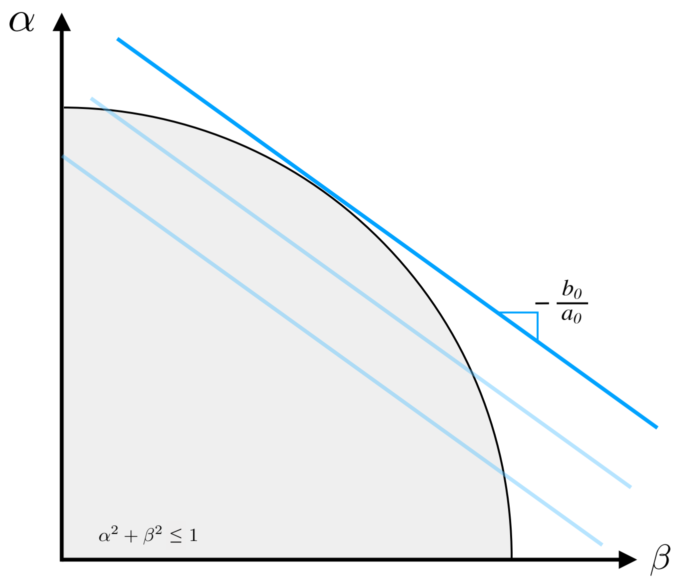

Definition. The agent's *trust-maximizing source* is one with accuracy and slant

$$(\alpha^{max}, \beta^{max}) = \left(\frac{a_0}{\sqrt{a_0^2 + b_0^2}}, \frac{b_0}{\sqrt{a_0^2 + b_0^2}} \right).$$

Figure 1 provides a graphical illustration of the forces that determine trust in our model. The gray shaded area shows the set of all feasible signals—i.e., the (α, β) satisfying the constraint that $\alpha^2 + \beta^2 \leq 1$. The curved boundary of this area is the set of frontier sources which have maximum possible accuracy given their bias. The blue lines in the figure plot the set of *iso-trust curves*: combinations of α and β that yield the same level of asymptotic trust. The slope of these lines is $-\frac{b_0}{a_0}$. Sources that fall on higher iso-trust curves are trusted more.

From this graphical analysis, it is immediately apparent that the trust-maximizing source $(\alpha^{max}, \beta^{max})$ will be the point on the frontier tangent to the iso-trust curves. For an unbiased agent ($b_0 = 0$), the curves are horizontal and this point will lie on the y axis—such agents' trust will be maximized by an unbiased source with accuracy $\alpha = 1$. As bias increases the trust-maximizing source shifts to the right as the agent effectively trades off accuracy in favor of bias. If b_0 is sufficiently high relative to a_0, the agent will prefer a source with bias close to one and accuracy close to zero—a source that essentially just reports the ideological state r_t. Here again we can see the amplification mechanism at work: when a_0 is small, small changes in bias b_0 translate into large changes in the agent's trust-maximizing source.

We can now derive the way trust depends on the bias of the sources. Fix any positive level of bias $b_0 > 0$ and suppose at least one source with $\alpha_0 > 0$ is located on the frontier. It follows from Corollary 1 that $a^* = \sqrt{a_0^2 + b_0^2} > a_0$. For an unbiased source ($\beta_{0j} = 0$), we have $\alpha_j^* = \alpha_{0j} a_0 / \sqrt{a_0^2 + b_0^2}$, and so the agent's trust underestimates the source's true accuracy and is decreasing in the agent's own bias. For a like-minded source ($\beta_{0j} > 0$), Remark 2 implies that her trust will be greater than her trust in an equally accurate unbiased source. Provided that a_0 is sufficiently small and the source's bias is sufficiently large, the agent will over-estimate this source's accuracy in the sense that $\alpha_j^* > \alpha_{0j}$. For an opposite-minded source ($\beta_{0j} < 0$), the agent's trust will be less than in an equally accurate unbiased source, and provided that a_0 is sufficiently small, the agent will believe the source is perverse in the sense that $\alpha_j^* < 0$.

Finally, suppose that at least one source in the market is the agent's trust-maximizing source. Plugging her confidence $a^* = \sqrt{a_0^2 + b_0^2}$ and the values of $(\alpha^{max}, \beta^{max})$ into Remark 2 shows that her confidence in a trust maximizing source will in fact be $\alpha_j^* = 1$. She will come to believe that a trust maximizing source is *perfectly* accurate, reporting $s_{jt} = \omega_t$ with probability one.

Corollary 2. *Suppose the agent has bias $b_0 > 0$ and at least one source with $\alpha_{0j} > 0$ is located on the frontier. The agent's trust in unbiased sources will be less than their true accuracies ($\alpha_j^* < \alpha_{0j}$). Provided a_0 is sufficiently small, the agent will overestimate the accuracy of sufficiently biased like-minded sources ($\alpha_j^* > \alpha_{0j}$), and will believe sources with opposite-minded biases are perverse ($\alpha_j^* < 0$). Regardless of the value of a_0, she will believe a trust-maximizing source is perfectly accurate ($\alpha_j^* = 1$).*

4.3 Ideology and Perceived Bias

As a consequence of bias, the agent will also come to believe that ω_t is correlated with r_t, even though they are in fact uncorrelated. That is, the agent's limiting posterior is $\gamma^* \neq 0$, even though $\gamma_0 = 0$. To see this, note that in the agent's model, $\gamma^* = \rho_{xr}/a^*$. If the agent has bias $b_0 > 0$, the true correlation between x_t and r_t is b_0 (recalling again that we focus on the case $\gamma_0 = 0$). Therefore, we have that $\gamma^* = b_0/a^* > 0$. The agent becomes right-leaning in her ideology, in the sense that she believes the truth is correlated with the conservative position on issues.

Corollary 3. *Suppose the agent has bias $b_0 > 0$. Then her ideology is right-leaning with $\gamma^* = b_0/a^* > 0$.*

We can then apply Proposition 3 to derive the agent's limiting posterior on source j's bias as a function of her confidence a^*. The following corollary follows immediately.

Corollary 4. *Suppose the agent has bias $b_0 > 0$. The agent perceives an unbiased source with $\alpha_{0j} > 0$ as oppositely biased ($\beta_j^* < 0$). She also perceives a like-minded biased source with $\alpha_{0j} > 0$ as less right-biased than it actually is ($\beta_j^* < \beta_{0j}$).*

4.4 Accuracy

We now turn to the accuracy of the agent's asymptotic beliefs about ω_t. We focus for simplicity on the agent's beliefs in periods where (x_t, r_t) is not observed, so that bias only influences accuracy via its influence on trust, and we focus on the limiting distribution of beliefs characterized by Proposition 4. We denote the source the agent observes in the period of interest by k to distinguish it from the generic index j.

To quantify the distance between the agent's posterior and the true states, we will define the agent's limiting *mean squared error (MSE)* to be $\phi^* = \mathrm{E}\left[(\overline{\omega}_t - \omega_t)^2\right]$, where the expectation is taken over the joint distribution of $(\overline{\omega}_t, \omega_t)$ given by Proposition 4. Her posterior mean will be given by $\overline{\omega}_t = \alpha_k^* (\alpha_{0k}\omega_t + \beta_{0k} r_t + \varepsilon_{kt})$. Recalling that we focus on the case where ω_t and r_t are orthogonal ($\gamma_0 = 0$), this implies that her MSE will be $\phi^* = \alpha_k^{*2} - 2\alpha_k^* \alpha_{0k} + 1$. We consider comparative statics with respect to α_{0k} and β_{0k}, assuming that $\alpha_{0k} \geq 0$ and holding the other parameters of the problem constant.

One natural conjecture would be that the agent's beliefs are more accurate when she observes a more accurate source—i.e., MSE is decreasing in α_{0k} holding β_{0k} constant.[6] This turns out to be mostly correct. It is clear in the benchmark case where the agent is unbiased, since then $\alpha_k^* = \alpha_{0k}$ and $\phi^* = 1 - \alpha_{0k}^2$. Plugging in the value of α_k^* from Remark 2 shows that it is true more generally when $\alpha_k^* \geq 0$.

[6] Note that increasing either α_{0k} or β_{0k} holding the other constant is only possible in the case where k is not on the frontier.

A second conjecture might be that agent's beliefs are more accurate when she observes a less biased source—i.e., MSE is increasing in $|\beta_{0k}|$ holding α_{0k} constant. Here, things turn out to be more subtle. To provide a relatively simple characterization, we focus on the case in which there is at least one frontier source other than source k in the market, which means that the agent's confidence a^* is held constant as we vary β_{0k}. From the expression for ϕ^* above, we have $\frac{\partial \phi^*}{\partial \beta_{0k}} = 2\left(\alpha_k^* - \alpha_{0k}\right)\frac{\partial \alpha_k^*}{\partial \beta_{0k}}$. Changing the source's bias affects MSE only via its effect on the agent's trust α_k^*. Using Remark 2, we can see that $\frac{\partial \alpha_k^*}{\partial \beta_{0k}} \geq 0$ (recalling that we focus on the case $b_0 \geq 0$). If the agent starts out trusting a like-minded source k excessively ($\beta_{0k} \geq 0$ and $\alpha_k^* > \alpha_{0k}$), increasing β_{0k} increases MSE because it exacerbates the agent's excessive trust. If the agent starts out trusting a like-minded source k too little ($\beta_{0k} \geq 0$ and $\alpha_k^* < \alpha_{0k}$), increasing β_{0k} *reduces* MSE, because it mitigates the agent's underestimation of the source. Note that we always have $\alpha_k^* < \alpha_{0k}$ if $\beta_{0k} = 0$, so giving an unbiased source a small like-minded bias always increases accuracy. If the source is opposite-minded ($\beta_{0k} < 0$), increasing $|\beta_{0k}|$ increases MSE, because we know $\alpha_k^* < \alpha_{0k}$ and α_k^* is decreasing in $|\beta_{0k}|$.[7]

A third conjecture might be that the agent's beliefs are more accurate when she herself is less biased—i.e., MSE is increasing in b_0. This case turns out to be more ambiguous. We have $\frac{\partial \phi^*}{\partial b_0} = 2\left(\alpha_k^* - \alpha_{0k}\right)\frac{\partial \alpha_k^*}{\partial b_0}$, so changing bias again only affects MSE via its effect on the agent's trust. We thus have similar ambiguity where the net effect of b_0 on MSE depends on whether the agent starts out trusting source k too much or too little and whether trust increases or decreases with b_0. The added wrinkle is that changing b_0 affects α_k^* both directly through the observed correlation ρ_{xsk} and indirectly via changes in confidence a^*. For a like-minded source with small bias, increasing b_0 tends to decrease α_k^* because the increase in ρ_{xsk} is relatively small compared to the increase in a^*. Since the agent generally trusts such sources too little, the overall effect is an increase in MSE. Similarly, the agent tends to trust highly-biased like-minded sources too much, and increasing b_0 tends to increase this trust even further—again resulting in an increase in MSE. For sources positioned between these two extremes, the net result from the two effects can be ambiguous.

[7] An alternative comparative static is to move source k along the frontier—increasing β_{0k} and decreasing α_{0k} while maintaining $\alpha_{0k}^2 + \beta_{0k}^2 = 1$. Since we know $\partial \phi^* / \partial \alpha_{0k} \leq 0$, increasing $|\beta_{0k}|$ in this way will make the effect on MSE more positive than in the comparative statics above. Thus, when $\alpha_k^* > \alpha_{0k}$ or the source is opposite-minded, increasing $|\beta_{0k}|$ still leads to higher MSE. When $\alpha_k^* < \alpha_{0k}$ for a like-minded source, the comparison becomes ambiguous. However, if we begin with $\beta_{0k} \geq 0$ sufficiently close to zero, it will still be true that increasing β_{0k} reduces MSE, since the change in α_{0k} along the frontier local to $\beta_{0k} = 0$ is close to zero.

Corollary 5. *Suppose the agent has bias $b_0 > 0$. Consider a period in which the agent observes source k with $\alpha_{0k} \geq 0$, and suppose at least one other source is located on the frontier. The agent's MSE is decreasing in the accuracy α_{0k} of source k if $\alpha_k^* \geq 0$. If source k is like-minded ($\beta_{0k} \geq 0$), MSE is increasing in the magnitude of the source's bias β_{0k} if and only if the agent initially trusts source k too much ($\alpha_k^* \geq \alpha_{0k}$). If source k is opposite-minded ($\beta_{0k} < 0$), MSE is increasing in the magnitude of the source's bias $|\beta_{0k}|$. Finally, if a source k is like-minded, MSE is increasing in the agent's bias b_0 if k is sufficiently unbiased (β_{0k}/α_{0k} is close to zero) or highly biased (β_{0k}/α_{0k} is large).*

4.5 Polarization

Finally, we characterize the extent to which agents with opposite biases come to disagree about the value of ω_t asymptotically. We consider an R-biased agent with bias $b > 0$ and a corresponding posterior $\overline{\omega}_t^R$, and an L-biased agent with bias $-b$ and a corresponding posterior $\overline{\omega}_t^L$. We assume the agents' a_0 values are the same.

We again focus on periods in which (x_t, r_t) is not observed and on the limiting distribution of beliefs characterized by Proposition 4. We consider a period where the R-biased agent observes a like-minded source with accuracy $\alpha > 0$ and bias equal to $\beta \geq 0$ and the L-biased agent observes a like-minded source with accuracy α and bias $-\beta$.

We define the agents' *expected disagreement* to be $\pi^* = \mathrm{E}\left[\frac{1}{4}\left(\overline{\omega}_t^R - \overline{\omega}_t^L\right)^2\right]$. Scaling by one fourth here ensures that $\pi^* \in [0,1]$.

Consider, first, the case where the agents observe the same source and it is unbiased ($\beta = 0$). We know that $\overline{\omega}_t^R$ and $\overline{\omega}_t^L$ will be equal to the source's signal s_{jt} multiplied by the agents' trusts α_j^{*R} and α_j^{*L} respectively. Expected disagreement could be positive in this case if the agents differ in trust. If at least one source in the market is on the frontier, however, both agents' trust will be equal to $\alpha_{0j} a_0 / \sqrt{a_0^2 + b^2}$ and so $\pi^* = 0$.

Consider, next, the case where the agents observe distinct biased sources ($\beta > 0$). Here, expected disagreement will generally be positive. Our main question of interest is whether it can be large even when the agent's bias b_0 is small. It turns out that the answer is affirmative whenever the agent observes a trust-maximizing source and a_0 is relatively small. To see this, recall that the bias of a trust-maximizing source is $\beta_{0j} = \beta^{max} = b_0 / \sqrt{a_0^2 + b_0^2}$ and the agent's trust in such a source is $\alpha_j^* = 1$ by Corollary 2. Thus, expected disagreement will be increasing in the trust-maximizing source's bias: $\pi^* = (\beta^{max})^2$. This will be at least $1/2$ provided that $a_0 \leq b_0$, and it will approach one in the limit as a_0 becomes small. Thus, even arbitrarily small biases can lead agents to maximal polarization.[8]

Corollary 6. *Suppose there are agents with opposite biases of magnitude $b > 0$ and same values for a_0. Suppose furthermore that there is least one source with $\alpha_{0j} > 0$ located on the frontier. In any period in which the agents both observe a single unbiased source, expected disagreement will be $\pi^* = 0$. In any period in which they each observe like-minded frontier sources with opposite biases, expected disagreement will be strictly positive. In any period where they each observe their trust-maximizing source, expected disagreement will be at least $\pi^* = \frac{1}{2}$ if $a_0 \leq b$, and will approach $\pi^* = 1$ in the limit as $a_0 \to 0$.*

5 Multi-Homing

A common intuition is that divergent trust and polarization could be reduced or eliminated if agents were exposed to an ideologically diverse set of information sources. In any given period, agents might observe both biased and unbiased sources and so have less extreme beliefs than if they observed their preferred biased source alone. Over time, the ability to compare the reports of different outlets might help them more accurately identify trustworthy sources. In this section, we show that it is possible for multihoming to have beneficial effects consistent with this intuition, but also that this need not be the case. Multi-homing may leave trust and polarization unchanged, or even exacerbate them.

We analyze the beliefs of a multi-homing agent who observes the signals s_t of all J sources in every period. We can derive the multi-homing agent's limiting posterior beliefs θ^* in a similar manner to the single-homing case. The main difference is that the multi-homing agent not only learns the correlations $(\rho_{xs}, \rho_{xr}, \rho_{rs})$, but also the matrix of correlations among the elements of s_t, which we denote as Σ. We focus on the case where Σ is not singular. The observed correlations

[8]Note that the requirement that the source be trust-maximizing is not knife-edge: π^* is continuous in α and β, so the result holds approximately when these are close to the trust-maximizing values.

$R^M = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$ thus place an additional restriction on the agent's identified set, resulting in a tighter lower bound on the accuracy a of the agent's own feedback x_t than in the single-homing case. The exact condition is that a cannot be smaller than $\underline{\underline{a}}^R = \sqrt{\zeta}$, where ζ is the population R^2 from a regression of x_t on r_t and all of the elements of s_t.[9]

Proposition 6. *Suppose the true correlations of the observed data are $R_0^M = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$, where Σ is nonsingular. The limiting posterior μ_∞^* exists. Its support contains a unique value θ^* whose elements are $a^* = \max\left\{a_0, \underline{\underline{a}}^R\right\}$, $\gamma^* = \frac{\rho_{xr}}{a^*}$, $\alpha_j^* = \frac{\rho_{xsj}}{a^*}\ \forall j$, and $\beta_j^* = \frac{1}{\sqrt{1-\gamma^{*2}}}\left(\rho_{rs} - \gamma^* \alpha_j^*\right) \forall j$. In periods where the agent does not observe (x_t, r_t), her posterior belief about ω_t will be a normal distribution with mean $\overline{\omega}_t = \alpha^* \Sigma^{-1} s_t$ and variance $v_t = 1 - \alpha^* \Sigma^{-1} \alpha^*$.*

Proof. See appendix. □

The multi-homing agent's limiting posterior θ^* closely resembles the single-homing agent's. As shown in the above proposition, the multi-homing agent's ideology γ^*, trust α^*, and perceived biases β^* are related to the observed correlations $(\rho_{xs}, \rho_{xr}, \rho_{rs})$ and the agent's confidence a^* in the same way as under single-homing. The multi-homing and single-homing cases only differ if observing multiple sources changes the agent's confidence a^*.

The intuition that multi-homing might reduce divergent trust can be correct in our model. Since the multi-homing bound on confidence $\underline{\underline{a}}^R$ is greater than the single-homing bound $\underline{a}^R$, a multi-homing agent will be weakly more confident than a single-homing agent. This means that the difference in trust $\left|\alpha_j^* - \alpha_k^*\right|$ between any two sources will be weakly smaller under multi-homing, and that ideology γ^* will tend to be less extreme. Thus, multi-homing may dampen divergent trust and ideology (while also increasing overconfidence).

Corollary 7. *The difference in trust $\left|\alpha_j^* - \alpha_k^*\right|$ between any two sources and the magnitude of ideology $|\gamma^*|$ are both weakly smaller under multi-homing.*

While such an effect of multi-homing is possible, it need not be large, and it is possible for the limiting posterior θ^* of a multi-homing agent to be exactly the same as a single-homing agent's. A leading case is when at least one source with $\alpha_{0j} \neq 0$ is located on the frontier. In this case,

[9] That is, $\zeta = \tilde{\rho}' \tilde{\Sigma}^{-1} \tilde{\rho}$ where $\tilde{\rho} = \begin{pmatrix} \rho_{xr} \\ \rho_{xs} \end{pmatrix}$ and $\tilde{\Sigma} = \begin{pmatrix} 1 & \rho_{rs}' \\ \rho_{rs} & \Sigma \end{pmatrix}$.

a^* already achieves its maximal value under single-homing, so the limiting posterior under multi-homing is unchanged.

Corollary 8. *Suppose the agent has bias* $b_0 > 0$ *and at least one source with* $\alpha_{0j} > 0$ *is located on the frontier. Then the multi-homing agent's confidence* a^**, ideology* γ^**, trust* α^**, and perceived biases* β^* *are the same as the single-homing agent.*

We now turn to characterizing disagreement about ω_t in the multi-homing case. Consistent with the intuition above, it is possible for multi-homing to make beliefs more accurate and less polarized. To see this, suppose all sources are on the frontier and consider a period where the single-homing agent observes the most biased source.[10] Then, because Proposition 6 implies that her posterior beliefs will be a weighted average of the signals she observes, her beliefs under multi-homing will be more correlated with ω_t and less correlated with r_t than under single-homing.

In spite of this force being present, we can show that multi-homing does not in general reduce expected disagreement, and may in fact make it worse. We begin with three lemmas that together will provide the proof of our main result in this section.

The first lemma is an immediate implication of the result in Corollary 2 that an agent's trust in her trust maximizing source must be $\alpha_j^* = 1$. Because she believes her trust-maximizing source is perfectly accurate, her posterior belief about ω_t is simply a point mass on the report of her trust-maximizing source.

Lemma 2. *Suppose at least one source is the agent's trust-maximizing source. Then the multi-homing agent's posterior belief about* ω_t *is invariant to what other sources are in the market, and in every period it is degenerate at the value equal to the signal of her trust-maximizing source.*

Proof. See appendix. □

The second lemma shows that observing *any two* distinct sources on the frontier is equivalent for a multi-homing agent to observing her trust-maximizing source. This perhaps surprising result follows from showing that for any two frontier signals s_{jt} and s_{kt}, there exists a linear combination of the two that is equal to the trust-maximizing signal with probability one. This is easy to see in

[10]This would occur, for example, if a_0 is sufficiently small relative to b_0 and the agent observes the source they trust the most asymptotically (i.e., with the highest α_j^*).

the event where j is an unbiased source ($s_{jt} = \omega_t$) and k is a perfectly biased source ($s_{kt} = r_t$); in this case we can take the linear combination $\alpha^{max} s_{jt} + \beta^{max} s_{kt}$. The proof of the lemma shows that such a linear combination always exists. Note that this holds even if the biases of sources j and k are both *opposite* to that of the agent.

Lemma 3. *Suppose the market contains two frontier sources with distinct biases. Then the multi-homing agent's confidence and her posterior belief about ω_t are the same as in the case where the market contains the agent's trust-maximizing source.*

Proof. See appendix. □

The third lemma addresses the case in which sources are not located on the frontier. Here, we focus on the case where the noise components ε_{jt} of the signals s_{jt} are mutually independent. We show that provided the number of sources is large, and provided there is at least a minimal amount of diversity in their slants, this case is also equivalent to observing her trust-maximizing source. We formalize this notion of a "large and diverse" set of sources by considering a large random market as follows.

Definition 1. A sequence of **random markets** is indexed by $J = 1, 2, ..., \infty$. Random market J has J sources, indexed by $j = 1, ..., J$, each with accuracy and bias (α_j, β_j) drawn i.i.d. from some distribution F. The noise components of the sources' signals, i.e., $\{\varepsilon_{jt}\}_{j=1,..,J}$, are mutually independent. Furthermore, under F,

1. Both $\alpha_j \neq 0$ and $\beta_j \neq 0$ have nonzero probability;
2. α_j and β_j are not perfectly correlated; and
3. $\alpha_j^2 + \beta_j^2 < 1$ with probability one.

The proof of the lemma shows that a multi-homing agent in a random market can construct a linear combination of the sources' signals whose value will be close to to the signal of the agent's trust-maximizing source in the limit as the number of sources grow large.

Lemma 4. *Consider any sequence of random markets indexed by $J = 1, 2, ..., \infty$. As $J \to \infty$, the multi-homing agent's confidence converges in probability and her posterior on ω_t converges in distribution to those in the case where the market contains the agent's trust-maximizing source.*

Proof. See appendix. □

Under single-homing, posterior beliefs depend crucially on the observed source in each period. In contrast, under multi-homing the exact composition of the market often has no impact because agents are always able to back out the trust-maximizing signal. The following proposition is then immediate.

Proposition 7. *Suppose one of the following is true: (i) one of the sources is the agent's trust-maximizing source; (ii) the market contains two frontier sources with distinct biases; (iii) we focus on the confidence a^* and the posterior distribution on ω_t corresponding to the probability limit of a sequence of random markets. Then a multi-homing agent with any $b_0 > 0$:*

- *Is overconfident with $a^* = \sqrt{a_0^2 + b_0^2}$*
- *Underestimates the accuracy of unbiased sources*
- *Believes her trust-maximizing source (if it exists) is perfectly accurate*
- *Holds a posterior belief about ω_t in each period degenerate at the value $\alpha^{max}\omega_t + \beta^{max} r_t$.*

If condition (ii) or (iii) holds but not (i), and the accuracies $|\alpha_j|$ are sufficiently large and the biases $|\beta_j|$ are sufficiently small for all sources j, then the agent's error ϕ^* and expected disagreement π^* are greater under multi-homing than under single-homing.

Proof. See appendix. □

6 Extensions

Two extensions to our model examine the implications of ideological bias for media competition and political behavior, respectively. The first shows that media competition can intensify disagreements in a population with biases. The second shows that interpersonal mistrust arises when agents underaccount for both their own and others' biases and results in welfare losses in strategic games of collective decision-making.

6.1 Endogenous Media Slant

Our first extension endogenizes the accuracies and slants of the information sources in a sequential positioning game to explore how media competition affects ideological disagreement.

We consider a unit mass of agents. The agents are divided into three types $\iota \in \{L, U, R\}$. Mass $\mu_R > 0$ are R-types with bias $b > 0$. Mass $\mu_L > 0$ are L-types with bias $-b$. The residual mass $1 - \mu_R - \mu_L > 0$ are U-types with bias equal to 0. All three types share accuracy a_0. Our primary interest remains the case where b and a_0 are both small. We assume that $\mu_L = \mu_R$.

A possibly infinite set of E identical potential entrants sequentially choose whether or not to enter. If they enter, they may choose any accuracy α_j and slant β_j consistent with $\alpha_j^2 + \beta_j^2 \leq 1$. Prior to entry, each entering outlet observes all preceding entrants' choices of (α_j, β_j). We use subgame perfect equilibrium as our solution concept.

We focus on media viewership choices assuming that agents have beliefs about the accuracies of the outlets corresponding to the limiting posterior μ_∞^*.[11] All agents are single-homers who choose a single outlet j to observe in a given period t. We assume that, to maximize utility, an agent always chooses to observe a outlet j for which their trust α_j^* is highest, and randomize with equal probability among the sources that they trust most.[12]

We assume that the revenue of a media outlet is increasing in both the size of its viewership and the trust of its viewers. This is consistent with advertising-supported media where conditional on viewing an outlet a customer spends more time viewing when trust is high. It could also be consistent with paid media where the revenue an outlet can earn from a customer who chooses to view is greater when trust is high.

Let $\alpha_j^{*\iota}$ be ι-type agent's trust for the source j and let $\mathscr{J}_\iota$ be the set of outlets for which an ι-type agent's trust $\alpha_j^{*\iota}$ is highest. Let $\xi\left(\alpha_j^{*\iota}\right)$ denote revenue per viewer of type ι. We assume that $\xi(\cdot)$ is positive, strictly increasing, continuously differentiable, and concave, to capture the idea that firms make additional revenue from higher trust, but with declining marginal revenue.

[11] This can be motivated by considering media viewership choices after a large number of exploration periods, in which beliefs about the sources' accuracies converged to the limit α^*.

[12] Note our focus on the limiting posterior μ_∞^* implies agents hold the continuous and full support priors over (α_j, β_j) from Section 2.2. Strictly speaking, this is a behavioral assumption that the agents' inferences about media outlet accuracy do not condition on the equilibrium strategies chosen by the outlets, but only on the signals the outlets produce.

Firms also pay an entry cost λ. Each firm j thus has expected profit:

$$\Pi_j \quad = \sum_{\iota \in \{L,U,R\}} \mathbf{1}\{j \in \mathscr{J}_\iota\} \frac{\mu_\iota}{|\mathscr{J}_\iota|} \xi\left(\alpha_j^{*\iota}\right) - \lambda,$$

where $\mathbf{1}\{j \in \mathscr{J}_\iota\}$ is an indicator for whether outlet j is in the set of outlets that type-ι agents observe, and $\mu_\iota / |\mathscr{J}_\iota|$ measures the probability of observing j within that set. Note that both $\mathscr{J}_\iota$ and $\alpha_j^{*\iota}$ are equilibrium outcomes that depend on the accuracy and slant choices of all media outlet entrants.

We can now solve for the media outlets' equilibrium choice of accuracies and slant via backward induction. We first consider outcomes in a monopoly market.

Proposition 8. *Suppose there is only one potential entrant ($E = 1$). Then for λ sufficiently low, this firm enters and becomes a monopolist with $\alpha_j = 1$ and $\beta_j = 0$. All biased agents then:*

- *Are overconfident with $a^* = \sqrt{a_0^2 + b^2}$*
- *Trust the monopoly outlet with $\alpha_j^* = \frac{a_0}{\sqrt{a_0^2+b^2}}$*
- *Have expected disagreement $\pi^* = 0$.*

Proof. See Appendix. □

Proposition 8 shows that the monopolist becomes a completely accurate and unbiased source of information. Even though the monopolist has a captive audience, it still seeks to capture rising profits from trust. Since it faces a linear trade-off in trust between the L and R agents when it adds slant, the optimal choice under equal proportions of L and R agents and a concave revenue function ξ is to simply focus on accuracy instead and choose $\alpha_j = 1$ and $\beta_j = 0$. This results in no expected disagreement in the population, as agents observe a common outlet and have a common level of trust. Note that this trust is still suboptimal, however, and so beliefs are less than perfectly accurate. Note also that this result is not knife edge: If the proportions of L and R agents are slightly unequal, the resulting optimal position remains close to unbiased, and confidence, trust, and beliefs remain close to the characterization above.

Turning to the competitive case, we can see from Corollary 2 that sources gain maximum trust from biased agents by choosing those agents' trust-maximizing level of slant. It is then unsurprising that in the case of competition, some sources choose to be biased and successfully retain a large audience.

Proposition 9. *Suppose the set of potential entrants is large ($E = \infty$). Then for λ sufficiently low, all outlets locate at positions on the frontier with $\beta_j \in \{\beta^L, 0, \beta^R\}$, where β^L and β^R are the trust-maximizing slants for type L and type R agents respectively. At least one outlet chooses each of these positions.*

Proof. See Appendix. □

The next Corollary then follows immediately from Corollary 6 and Proposition 7.

Corollary 9. *All biased agents:*

- *Are overconfident with $a^* = \sqrt{a_0^2 + b^2}$*
- *Underestimate the accuracy of unbiased sources*
- *Believe their trust-maximizing sources to be perfectly accurate*
- *Hold a posterior belief about ω_t in each period degenerate at the value $\alpha^{max}\omega_t + \beta^{max} r_t$.*

Furthermore, expected disagreement will be at least $\pi^* = \frac{1}{2}$ if $a_0 \leq b$, and will approach $\pi^* = 1$ in the limit as $a_0 \to 0$. Thus, the entry of partisan media leads to greater divergence in beliefs.

In contrast to the monopoly case, there is now significant disagreement in the population. This stems directly from their complete faith in the accuracy of like-minded outlets and their undivided attention to such outlets. Their beliefs about ω_t are simply degenerate at the signal s_{jt} of their trust-maximizing source. Since such outlets adopt the trust-maximizing slant and this slant approaches ± 1 as the ratio of b to a_0 increases, competition can potentially give rise to maximal disagreement and perfectly negatively correlated beliefs.

6.2 Mistrust of Motives and Partisan Conflict

Our second extension shows how ideological bias leads to mistrust of motives across ideological divides and results in intensified conflict in political settings. This exploration is motivated by studies that show rising numbers of Americans hold negative views towards people on the other side of the partisan divide, for example, seeing them as unintelligent and selfish (Iyengar et al. 2012; Iyengar et al. 2019), with potentially important consequences such as reducing the efficacy of government (Hetherington and Rudolph 2015).

We augment our model by adding an observable policy decision d_t to be made by one of two agents, R and L, and allow for ulterior motives B in decision making. We then characterize the agents' beliefs about the others' motive B when the agents assume both their and others' biases are $b = 0$. We show that agents mistakenly learn that $B \neq 0$ even when in fact $B = 0$.

The setup is as follows. Suppose R and L are multi-homing agents that observe the same set of sources, which includes at least two frontier sources with distinct biases. We assume that agents have beliefs about the sources' accuracies corresponding to the limiting posterior μ^*_∞. After observing the sources' signals in some period t, R makes an observable policy decision d_t to maximize the social welfare function, given by $-(\omega_t - d_t)^2$.

Importantly, we assume that agents fail to appreciate both their own and others' ideological bias and instead believe that $b = 0$ for all agents. Consequently, they believe that others have the same belief about the state ω_t as they do. At the same time, agents entertain the possibility that others may have ulterior motives. Specifically, we assume that L believes that R maximizes $-(\omega_t + B_R r_t - d_t)^2$, where B_R parameterizes R's ulterior motive and may not be equal to zero.

We characterize L's beliefs about R's ulterior motive B_R after observing R's decision d_t.

Proposition 10. *Suppose R and L are multi-homing agents observing the same sources including at least two frontier sources with distinct biases. L assumes that both they and the other agent have biases of zero, but entertains the possibility that R may have an ulterior motive B_R when deciding on d_t. If L observes R's decision d_t in any period when $r_t \neq 0$, then L infers that $B_R = 2\beta^{max} > 0$.*

Proof. See appendix. □

Similarly, if R were to observe L's decisions, R would also conclude that L had an ulterior motive $B_L = -2\beta^{max} < 0$. In other words, mistrust of motives arises when well-meaning agents fail to see how ideological bias colors inference about facts by both themselves and others. The magnitude of mistrust in other's motives is proportional to $\beta^{max} = b/\sqrt{a_0^2 + b^2}$, which is increasing in ideological bias b of the agents.

The political behavior of such well-meaning agents mimics that of self-interested agents with actual conflicts of interest. For example, suppose that the above two agents engage in a contest for the power to decide d_τ for some $\tau > t$ after learning about the bias in each other's preferences. Tullock (1980) provides an elemental model of such a contest. R and L simultaneously invest in arms a_L and a_R to obtain decision-making power, where the probability that R has power to decide d_τ is $a_R/(a_R + a_L)$. In L's eyes, the payoff of obtaining decision-making power is zero when $B_R = 0$, since the two agents would choose the same decision. However, the gain from winning the contest becomes positive if L either perceives R to have a nonzero ulterior motive B_R or believes R's inference about ω_t to be biased. The symmetric Nash equilibrium therefore has positive expenditures on arms, even though in equilibrium the contestants have the same probabilities of winning as if neither had spent anything.

Other types of inefficient strategic behavior also arise from conflicts of interest in elemental game theoretic models of organizational behavior, including costly signaling, signal jamming, obfuscation and uninformative cheap talk (see Gibbons, Matouschek, and Roberts 2013). Ideological differences may therefore lead to welfare losses from uninformative communication across ideological divides, poor decision-making, as well as inefficient expenditures in the battle for power.

7 Conclusion

We present a model in which agents learn about policy-relevant states through observing signals from multiple information sources over time. In the model, agents learn about the accuracies of the information sources based on prior information that they assume to be unbiased. Arbitrarily small biases in the prior information can result in large and persistent divergence in both trust and beliefs about facts. Partisans end up trusting unreliable ideologically aligned sources more than accurate neutral sources, and also becoming overconfident in their own prior information.

Divergent trust and beliefs can arise to a similar extent whether agents selectively view only ideologically aligned sources or are exposed to a diverse range of sources. These predictions of the model are consistent with recent evidence suggesting that the magnitude of selective exposure has generally been limited, even as individuals persistently disagree about both objective facts and which sources can be trusted to provide reliable information about those facts. Moving from a monopoly to a competitive market can deepen rather than mitigate ideological disagreement.

References

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz. 2016. "Fragility of asymptotic agreement under bayesian learning." *Theoretical Economics* 11(1): 187–225.

Andreoni, James, and Tymofiy Mylovanov. 2012. "Diverging Opinions." *American Economic Journal: Microeconomics*, 4 (1): 209-32.

Baliga, Sandeep, Eran Hanany, and Peter Klibanoff. 2013. "Polarization and Ambiguity." *American Economic Review*, 103 (7): 3071-83.

Banerjee, Abhijit. 1992. "A simple model of herd behavior." *Quarterly Journal of Economics* 107(3): 797-817.

Benedictis-Kessner, Justin De, Matthew A Baum, Adam J Berinsky, and Teppei Yamamoto. 2019. "Persuading the enemy: Estimating the persuasive effects of partisan media with the preference-incorporating choice and assignment design," *American Political Science Review*, 113(4): 902-916

Benoit, Jean-Pierre and Juan Dubra. 2019. "Apparent Bias: What does Attitude Polarization Show?" *International Economic Review*, 60: 1675-1703.

Berk, Robert. 1966. "Limiting behavior of posterior distributions when the model is incorrect." *Annals of Mathematical Statistics* 37(1): 51-58.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy* 100(5): 992-1026.

Blackwell, David and Lester Dubins. 1962. "Merging of opinions with increasing information." *Annals of Mathematical Statistics* 33(3): 882-886.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience theory of choice under risk." *Quarterly Journal of Economics,* 127(3): 1243-1285.

Boyd, Stephen and Lieven Vandenberghe. 2004. *Convex Optimization*. New York: Cambridge U Press.

Campante, Filipe and Hojman, Daniel. 2013. "Media and polarization. Evidence from the introduction of broadcast TV in the United States." *Journal of Public Economics* 100: 79–92.

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

Cheng, Ing-Haw and Hsiaw, Alice. 2019. "Distrust in Experts and the Origins of Disagreement." Working paper.

Cotton, John L. 1985. "Cognitive dissonance in selective exposure." In D. Zillmann and J. Bryant eds, *Selective Exposure to Communication*. Hillsdale, JM: Lawrence Erlbaum.

DeGroot, Morris. 1974. "Reaching a consensus." *Journal of American Statistical Association* 69:

118- 121.

DeMarzo, Peter, Dimitri Vayanos, Jeffrey Zweibel. 2003. "Persuasion bias, social influence, and unidimensional opinions." *Quarterly Journal of Economics* 118: 909-968.

Dixit, Avinash K. and Jorgen W. Weibull. 2007. "Political polarization." *Proceedings of the National Academy of Sciences* 104: 7351-7356.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*. Stanford: Stanford University Press.

Flaxman, Seth, Sharad Goel, and Justin M. Rao. 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly* 80(1): 298–320.

Fryer, Roland G., Jr., Philipp Harms, and Matthew O. Jackson. 2019. "Updating beliefs when evidence is open to interpretation: Implications for bias and polarization." *Journal of the European Economic Association*, 17(5):1470-1501.

Gentzkow, Matthew and Jesse M. Shapiro. 2006. "Media bias and reputation." *Journal of Political Economy* 115(2): 280–316.

Gentzkow, Matthew and Jesse M. Shapiro. 2011. "Ideological segregation online and offline." *Quarterly Journal of Economics* 126(4): 1799–1839.

Gentzkow, Matthew, Jesse M. Shapiro and Daniel Stone. 2016 "Media bias in the marketplace: Theory." In S. Anderson, J. Waldfogel, and D. Stromberg eds., *Handbook of Media Economics*. Volume 2. Amsterdam and Oxford: Elsevier.

Gibbons, Robert, Niko Matouschek and John Roberts. 2013. "Decisions in Organizations." Chapter 10 in *Handbook of Organizational Economics* (R. Gibbons and J. Roberts, eds.), Princeton University Press.

Glaeser, Edward L. and Cass R. Sunstein. 2014. "Does More Speech Correct Falsehoods?," *Journal of Legal Studies*, 43 (1), 65–93.

Glaeser, Edward L. and Bryce A. Ward. 2006. "Myths and realities of American political geography." *Journal of Economic Perspectives* 20: 119–144.

Golub, Benjamin and Matthew O. Jackson. 2010. "Naive Learning in social networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics* 2(1): 112-149.

Golub, Benjamin and Matthew O. Jackson. 2012. "How Homophily Affects the Speed of Learning and Best-Response Dynamics." *Quarterly Journal of Economics* 127(3): 1287-1338.

Hetherington, Marc J. and Thomas J. Rudolph. 2015. *Why Washington Won't Work: Polarization,Trust, and the Governing Crisis.* Chicago: University of Chicago Press.

IPCC. 1990. First assessment report. Accessed at <http://www.ipcc.ch/publications_and_data/publications_and_data_reports.shtml> on June 20, 2011.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. Affect, not ideology: A social identityperspective on polarization. *Public Opinion Quarterly*. 76(3): 405–431.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. The origins and consequences of affective polarization in the United States. *Annual Review of Political Science.* 22: 7.1–7.18.

Jo, Donghee. 2020. "Better the Devil You Know: Selective Exposure Alleviates Polarization in an Online Field Experiment." Working paper.

Kartik, Navin, Frances Xu Lee, and Wing Suen. 2020 "Information Validates the Prior: A Theorem on Bayesian Updating and Applications." Working paper.

Keyes, Ralph. 2004. *The Post-Truth Era: Dishonesty and Deception in Contemporary Life.* New York: St Martin's Press.

Kondor, Péter. 2012. "The More We Know about the Fundamental, the Less We Agree on the Price." *The Review of Economic Studies*, 79(3): 1175–1207.

Kunda, Ziva. 1990. "The case for motivated reasoning." *Psychological Bulletin,* 108(3): 480-498.

Levendusky, Matthew S. 2013. "Why Do Partisan Media Polarize Viewers?" *American Journal of Political Science*, 57: 611-623.

Levy, Gilat, and Ronny Razin. 2015. "Correlation Neglect, Voting Behavior, and Information Aggregation." *American Economic Review*, 105 (4): 1634-45.

Lian, Heng. 2009. "On rates of convergence for posterior distributions under misspecification." *Communications in Statistics - Theory and Methods,* 38(11): 1893-1900.

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* 37: 2098–2109.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches.* Cambridge, MA: MIT Press.

Mullainathan, Sendhil and Andrei Shleifer. 2005 "The Market for News." *American Economic Review* 95(4): 1031–1053.

Nisbet, Erik C., Kathryn E. Cooper, and R. Kelly Garrett. 2015. "The Partisan Brain: How Dissonant Science Messages Lead Conservatives and Liberals to (Dis)Trust Science." *The ANNALS of the American Academy of Political and Social Science* 658(1): 36–66.

Ortoleva, Pietro and Erik Snowberg. 2015. "Overconfidence in Political Behavior." *American Economic Review* 105(2): 504–35.

Pew Research Center. 2020. "U.S. Media Polarization and the 2020 Election: A Nation Divided." Accessed at <https://www.journalism.org/2020/01/24/u-s-media-polarization-and-the-2020-election-a-nation-divided//> on June 17, 2020.

Prior, Markus. 2013. "Media and political polarization." *Annual Review of Political Science* 16: 101–27.

Public Policy Polling. 2013. "National Survey Results." Accessed at <http://www.publicpolicypolling.com/pdf/20

on September 28, 2014.

Rabin, Matthew and Joel L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *The Quarterly Journal of Economics*, 114(1): 37–82.

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

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

Suen, Wing. 2004. "The self-perpetuation of biased beliefs." *The Economic Journal*, 114: 377–396.

Thaler, Michael. 2020. "The "Fake News" Effect: Experimentally Identifying Motivated Reasoning Using Trust in News." Working paper.

Tversky, Amos and Daniel Kahneman. 1973. "Availability: A heuristic for judging frequency and probability." *Cognitive Psychology,* 5(2): 207-232.

van der Vaart, A. W. 1998. *Asymptotic Statistics*. New York: Cambridge University Press.

Appendices

A Proofs

A.1 Proof of Proposition 1

We first show that any $\theta \in I(R;b=0)$ satisfies the conditions in the proposition. By definition, $b=0$. Furthermore since $R=R(\theta)$ we have $\alpha=\frac{\rho_{xs}}{a}$, $\gamma=\frac{\rho_{xr}}{a}$, and $\beta=\frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs}-\gamma\alpha)$ by remark 1. Since θ must be feasible, we also have, for all j, $\alpha_j^2+\beta_j^2\leq 1$. Substitution and rearrangement shows that

$$\alpha_j^2+\beta_j^2\leq 1 \iff a^2\geq\zeta_j=\frac{\rho_{xr}^2+\rho_{xsj}^2-2\rho_{rsj}\rho_{xsj}\rho_{xr}}{1-\rho_{rsj}^2}.$$

It is straightforward to verify that ζ_j is the vector of population R^2 from a regression of x_t on r_t and s_{jt}. We conclude that $\theta \in I(R;b=0)$ implies that $a\geq\underline{a}^R=\max_j\sqrt{\zeta_j}$.

For the reverse direction, take θ such that $a\in\left[\underline{a}^R,1\right]$, $b=0$, $\alpha=\frac{\rho_{xs}}{a}$, $\gamma=\frac{\rho_{xr}}{a}$, and $\beta=\frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs}-\gamma\alpha)$. Note that $R(\theta)=R$ by remark 1. Since $a\geq\sqrt{\zeta_j}$ for all j, we also have $\alpha_j^2+\beta_j^2\leq 1$ for all j, so θ is feasible. Together this implies $\theta\in I(R)$. Since additionally $b=0$, $\theta\in I(R;b=0)$.

A.2 Proof of Proposition 2

Let μ_0^R denote the distribution of $R(\theta)$ under the assumption that θ is distributed as μ_0.[13] Further, let $P_{\theta|D_1,\ldots,D_t}$ and $P_{R|D_1,\ldots,D_t}$ denote the posterior distributions of θ and R given the priors μ_0 and μ_0^R and the data $D_1,\ldots,D_t$, respectively. Additionally, let $P_{\theta|R}$ be the conditional distribution of θ given the prior μ_0 and R. That is, for all $\vartheta\in\mathscr{L}_\Theta$, $P_{\theta|R}(\vartheta)=\mu_0(\vartheta|R=R(\theta))$. Note that $D_1,\ldots,D_t$ is independent of θ conditional on R. Thus, we can see that for all $\vartheta\in\mathscr{L}_\Theta$, $P_{\theta|D_1,\ldots,D_t}(\vartheta|R)=$

[13] Formally, let $(\mathscr{R},\mathscr{L}_{\mathscr{R}},\nu_{\mathscr{R}})$ denote the Lebesgue space on $\mathscr{R}$, where $\mathscr{L}_{\mathscr{R}}$ is the σ-algebra of Lebesgue measurable sets, and $\nu_{\mathscr{R}}$ is the corresponding Lebesgue measure. As the mapping $\theta\mapsto R(\theta)$ is measurable, if θ is distributed as μ_0 then $R(\theta)$ is a $(\mathscr{R},\mathscr{L}_{\mathscr{R}})$-valued random variable. It is straightforward to see that, for all $\rho\in\mathscr{L}_R$

$$\mu_0^R(\rho)=\mu_0\left(\bigcup_{R\in\rho}I(R)\right).$$

$P_{\theta|R}(\vartheta)$ and hence that

$$P_{\theta|D_1,\ldots,D_t}(\vartheta) = \int_{\mathscr{R}} P_{\theta|D_1,\ldots,D_t}(\vartheta|R)\, dP_{R|D_1,\ldots,D_t}$$
$$= \int_{\mathscr{R}} P_{\theta|R}(\vartheta)\, dP_{R|D_1,\ldots,D_t}.$$

We first characterize the limit of $P_{R|D_1,\ldots,D_t}$. Begin with the case where the market contains a single source j. Consider only the subset of periods where the agent observes all components of $D_t = \left(s_{jt}, x_t, r_t\right)$ and relabel t to only index such periods. (Note that the remaining periods, where the agent only observes s_{jt} but not (x_t, r_t), provide no information about R.) Let $P_{D|R}$ denote the distribution of D_t conditional on R. The experiment $\left(P_{D|R} : R \in \mathscr{R}\right)$ is Gaussian with known mean zero and known variance one, and its parameter space $\mathscr{R}$ is compact. It is straightforward to verify the following regularity conditions: (i) $P_{D|R} \neq P_{D|R'}$ for any $R \neq R^{'}$; (ii) the mapping $R \mapsto P_{D|R}$ is continuous in total variation norm; (iii) $P_{D|R}$ has a nonsingular information matrix I_{R_0} at R_0 (recalling from Remark 1 we focus on θ_0 such that $R_0 \in \text{int}(\mathscr{R})$); (iv) $\left(P_{D|R} : R \in \mathscr{R}\right)$ is differentiable in quadratic mean at R_0. Then by van der Vaart (1998) Lemma 10.6 and Theorem 10.1 (the Bernstein-von Mises Theorem), the limit of $P_{R|D_1,\ldots,D_t}$ as $t \to \infty$ is a distribution degenerate at the true correlations R_0.

Now consider the general case with multiple sources. Reorder periods so those where the agent observes (s_{1t}, x_t, r_t) occur first, those where the agent observes (s_{2t}, x_t, r_t) occur second, through those where the agent observes (s_{Jt}, x_t, r_t). Denote these respective subsequences of data by $D^1, \ldots, D^J$. Since posterior beliefs are invariant to the order of data, this reordering does not affect the limit of $P_{R|D_1,\ldots,D_t}$. The logic above implies that as $t \to \infty$, the agent's posterior belief $P_{R|D^1}$ at the end of the first set of periods converges to a limit whose marginal distribution on $(\rho_{xs1}, \rho_{xr}, \rho_{rs1})$ is degenerate at the true values of these correlations. Note that for every finite t, $P_{R|D^1}$ has continuous density on $\mathscr{R}$ and so is a valid prior under our model. Applying the same logic again then implies that the agent's posterior belief $P_{R|D^1,D^2}$ at the end of the second set of periods converges to a limit whose marginal distribution on $(\rho_{xs1}, \rho_{xs2}, \rho_{xr}, \rho_{rs1}, \rho_{rs2})$ is degenerate at the true value of these correlations. Iterating this logic repeatedly shows that $P_{R|D^1,\ldots,D^J} = P_{R|D_1,\ldots,D_t}$ converges to a distribution degenerate at the full vector of true correlations R_0.

Finally, note that for all $\vartheta \in \mathscr{L}_\theta$,

$$\begin{aligned}\mu_\infty(\vartheta) &= \lim_{t\to\infty} P_{\theta|D_1,\ldots,D_t}(\vartheta)\\ &= \lim_{t\to\infty}\int_{\mathscr{R}} P_{\theta|R}(\vartheta)\, dP_{R|D_1,\ldots,D_t}\\ &= P_{\theta|R_0}(\vartheta)\\ &= \mu_0(\vartheta|R=R_0)\\ &= \frac{\mu_0(\vartheta\cap I(R_0))}{\mu_0(I(R_0))} = \frac{\mu_0(\vartheta\cap I(R_0:b=0))}{\mu_0(I(R_0:b=0))}.\end{aligned}$$

where the third equality uses the convergence of $P_{R|D_1,\ldots,D_t}$, and the final equality follows from the fact that the support of μ_0 consists only of $\theta\in\Theta$ with $b=0$.

A.3 Proof of Proposition 3

Fix $\vartheta \in \mathscr{L}_\Theta$. Observe that

$$\begin{aligned}\lim_{n\to\infty}\mu_{\infty,n}(\vartheta) &= \lim_{n\to\infty}\frac{\mu_{0,n}(\vartheta\cap I(R_0:b=0))}{\mu_{0,n}(I(R_0:b=0))}\\ &= \frac{\mu_0^*(\vartheta\cap I(R_0:b=0))}{\mu_0^*(I(R_0:b=0))},\end{aligned}$$

where the first equality follows from Proposition 2 and the second proposition follows as $\mu_{0,n}\xrightarrow{d}\mu_0^*$. Therefore, setting $\mu_\infty^*(\vartheta)=\lim_{n\to\infty}\mu_{\infty,n}(\vartheta)$ for all $\vartheta\in\mathscr{L}_\Theta$ we have that the limiting posterior μ_∞^* exists and is well-defined.

Let $a^*=\max\{a_0,\underline{a}^R\}$. By Proposition 1, there is a unique θ^* in $I(R;b=0)$ that minimizes $|a-a_0|$, and the a component of θ^* is a^*. Consider sets $\vartheta',\vartheta''\in\mathscr{L}_\Theta$ such that $\theta^*\in\vartheta'$ but $\theta^*\notin\vartheta''$. As $\{\mu_{0,n}\}_{n=1}^\infty$ becomes concentrated at a_0, $d(\vartheta',a_0)<d(\vartheta'',a_0)$, and so

$$\frac{\mu_\infty^*(\vartheta'')}{\mu_\infty^*(\vartheta')} = \lim_{n\to\infty}\frac{\mu_{0,n}(\vartheta'')}{\mu_{0,n}(\vartheta')}=0.$$

Thus, μ_∞^* is degenerate at θ^*.The characterization of b^*,γ^*,α_j^*, and β_j^* then follow from Proposition 1.

A.4 Proof of Proposition 4

Under the limiting posterior μ_∞^*, state ω_t and signal s_{jt} are joint distributed

$$\begin{pmatrix} \omega_t \\ s_{jt} \end{pmatrix} \sim N\left(0, \begin{bmatrix} 1 & \alpha_j^* \\ \alpha_j^* & 1 \end{bmatrix}\right).$$

By the properties of the multivariate normal distribution, the conditional distribution of ω_t given s_{jt} is therefore $N\left(\alpha_j^* s_{jt}, 1-\alpha_j^{*2}\right)$.

A.5 Proof of Proposition 5

The identified set $I\left(R^{\emptyset}; b=0\right)$ includes all θ such that $\rho_{rs} = \alpha\gamma + \beta\sqrt{1-\gamma^2}$. Let $\hat{j} = \arg\max_j \left|\rho_{rsj}\right|$. To see (i), let $\gamma = \rho_{rs\hat{j}}$, $\alpha = \rho_{rs}/\gamma$, and $\beta = 0$. To see (ii), let $\gamma = \sqrt{1-\rho_{rs\hat{j}}^2}$, $\alpha = 0$, and $\beta = \rho_{rs}/\left|\rho_{rs\hat{j}}\right|$.

A.6 Proof of Lemma 1

From the proof of Proposition 1, we have

$$\zeta_j = \frac{\rho_{xr}^2 + \rho_{xsj}^2 - 2\rho_{rsj}\rho_{xsj}\rho_{xr}}{1-\rho_{rsj}^2}.$$

Plugging in values of ρ_{xr}, ρ_{xsj}, and ρ_{rsj} from Remark 1 and setting $\gamma_0 = 0$ yields $b_0^2 + a_0^2\left(\frac{\alpha_{0j}^2}{1-\beta_{0j}^2}\right)$. The final step follows by noting that $\max_j\left\{b_0^2 + a_0^2\left(\frac{\alpha_{0j}^2}{1-\beta_{0j}^2}\right)\right\} = b_0^2 + a_0^2 \max_j\left\{\frac{\alpha_{0j}^2}{1-\beta_{0j}^2}\right\}$.

A.7 Proof of Proposition 6

Lemma A1. *Let $\alpha \in \mathbb{R}^J$ and let $\Sigma \in \mathbb{R}^J \times \mathbb{R}^J$ be a positive-definite symmetric matrix. $\Sigma - \alpha\alpha'$ is positive semi-definite if and only if $\alpha'\Sigma^{-1}\alpha \leq 1$.*

Proof. Define the symmetric matrix $X = \begin{bmatrix} 1 & \alpha' \\ \alpha & \Sigma \end{bmatrix}$. Note that $\Sigma - \alpha\alpha'$ is the Schur complement of 1 in X and $1 - \alpha'\Sigma^{-1}\alpha$ is the Schur complement of Σ in X. Given that 1 and Σ are both positive definite, standard matrix results (e.g., Boyd and Vandenberghe 2004 Appendix A.5.5) imply that (i) X is positive semi-definite if and only if $\Sigma - \alpha\alpha'$ is positive semi-definite; (ii) X is positive semi-definite if and only if $1-\alpha'\Sigma^{-1}\alpha$ is positive semi-definite. Together, these yield the desired result. □

Lemma A2. *Suppose the true correlations of the observed data are* $R_0^M = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$*, where* Σ *is nonsingular. The identified set* $I\left(R^M; b=0\right)$ *consistent with the multi-homing agent's observed data* R^M *and her belief that* $b=0$ *is non-empty and consists of all* $\theta \in \Theta$ *with* $a \in \left[\underline{\underline{a}}^R, 1\right]$*,* $b=0$*,* $\alpha = \frac{\rho_{xs}}{a}$*,* $\gamma = \frac{\rho_{xr}}{a}$*, and* $\beta = \frac{1}{\sqrt{1-\gamma^2}}\left(\rho_{rs} - \gamma\alpha\right)$*.*

Proof. The vector $\begin{bmatrix} r_t & s_t \end{bmatrix}'$ can be written as $\tilde{\alpha}\omega_t + \xi_t$, where $\tilde{\alpha} = \begin{pmatrix} \gamma & \alpha' \end{pmatrix}'$ and ξ_t is distributed $N(0,\Omega)$ for some positive semi-definite Ω. The covariance matrix of $\begin{bmatrix} r_t & s_t \end{bmatrix}'$ is

$$\tilde{\Sigma} = \begin{bmatrix} 1 & \rho_{rs}' \\ \rho_{rs} & \Sigma \end{bmatrix}.$$

We therefore have $\Omega = \tilde{\Sigma} - \tilde{\alpha}\tilde{\alpha}'$.

Take any $\theta \in I\left(R^M; b=0\right)$. Since Remark 1 applies to the multi-homing case as well, we have $\alpha = \rho_{xs}/a$, $\gamma = \rho_{xr}/a$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}\left(\rho_{rs} - \gamma\alpha\right)$. Since Ω must be positive semi-definite, Lemma A1 implies $\tilde{\alpha}'\tilde{\Sigma}^{-1}\tilde{\alpha} \leq 1$. Plugging the values of γ and α into $\tilde{\alpha}$ then implies $a^2 \geq \tilde{\rho}'\tilde{\Sigma}^{-1}\tilde{\rho} = \zeta$, where $\tilde{\rho} = \begin{pmatrix} \rho_{xr} \\ \rho_{xs} \end{pmatrix}$. Thus, $a \geq \underline{\underline{a}}^R = \sqrt{\zeta}$.

Now take θ such that $a \in \left[\underline{\underline{a}}^R, 1\right]$, $b=0$, $\alpha = \frac{\rho_{xs}}{a}$, $\gamma = \frac{\rho_{xr}}{a}$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}\left(\rho_{rs} - \gamma\alpha\right)$. Note that $R(\theta) = (\rho_{xs}, \rho_{xr}, \rho_{rs})$ by Remark 1. Since $a \geq \sqrt{\zeta} \geq \sqrt{\zeta_j}$ for all j, we also have $\alpha_j^2 + \beta_j^2 \leq 1$ for all j, so θ is feasible. Furthermore, because $a \geq \sqrt{\zeta}$, we have that $\tilde{\alpha}'\tilde{\Sigma}^{-1}\tilde{\alpha} \leq 1$, which implies that $\tilde{\Sigma} - \tilde{\alpha}\tilde{\alpha}' \succeq 0$ is positive semi-definite by Lemma A1. We thus know that these parameters imply a valid covariance matrix for ξ_t as defined above. Together, these facts establish $\theta \in I\left(R^M; b=0\right)$. □

Lemma A2 proves that the identified set $I(R; b=0)$ consistent with the observed data $R = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$ and the agent's belief that $b=0$ is non-empty and consists of all $\theta \in \Theta$ with $a \in \left[\underline{\underline{a}}^R, 1\right]$, $b=0$, $\alpha = \frac{\rho_{xs}}{a}$, $\gamma = \frac{\rho_{xr}}{a}$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}\left(\rho_{rs} - \gamma\alpha\right)$. The steps in the proofs of propositions

2 and 3 then show that the limiting posterior μ_∞^* exists and its support contains a unique value θ^* whose elements are $a^* = \max\left\{a_0, \underline{\underline{a}}^R\right\}$, $\gamma^* = \frac{\rho_{xr}}{a^*}$, $\alpha_j^* = \frac{\rho_{xsj}}{a^*}\,\forall j$, and $\beta_j^* = \frac{1}{\sqrt{1-\gamma^{*2}}}\left(\rho_{rs} - \gamma^*\alpha_j^*\right)\forall j$. Under the limiting posterior μ_∞^*, the state ω_t and the signals s_{jt} are jointly distributed

$$\begin{pmatrix} \omega_t \\ s_t \end{pmatrix} \sim N\left(0, \begin{bmatrix} 1 & \alpha^{*\prime} \\ \alpha^* & \Sigma \end{bmatrix}\right).$$

By the properties of the multivariate normal distribution, the conditional distribution of ω_t given s_t is therefore $N\left(\alpha^{*\prime}\Sigma^{-1}s_t, 1-\alpha^{*\prime}\Sigma^{-1}\alpha^*\right)$.

A.8 Proof of Lemma 2

When at least one source is the agent's trust maximizing source, confidence under single-homing is $\sqrt{a_0^2+b_0^2}$ by Corollary 1. Since $\underline{\underline{a}}^R \geq \underline{a}^R$, confidence a^* under multi-homing must be weakly greater than $\sqrt{a_0^2+b_0^2}$. However, we also have $a^* \leq \underline{\underline{a}}^R \leq \sqrt{a_0^2+b_0^2}$, because r_t and all s_{jt} are linear functions of $\tilde{r}_t$ and ω_t and thus the R^2 from a regression of x_t on r_t and the vector of all s_{jt} cannot exceed the R^2 from a regression of x_t on $\tilde{r}_t$ and ω_t which is $a_0^2+b_0^2$. Thus, $a^* = \sqrt{a_0^2+b_0^2}$. The agent's trust in the trust-maximizing source must therefore be $\alpha_j^* = 1$, and so her posterior belief about ω_t is a point mass at $\overline{\omega}_t = \alpha^{max}\omega_t + \beta^{max}r_t$, by Proposition 6.

A.9 Proof of Lemma 3

Lemma A3. *Suppose a multi-homing agent observes the vector of realizations s_t and neither x_t nor r_t in period t. Under the limiting posterior μ_∞^* associated with $F \in \mathscr{F}$, the conditional distribution of the state ω_t given signals s_t is $N(\overline{\omega}_t, v)$, where*

$$\begin{aligned} \overline{\omega}_t &= \frac{1}{a^*}y'Z\left(Z'Z+K\right)^{-1}\left(Z'w_t+\varepsilon_t\right) \\ v &= 1-\frac{1}{a^{*2}}y'Z\left(Z'Z+K\right)^{-1}Z'y \end{aligned}$$

where $y = \begin{bmatrix} a_0 & b \end{bmatrix}'$, Z is the $2\times J$ matrix where the jth column is $\begin{bmatrix} \alpha_j & \beta_j \end{bmatrix}'$, K is a diagonal matrix such that the jth diagonal is $\kappa_j^2 = 1-\alpha_j^2-\beta_j^2$ and *$w_t = \begin{bmatrix} \omega_t & r_t \end{bmatrix}'$.*

Proof. By Proposition 4, $\overline{\omega}_t = \frac{1}{a^*}\rho_{xs}'\Sigma^{-1}s_t$ and $v = 1 - \frac{1}{a^{*2}}\rho_{xs}'\Sigma^{-1}\rho_{xs}$. In our special parameterization of ideological bias, we have that $\rho_{xs} = Z'y$ and $\Sigma = Z'Z + K$. The lemma follows from noting that $s_t = Z'w_t + \varepsilon_t$. $\square$

When the market contains two frontier sources with distinct biases, confidence under single-homing is $\sqrt{a_0^2 + b_0^2}$ by Corollary 1. By the same logic as the proof of Lemma 2, this implies confidence a^* under multi-homing must be equal to $\sqrt{a_0^2 + b_0^2}$ as well.

First consider the case where there are exactly two sources in the market. Since both are frontier and have distinct biases, $K = 0$ and Z spans $\mathbb{R}^2$, so $Z(Z'Z+K)^{-1}Z' = I$. Lemma A3 then implies that $\overline{\omega}_t = \frac{1}{a^*}y'w_t = \alpha^{max}\omega_t + \beta^{max}r_t$ and $v = 1 - \frac{1}{a^{*2}}y'y = 0$.

Now consider the case with more than two sources. Because we assumed that Σ is nonsingular, there cannot be more than two frontier sources. Without loss of generality, let the frontier sources be s_{1t} and s_{2t}. Note that the elements of $\rho_{xs}'\Sigma^{-1}$ are the coefficients from a population regression of x_t on the elements of s_t. Further note that $x_t = a_0\omega_t + b_0\tilde{r}_t + \eta_t$ where η_t is orthogonal to ε_{jt} for all j, while s_{1t} and s_{2t} are linearly independent linear combinations of ω_t and $\tilde{r}_t$. Thus x_t is orthogonal to s_{jt} conditional on s_{1t} and s_{2t} for all $j \geq 3$. By the Frisch–Waugh–Lovell theorem, the elements of $\rho_{xs}'\Sigma^{-1}$ corresponding to non-frontier sources must be equal to zero, and the elements of $\rho_{xs}'\Sigma^{-1}$ corresponding to frontier sources are the same as in the two source case. We conclude that $\overline{\omega}_t = \alpha^{max}\omega_t + \beta^{max}r_t$ and $v = 0$.

A.10 Proof of Lemma 4

Let $d_{\alpha\alpha} = \frac{1}{J}\sum_{j=1}^{J}\alpha_j^2/\kappa_j^2$, $d_{\beta\beta} = \frac{1}{J}\sum_{j=1}^{J}\beta_j^2/\kappa_j^2$, and $d_{\alpha\beta} = \frac{1}{J}\sum_{j=1}^{J}\alpha_j\beta_j/\kappa_j^2$, where $\kappa_j^2 = 1 - \alpha_j^2 - \beta_j^2$. By Woodbury's matrix identity, we can write $Z(Z'Z+K)^{-1} = (I-R)ZK^{-1}$, where $R = Q(I+Q)^{-1}$ and $Q = ZK^{-1}Z' = J\begin{bmatrix} d_{\alpha\alpha} & d_{\alpha\beta} \\ d_{\alpha\beta} & d_{\beta\beta} \end{bmatrix}$. It is easy to check that

$$R = \begin{bmatrix} \frac{\frac{1}{J}d_{\alpha\alpha} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2} & \frac{-\frac{1}{J}d_{\alpha\beta}}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2} \\ \frac{-\frac{1}{J}d_{\alpha\beta}}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2} & \frac{\frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2} \end{bmatrix}.$$

It is also easy to check that $Z\left(Z'Z+K\right)^{-1}Z' = (I-R)Q = R$. By Definition 1, $\mathrm{E}\left[\alpha_j^2/\kappa_j^2\right]$, $\mathrm{E}\left[\beta_j^2/\kappa_j^2\right]$, and $\mathrm{E}\left[\alpha_j\beta_j/\kappa_j^2\right]$ exist and are finite. Therefore, $d_{\alpha\alpha}$, $d_{\beta\beta}$, and $d_{\alpha\beta}$ converge in probability in the limit as $J \to \infty$ by the weak law of large numbers. Furthermore, Definition 1 implies that α_j and β_j are not linearly dependent, so neither are $\alpha_j/\sqrt{\kappa_j^2}$ and $\beta_j/\sqrt{\kappa_j^2}$. By Cauchy-Schwarz, we have that $\mathrm{E}\left[\alpha_j\beta_j/\kappa_j^2\right]^2 < \mathrm{E}\left[\alpha_j^2/\kappa_j^2\right]\mathrm{E}\left[\beta_j^2/\kappa_j^2\right]$. This implies that plim $\left(d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2\right) > 0$. Therefore, $Z\left(Z'Z+K\right)^{-1}Z' = R \to_p I$.

Further algebraic manipulation shows that

$$y'Z\left(Z'Z+K\right)^{-1}\varepsilon_t = \frac{\left(a_0 d_{\beta\beta} + b_0 d_{\alpha\beta}\right)\left(\frac{1}{J}\sum_{j=1}^{J}\frac{\alpha_j}{\sqrt{\kappa_j^2}}\tilde{\varepsilon}_{jt}\right) + \left(a_0 d_{\alpha\beta} + b_0 d_{\alpha\alpha}\right)\left(\frac{1}{J}\sum_{j=1}^{J}\frac{\beta_j}{\sqrt{\kappa_j^2}}\tilde{\varepsilon}_{jt}\right)}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2} + o_p(1)$$

where $\tilde{\varepsilon}_{jt} = \varepsilon_{jt}/\sqrt{\kappa_j^2} \sim N(0,1)$ are mututally independent across j as well as independent of α_j and β_j. Therefore, by the weak law of large numbers, $y'Z\left(Z'Z+K\right)^{-1}\varepsilon_t \to_p 0$.

Let a_J^* denote the agent's confidence in market J. Let $\overline{\omega}_t^J$ and v^J denote the mean and variance of the agent's distribution on ω_t under her limiting posterior on θ in market J. Note that the R^2 of the population regression of x_t on s_t is $\rho_{xs}'\Sigma^{-1}\rho_{xs}$. Furthermore, $\rho_{xs}'\Sigma^{-1}\rho_{xs} = y'Z\left(Z'Z+K\right)^{-1}Z'y \to_p a_0^2 + b_0^2$. Therefore, the R^2 of the population regression of x_t on s_t and r_t also converges in probability to $a_0^2 + b_0^2$. This implies that $a_J^* \to_p \sqrt{a_0^2 + b^2}$. Her trust is $\rho_{xs}/\text{plim } a_J^*$ by Proposition 3. It follows from Lemma A3 and derivations in the previous paragraph that

$$\overline{\omega}_t^J = \frac{1}{a_J^*}y'Z\left(Z'Z+K\right)^{-1}\left(Z'w_t + \varepsilon_t\right) \to_p \frac{y'w_t}{\text{plim } a_J^*} = \alpha^{max}\omega_t + \beta^{max}r_t$$

and

$$v^J = 1 - \frac{1}{a_J^{*2}}y'Z\left(Z'Z+K\right)^{-1}Z'y \to_p 1 - \frac{y'y}{\text{plim } a_J^{*2}} = 0.$$

Since convergence in probability implies convergence in distribution, we have that $\mathrm{E}f\left(\left(\overline{\omega}_t^J, v^J\right)\right) \to \mathrm{E}f\left(\overline{\omega}_t, 0\right)$ for all bounded, continuous functions f by the portmanteau lemma, where $\overline{\omega}_t = \alpha^{max}\omega_t + \beta^{max}r_t$. It follows that the agent's posterior probability distribution function on ω_t converges pointwise to the posterior in the case where the market contains the agent's trust-maximizing source as $J \to \infty$.

A.11 Proof of Proposition 7

Everything except the two last sentences of the proposition is immediate from lemmas 2, 3, and 4. The final sentence can be proved by noting that as $|\alpha_j| \to 1$ and as $|\beta_j| \to 0$ for all j, the single-homing agent has posterior mean $\overline{\omega}_t \to_p \alpha^{max}\omega_t$, whereas the multihoming agent has $\overline{\omega}_t \to_p \alpha^{max}\omega_t + \beta^{max}r_t$ by lemma 4.

A.12 Proof of Proposition 8

Because all agents will observe the monopolist's signal in every period, the monopolist's profit maximization problem simplifies to choosing accuracy α_0 and bias β_0 to maximize $\Pi = \sum_{\iota\in\{L,U,R\}} \mu_\iota \xi(\alpha^{*\iota}) - \lambda$, where $\alpha^{*\iota}$ is type-ι consumers' trust in the monopolist.

First we show that the optimal location for the monopolist is on the frontier. When agents are not overconfident, $\frac{\partial\alpha^{*\iota}}{\partial\alpha_0} = \frac{a_0}{a^*} > 0$ and it's strictly profitable to increase α_{0j}. When biased agents are overconfident, by Lemma 1 $a_\iota^* = \sqrt{b_{0\iota}^2 + a_{0\iota}^2\left(\frac{\alpha_0^2}{1-\beta_0^2}\right)}$ and the effect of increasing accuracy α_0 on trust is

$$
\begin{aligned}
\frac{\partial\alpha^{*\iota}}{\partial\alpha_0} &= \frac{a_{0\iota}}{a_\iota^*} - (a_{0\iota}\alpha_0 + b_{0\iota}\beta_0)\frac{1}{a_\iota^{*2}}\frac{\partial a_\iota^*}{\partial\alpha_0} \\
&= \frac{a_{0\iota}}{a_\iota^{*3}}\left(b_{0\iota}^2 - \frac{a_{0\iota}\alpha_0 b_{0\iota}\beta_0}{1-\beta_0^2}\right).
\end{aligned}
$$

Unbiased agents are never overconfident and hence $\frac{\partial\alpha^{*U}}{\partial\alpha_0} = 1$ regardless of confidence. The derivative of profit with respect to accuracy when biased agents are overconfident is then

$$
\frac{\partial\Pi}{\partial\alpha_0} = (1-2\mu)\xi'(\alpha^{*U}) + \mu\frac{a_0}{a_\iota^{*3}}\left[\left(\xi'(\alpha^{*R}) + \xi'(\alpha^{*L})\right)b^2 + \left(\xi'(\alpha^{*L}) - \xi'(\alpha^{*R})\right)\frac{a_0\alpha_0 b\beta_0}{1-\beta_0^2}\right],
$$

where we define $\mu = \mu_R = \mu_L$.

When $\beta_0 \geq 0$, we have $\alpha^{*R} \geq \alpha^{*L}$ and so $\xi'(\alpha^{*R}) \leq \xi'(\alpha^{*L})$, since $\xi(\cdot)$ is assumed to be concave. Hence all terms are positive and $\frac{\partial\Pi}{\partial\alpha_0}$ is strictly positive. A similar argument shows that $\frac{\partial\Pi}{\partial\alpha_0}$ is also strictly positive for $\beta_0 < 0$. Thus, the monopolist finds it optimal to increase accuracy until it reaches the frontier.

From Remark 2, we can see that the derivative of a type-ι agent's trust with respect to β_0

along the frontier is:

$$\delta^{\iota}(\beta_0) = \left.\frac{\partial \alpha^{*\iota}}{\partial \beta_0}\right|_{\alpha_0^2+\beta_0^2=1} = \frac{1}{a_\iota^*}\left(b_{0\iota} - a_{0\iota}\frac{\beta_0}{\sqrt{1-\beta_0^2}}\right).$$

The optimal frontier location must satisfy the first order condition that

$$\left.\frac{\partial \Pi}{\partial \beta_0}\right|_{\alpha_0^2+\beta_0^2=1} = (1-2\mu)\,\xi'\left(\alpha^{*U}\right)\delta^U(\beta_0) + \mu\left[\xi'\left(\alpha^{*R}\right)\delta^R(\beta_0) + \xi'\left(\alpha^{*L}\right)\delta^L(\beta_0)\right] = 0. \quad (4)$$

This condition is satisfied at $\beta_0 = 0$ because $\delta^U(0) = 0$, $\alpha^{*R} = \alpha^{*L}$, and $\delta^R(\beta_0) = -\delta^L(\beta_0)$. Recall from earlier when $\beta_0 > 0$, $\xi'\left(\alpha^{*R}\right) \leq \xi'\left(\alpha^{*L}\right)$. Moreover it is straightforward to show that $\delta^R(\beta_0) + \delta^L(\beta_0) < 0$ and $\delta^L(\beta_0), \delta^U(\beta_0) < 0$. Thus the derivative in Equation 4 is strictly negative. Symmetric reasoning shows that this derivative is strictly positive when $\beta_0 < 0$. Thus, the unique solution is for the monopolist to choose $\beta_0 = 0$ and $\alpha_0 = 1$.

Since profits at this position are strictly positive, the monopolist will enter when λ is sufficiently low. The remaining results follow immediately from Propositions 3 and 4.

A.13 Proof of Proposition 9

First, suppose one of the positions $\left\{\beta^L, 0, \beta^R\right\}$ is not occupied by any outlet. Then a potential entrant j can enter into this position and become the unique outlet with maximum trust from the associated type of agent. This outlet will have strictly positive trust and revenue from the associated agent type and so entry will be profitable for sufficiently low λ. Thus, when λ is sufficiently low at least one outlet will enter and occupy each of these positions. Furthermore, when these positions are occupied, an outlet at any other position earns zero revenue and so strictly negative profit for any $\lambda > 0$. Thus, in any equilibrium all entrants must locate at one of these positions.

It remains to show that such an equilibrium exists. The above result reduces the problem to a standard sequential entry game with three possible locations. Let $\Pi_L(J_L, J_U, J_R)$ denote the profit earned by an outlet choosing position β^L in a market with J_L, J_U, J_R firms in the three positions respectively. Let Π_U and Π_R denote similar objects for the other two positions. Any tuple (J_L, J_U, J_R)

of outlets in each of the three positions is an equilibrium if the following conditions hold for Π_L:

$$\begin{aligned}
\Pi_L(J_L, J_U, J_R) &\geq 0 \\
\Pi_L(J_L+1, J_U, J_R) &< 0 \\
\Pi_L(J_L, J_U, J_R) &\geq \Pi_U(J_L-1, J_U+1, J_R) \\
\Pi_L(J_L, J_U, J_R) &\geq \Pi_R(J_L-1, J_U, J_R+1),
\end{aligned}$$

and similar conditions hold for Π_U and Π_R. By symmetry, we focus on the conditions for Π_L above. Since Π_L is strictly decreasing in J_L, there exists J_L^* where the first two conditions hold (and let J_U^* and J_R^* denote the corresponding objects for U and R). For the third condition, at (J_L^*, J_U^*, J_R^*) we have $\Pi_L(J_L^*, J_U^*, J_R^*) \geq 0 > \Pi_L(J_L^*+1, J_U^*, J_R^*)$ from the first two conditions and also $\Pi_U(J_L^*-1, J_U^*, J_R^*) \geq 0 > \Pi_U(J_L^*-1, J_U^*+1, J_R^*)$ by noting that Π_U does not depend on the number of outlets in the L position. Together, these imply $\Pi_L(J_L^*, J_U^*, J_R^*) \geq 0 > \Pi_U(J_L^*-1, J_U^*+1, J_R^*)$, giving us the third condition. The fourth condition follows similarly.

A.14 Proof of Proposition 10

As shown in Proposition 7, R has bias $b > 0$ and holds a posterior belief about ω_t in each period degenerate at the value $\overline{\omega}_t^R = \alpha^{max}\omega_t + \beta^{max}r_t$, while L has the opposite bias and a posterior belief degenerate at $\overline{\omega}_t^L = \alpha^{max}\omega_t - \beta^{max}r_t$. L observes R's policy choice d_t in period t, which is equal to $\overline{\omega}_t^R$, while believing that the true state to be $\overline{\omega}_t^L$. Since L believes R to have the same posterior belief as herself and can infer the true value of r_t from the set of observable signals, L concludes that R has a biased utility function with $B_R = 2\beta^{max}$.