

Ideological Bias and Trust in Information Sources

Matthew Gentzkow*
Stanford and NBER

Michael B. Wong
MIT

Allen T. Zhang
Harvard

August 2021

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 rely on noisy feedback about the state from direct experience to learn the accuracy of sources. Small biases in 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. Disagreements can be similar in magnitude 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, Harvard, the University of Michigan, MIT, Stanford, and Stony Brook 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, Office of Naval Research Army Research Office MURI Grant W911NF-20-1-0252, and the Stanford Cyber Initiative. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

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. However, such accounts generally require that individuals have substantial psychological biases towards cognitive consistency or confirmation (e.g., Lord, Ross and Lepper 1979; Cotton 1985; Rabin and Schrag 1999; Baliga, Hanany, and Klibanoff 2013) or have limited memory or attention (e.g., Fryer, Harms, and Jackson 2019; Che and Mierendorff 2019).

In this paper, we explore a different possibility, which is that rational Bayesian inference may magnify the influence of small biases when agents are uncertain which sources they can trust. Building on insights by Acemoglu, Chernozhukov, and Yildiz (2016) and Sethi and Yildiz (2016), among others, we show that small biases 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 evidence suggesting that the magnitude of selective exposure has generally been limited (Gentzkow and Shapiro 2011; Flaxman, Goel and Roth 2016), 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 effectiveness of masks at stopping disease transmission, in the next period the extent of fraud in a recent election, in a third period the magnitude of global warming due to human activity, and so on. In 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 sources 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, with $r_t = 0$ representing 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 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 an election, this might be procedures at polling places in her neighborhood. In the case of masks, it might be the disease outcomes of friends in her social network. While such feedback will typically be noisy and so not be very informative about ω_t on its own, its key feature is 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 may in fact be 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 $t \rightarrow \infty$. We show that in this case her beliefs about the accuracies of information sources eventually converge to a limiting distribution, the mean of which we define to be the agent’s asymptotic *trust* in the respective sources.

In a benchmark case in which the agent has no ideological bias and there is no uncertainty about the accuracy of the feedback x_t , her trust is a distribution degenerate at 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 in each period.

Introducing 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 valence 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 direct feedback, as well as on those of the observed sources. When agents all observe a common unbiased source, disagreement is generally small even when the agents themselves are biased. When biased sources are introduced to the market,

even small biases on the part of agents can lead to large disagreement.

The final section of our main results considers how these findings 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 model provides a unified framework to explain why large ideological disagreements and divergent trust in information sources can persist despite wide availability of accurate sources and limited information processing biases. Prior work has touched on elements of our model. Berk (1966) provides a general statement that beliefs need not converge in the long run under misspecified learning. Acemoglu, Chernozhukov, and Yildiz (2016) show that arbitrarily small differences in beliefs about the interpretation of signals can generate large disagreements about an underlying state.¹ Our model builds on these insights and links them to the literature on media bias and competition (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.²

¹Blackwell and Dubins (1962) show that Bayesian agents observing signals with increasing information eventually agree on the distribution of future signals. 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. Kartik, Lee, and Suen (2020) show that Bayesians expect disagreements to lessen upon the arrival of any new information. These results contrast 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. Andreoni and Mylovanov (2012), Kondor (2012), Glaeser and Sunstein (2014), and Benoit and Dubra (2019) present models wherein initial differences in the interpretation of signals can generate polarization.

²There is a large empirical literature on the relationship between media markets and political polarization (Glaeser and Ward 2006; McCarty, Poole, and Rosenthal 2006; Campante and Hojman 2013; Prior 2013). A growing number of papers provide experimental evidence on the link between trust in information sources and political beliefs

A key contribution of our theory is to demonstrate how endogenous trust (i.e., Bayesian inference about the accuracy of information sources) can amplify small biases. In contrast to some recent theories, persistent polarization in our model does not arise from non-Bayesian learning rules, such as those with confirmation bias (Rabin and Schrag 1999), ambiguity aversion (Baliga, Hanany, and Klibanoff 2013), or inferential mistakes about the credibility of sources (Cheng and Hsiaw 2019). The small biases in our model affect agents’ perception of the feedback they observe, not how they use this feedback to update their priors. Furthermore, agents in our model can process information from an arbitrarily large set of informative sources. As such, we differ from models where polarization arises from costly information (Suen 2004), limited memory (Fryer, Harms, and Jackson 2019), or limited attention (Che and Mierendorff 2019).

Our work relates to recent literature on learning and disagreement in social networks.³ Sethi and Yildiz (2012; 2016) study disagreement when Bayesian individuals must choose between observing well-informed and well-understood signals. Bowen, Dmitriev, and Galperti (2021) present a result closely related to ours, where belief polarization arises despite abundant information and small misperceptions among Bayesian agents. However, polarization in these models results from selective sharing or ideological segregation, whereas polarization can persist in our model even in the absence of these forces.

The paper proceeds as follows. Section 2 describes the model. Section 3 presents our results on overconfidence and trust. Section 4 presents our results on disagreement. Section 5 presents extensions. Section 6 concludes.

(Levendusky 2013; Nisbet, Cooper, and Garrett 2015; Benedictis-Kessner et al. 2019; Thaler 2020; Jo 2020).

³One 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.

2 Model

2.1 Setup

An agent learns about a sequence of unobservable states ω_t over time periods $t = 1, 2, \dots, \infty$. The conservative position on ω_t (i.e., the realization of ω_t most consistent with a conservative platform) is the ideological valence r_t . There are observable signals s_{jt} from information sources $j = 1, \dots, J$, such as media outlets or talkative neighbors. There is also observable feedback x_t , which represents the agent's information from direct investigation.

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 valence r_t can be written as

$$r_t = \gamma\omega_t + \left(\sqrt{1 - \gamma^2}\right)\tilde{r}_t$$

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 . Allowing r_t to be correlated with ω_t captures the idea that either the conservative or the liberal position on issues may be closer to the truth on average. As discussed in more detail below, we will define an agent's *ideology* to be her belief about this correlation γ .

Each source's signal s_{jt} is given by

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

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 , and ε_{jt} is independent of ω_t and \tilde{r}_t . Since $s_{jt} \sim N(0, 1)$, we have that $\alpha_j^2 + \beta_j^2 \leq 1$ and $\varepsilon_{jt} \sim N\left(0, 1 - \alpha_j^2 - \beta_j^2\right)$. We refer to α_j and β_j as the *accuracy* and *bias* of signal j respectively. Signals are more accurate when they are more correlated with ω_t , and they are more (ideologically) biased when they are more correlated with r_t conditional on ω_t . We let s_t , ε_t , α , and β denote the J -vectors of s_{jt} , ε_{jt} , α_j , and β_j respectively.

Feedback x_t is given by

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

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 , and η_t

is independent of ω_t , \tilde{r}_t , and all ε_{jt} . Importantly, we assume that $a > 0$, so x_t is always positively correlated with ω_t . Since $x_t \sim N(0, 1)$, we have that $a^2 + b^2 \leq 1$ and $\eta_t \sim N(0, 1 - a^2 - b^2)$. We refer to a and b as the agent's *accuracy* and *bias*, paralleling the definition of accuracy and bias of signals above. We are primarily interested in the case where the true value of a is small, so that directly-observed feedback is a noisy signal of the truth, and substantially less informative than what can be learned through media or other sources s_{jt} . We will focus on the case where the agent believes her own feedback to be unbiased, so her prior places probability one on $b = 0$. The only departure from rationality in our framework is that this dogmatic belief may in fact be incorrect.

The parameters of this model are $\theta = (a, b, \alpha, \beta, \gamma)$. The set of all such parameters is denoted as Θ and the Lebesgue space on Θ as $(\Theta, \mathcal{L}_\Theta, \nu)$.⁴ The true value of θ is denoted θ_0 , and where it adds clarity we use a_0 , α_{0j} , b_0 , β_{0j} , and γ_0 to refer to the true values of the individual components. We focus for simplicity on the case in which $\gamma_0 = 0$, so ω_t and r_t are in fact independent.

Example. (Masks) In period t , the state ω_t indexes the effectiveness of masks at preventing the transmission of disease during a pandemic. The ideological valence r_t is negative: conservatives believe masks are relatively ineffective, while liberals believe they are relatively effective. (Note that it is natural to think of the magnitude $|r_t|$ as a finite number—neither group typically claims that masks are *completely* effective or ineffective.) The signals s_{jt} could include reports on news or social media about scientific studies of mask effectiveness, anecdotes where people caught the disease despite wearing a mask, and so on. The feedback x_t could capture individuals' observations of people in their own social networks who either did or did not wear masks and may or may not have caught the disease.

Example. (Election) In period t , the state ω_t indexes the extent of fraud in a recent presidential election. The ideological valence r_t is positive: conservatives believe fraud was high while liberals believe fraud was low. (It is again natural to think of the magnitude $|r_t|$ as a finite number—neither group claims that *all* votes in the election were fraudulent nor that *no* votes in the election were fraudulent.) The signals s_{jt} could include reports on news or social media about scientific studies of voter fraud, court cases challenging election outcomes, claims made by politicians, and so on. The feedback x_t could capture individuals' own observations of the way polling places, mail-in

⁴ Θ is the set of all $\theta \in (0, 1] \times [-1, 1]^{2J+2}$ such that $a^2 + b^2 \leq 1$ and $\alpha_j^2 + \beta_j^2 \leq 1$. \mathcal{L}_Θ is the σ -algebra of Lebesgue measurable sets, and ν is the corresponding Lebesgue measure.

ballots, etc. were administered in their neighborhood.

2.2 The Agent's Problem

We are interested in the agent's beliefs about θ and ω_t in the limit as $t \rightarrow \infty$. We analyze two cases. In the *single-homing* case, the agent observes exactly one source j in each period, and each source is observed infinitely many times as $t \rightarrow \infty$. In the *multi-homing* case, the agent observes all j in each period. In both cases, the agent observes feedback x_t with some probability strictly between zero and one in each period. Whether the agent observes x_t is independent over time.

In some fraction of states, the agent also knows the ideological valence r_t . This captures the idea that the conservative and liberal positions on issues such as the outcome of recent elections are well-known, whereas for less hot-button issues they may not be known. We model this by assuming r_t is observed with a probability strictly between zero and one, drawn independently over time and independent of whether or not x_t is observed. The assumption that r_t is observed is not necessary for our main amplification result, but it allows agents to learn about β and γ over time, leading to dynamics we discuss in Section 3.1.

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 . Our central assumption is that all θ in the support of μ_0 have $b = 0$. To study the case where feedback is noisy and the agent knows this to be the case, we also assume that all θ in the support of μ_0 have $a \in (0, a^{max}]$ for some $a^{max} \in [\sqrt{a_0^2 + b_0^2}, 1]$. In other words, the agent *a priori* believes that her bias b is zero and her accuracy a is not greater than a^{max} . We say that the agent's prior is *tight* if $a^{max} = \sqrt{a_0^2 + b_0^2}$.⁵

We further assume that μ_0 has full support on the subset of Θ consistent with these two restrictions. Full support implies that the agent is *ex ante* uncertain about the values of the elements of θ other than b .

Assumption 1. *The support of μ_0 is the set $\Theta^{prior} \subset \Theta$ for which $b = 0$ and $a \in (0, a^{max}]$ for some $a^{max} \geq \sqrt{a_0^2 + b_0^2}$.*

To build intuition for how beliefs are updated in our model, we briefly consider some simple

⁵The condition that $a^{max} \geq \sqrt{a_0^2 + b_0^2}$ guarantees that there exists θ in the support of the agent's prior that is consistent with observed correlations, as we show formally in Section 3.1.

cases. First consider a single-homing agent who observes a single source j and who never observes feedback x_t or the ideological valence r_t .⁶ Note that the *marginal* distribution of s_{jt} is $N(0, 1)$ regardless of its accuracy, and so the agent’s beliefs about α_j will never change in this case. The mean of the agent’s posterior belief about ω_t after observing s_{jt} in any period t will simply be $\bar{\alpha}_j^0 s_{jt}$, where $\bar{\alpha}_j^0$ is the expectation of the source’s accuracy under her prior μ_0 . If $\bar{\alpha}_j^0$ is zero—the agent thinks the source is as likely to be a perverse “false news” source as an accurate “true news” source—the agent’s posterior belief about ω_t would be zero in all periods, no matter what value of s_{jt} is realized. The agent would therefore learn nothing from observing s_{jt} , even though s_{jt} is in fact positively correlated with ω_t .

Next, suppose that x_t is observed in some periods. Because the agent believes that x_t is positively correlated with ω_t (since $a > 0$) and uncorrelated with \tilde{r}_t (since $b = 0$ with probability one under μ_0), she expects the correlation between x_t and s_{jt} to be increasing in the source’s accuracy α_j . If the agent started with a prior mean $\bar{\alpha}_j^0 = 0$, and subsequently observed a positive correlation between x_t and s_{jt} , she would revise her estimate of the source’s accuracy upward. The mean of the agent’s posterior on ω_t given s_{jt} would in turn become positively correlated with s_{jt} . Feedback x_t thus functions as a reference point that allows the agent to learn the accuracy of source j . Having learned about source j ’s accuracy, the agent is then able to learn about ω_t from its signal s_{jt} .

If the agent knew the true value of a exactly (i.e., the agent had a degenerate prior belief that $a = a_0$) and she was unbiased ($b_0 = 0$), her posterior belief on α_j would converge to the true value α_{0j} as $t \rightarrow \infty$. If α_{0j} was close to one, or if the agent was multi-homing and there was a large set of such sources available, she would be able to learn the value of ω_t in every period in the limit as $t \rightarrow \infty$, even though feedback x_t is noisy and only occasionally observed, and the accuracies of the sources are *ex ante* unknown.

In our main case of interest, the true value of a is unknown and the agent’s bias b may be nonzero. The former will mean that the true parameters of the model will only be set identified rather than point identified from the agent’s point of view. The latter will mean that the identified set that the agent’s beliefs converge to need not include the true value θ_0 , and the agent may end up with mistaken beliefs about sources’ accuracies α_j and the states of the world ω_t . As we shall

⁶If x_t was not observed but r_t was, the agent would learn about the source’s bias β_j over time. This might in turn affect her beliefs about α_j if α_j and β_j are correlated under μ_0 . If α_j and β_j were independent under μ_0 , the agent would still never update her beliefs about α_j so long as the agent never observes x_t , whether r_t is observed or not.

show, these distortions can be large even when the magnitude of the agent’s bias is small.

2.3 Discussion: Feedback and Bias

What does feedback x_t represent in the real world? 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 mask behavior or election procedures as in the examples above. It could capture weather events in the agent’s locality (when ω_t relates to global warming), 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 two conditions. First, the agent believes that it is free from ideological bias ($b = 0$) with probability one. Second, contrary to the agent’s belief, x_t may in fact be subject to ideological bias ($b_0 \neq 0$). There are a large number of well-studied psychological phenomena that could provide a micro-foundation for this mistake. 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 would never be able reject that any source is either perfectly positively correlated, uncorrelated, or perfectly negatively correlated with the state. We make this precise in Proposition 7 below.

3 Learning about θ

We now derive the agent's beliefs about θ in the limit as $t \rightarrow \infty$. We begin by describing how the agent learns about θ , then derive comparative statics for the agent's asymptotic beliefs about θ as a function of the true parameters θ_0 .

3.1 Asymptotic Beliefs

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.

In the *single-homing* case, the agent observes a single source in each period and each source is observed infinitely many times in the limit as $t \rightarrow \infty$. Since 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 .

In the *multi-homing* case, the agent observes all sources in every period. The agent thus additionally learns the correlation among the elements of s_t as $t \rightarrow \infty$. The distribution of observable data can be parameterized (with slight abuse of notation) by $R = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$, where Σ is the matrix of correlations among elements of s_t .

We let $R(\theta)$ denote the correlations implied by θ and we let $\mathcal{R} = \{R : R = R(\theta), \theta \in \Theta\}$.

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

$$\begin{aligned}\rho_{xs} &= a\alpha + b\beta \\ \rho_{xr} &= a\gamma + b\sqrt{1 - \gamma^2} \\ \rho_{rs} &= \alpha\gamma + \beta\sqrt{1 - \gamma^2}.\end{aligned}$$

According to the agent's model of the world wherein $b = 0$, the agent believes $\rho_{xs} = a\alpha$, $\rho_{xr} = a\gamma$, and $\rho_{rs} = \alpha\gamma + \beta\sqrt{1 - \gamma^2}$. By contrast, under the true model of the world wherein b_0 may not be zero and we focus on our benchmark case of $\gamma_0 = 0$, we have that $\rho_{xs} = a_0\alpha_0 + b_0\beta_0$, $\rho_{xr} = b_0$, and $\rho_{rs} = \beta_0$. To avoid singular covariance matrices, we assume that none of x_t , s_t , and r_t are perfectly correlated with each other, so the vector of true correlations $R \in \text{int}(\mathcal{R})$.

Consider first the single-homing case. Given that there are $2J + 3$ parameters and $2J + 1$ observed moments, θ will be partially identified by the data. Let $I(R) = \{\theta \in \Theta^{\text{prior}} : R(\theta) = R\}$ denote the *identified set* of parameters in the support of the agent's prior that is consistent with correlations R . This set is always non-empty when $a^{\text{max}} = 1$, since for any $R \in \mathcal{R}$ we can choose parameters $a = 1$, $b = 0$, $\alpha = \rho_{xs}$, $\gamma = \rho_{xr}$, and $\beta = \frac{1}{\sqrt{1 - \gamma^2}}(\rho_{rs} - \alpha\gamma)$. It will also typically include elements with $a < 1$. We show in Lemma 1 below that it is in fact non-empty for any a^{max} consistent with Assumption 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 between x_t and 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} .⁷

Inference for the multi-homing agent proceeds 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 . The observed correlations $R = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$ 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. We focus on the case where Σ is not singular. Then the exact condition is that a cannot be smaller than $\underline{a}^R = \sqrt{\zeta}$, where ζ is the population R^2 from a regression of x_t on r_t and all of the elements of s_t .⁸ Other than the tighter lower bound on a , the multi-homing agent's identified set is the same as the single-homing agent's. In either case, we have that $\underline{a}^R \leq \sqrt{a_0^2 + b_0^2} \leq a^{max}$. Hence $I(R)$ is always nonempty, as shown in the following Lemma.

Lemma 1. *The identified set $I(R)$ consistent with the observed data R and the agent's prior μ_0 is non-empty and consists of all $\theta \in \Theta$ with $a \in [\underline{a}^R, a^{max}]$, $b = 0$, $\alpha = \frac{\rho_{xs}}{a}$, $\gamma = \frac{\rho_{xr}}{a}$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$. Under single-homing, $\underline{a}^R = \max_j \sqrt{\zeta_j}$, where ζ_j is the population R^2 from a regression of x_t on r_t and s_{jt} . Under multi-homing, $\underline{a}^R = \sqrt{\zeta}$, where ζ is the population R^2 from a regression of x_t on r_t and the elements of s_t .*

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, Lemma 1 shows that the data the agent observes can always be rationalized by some $\theta \in \Theta^{prior}$ that *does* fall within the

⁷Put differently, a must be sufficiently large so that the inferred correlation matrix among the variables $(x_t, r_t, s_t, \omega_t)$ is positive semi-definite.

⁸That is, $\zeta = \tilde{\rho}'\tilde{\Sigma}^{-1}\tilde{\rho}$ where $\tilde{\rho} = \begin{pmatrix} \rho_{xr} \\ \rho_{xs} \end{pmatrix}$ and $\tilde{\Sigma} = \begin{pmatrix} 1 & \rho'_{rs} \\ \rho_{rs} & \Sigma \end{pmatrix}$.

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.

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 an 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.

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.

Lemma 2. *Suppose the true correlations of the observed data are $R \in \mathcal{R}$. As $t \rightarrow \infty$, the agent's posterior distribution converges to a limit μ_∞ such that for all measurable $\vartheta \subseteq \mathcal{L}_\Theta$,*

$$\mu_\infty(\vartheta) = \frac{\mu_0(\vartheta \cap I(R))}{\mu_0(I(R))}.$$

Proof. See appendix. □

In the following three subsections, we use Lemmas 1 and 2 to study how the agent's asymptotic beliefs about θ vary with the true parameters θ_0 . We refer to sources that satisfy the condition $\alpha_{0j}^2 + \beta_{0j}^2 \leq 1$ as *feasible*. We refer to sources that satisfy this condition with equality as *frontier sources*. When b_0 and β_{0j} 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*.

3.2 Confidence

We first consider the agent's beliefs about the accuracy of her feedback. These can be interpreted broadly as the agent's confidence in her own information and judgment. A large literature in psychology, economics, and finance has documented overconfidence in many contexts, beginning with the pioneering study of Alpert and Raiffa (1982). Ortoleva and Snowberg (2015) explore

in detail the implications of overconfidence for political behavior.⁹ While overconfidence is a primitive in their model, the results of this section show that it may arise endogenously as a result of small biases in information processing.

We define an agent's *confidence* as the agent's limiting (marginal) posterior distribution on a , which we denote by μ_∞^a . We say that the agent is *overconfident* if $a > a_0$ for all $\theta \in I(R)$. The agent is *underconfident* if $a < a_0$ for all $\theta \in I(R)$. The following result follows from Lemma 1.

Proposition 1. *The agent is never underconfident. The agent is also not overconfident if $b_0 = 0$. However, the agent is overconfident whenever (i) $b_0 \neq 0$ and a_0 is sufficiently small or (ii) $b_0 \neq 0$ and there is some source j such that $|\alpha_{0j}|$ is sufficiently close to one.*

Proof. See appendix. □

Overconfidence emerges in our model in order to reconcile the dogmatic belief that the feedback x_t is unbiased with the observed correlations of the sources. Lemma 1 shows the agent will be overconfident if and only if $\underline{a}^R > a_0$. Under single-homing, 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 . Under multi-homing, this will occur if this will occur if ζ (the population R^2 from a regression of x_t on r_t and the elements of s_t) is greater than a_0^2 . Since it is always possible to rationalize the observed correlations with a sufficiently high value of a , the agent is never underconfident.

The following example shows the conditions under which $\underline{a}^R > a_0$ under single-homing.

Example 1. 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

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

⁹In Ortleva and Snowberg (2015), agents overestimate 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.

The example shows that biased single-homing 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 $b_0 \neq 0$ 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 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. If there is at least one frontier source with $\alpha_j \neq 0$, the agent is overconfident regardless of the value of a_0 , since $\underline{a}^R = \sqrt{a_0^2 + b_0^2}$.

Overconfidence also arises under multi-homing. In fact, since $\zeta \geq \zeta_j$ for all j , the lower bound on a is tighter under multi-homing than under single-homing (whenever the same sources are available to both). Therefore, a multi-homing agent is overconfident whenever a single-homing agent is overconfident.

We can characterize the agent's limiting posterior distribution on a directly. By Lemma 1, each value of a corresponds to a unique value of θ in the agent's identified set $I(R_0)$. Denote this value as $\theta(a) = \left(a, 0, \frac{\rho_{xs}}{a}, \frac{1}{\sqrt{1-\rho_{xr}^2}} \left(\rho_{rs} - \frac{\rho_{xs}\rho_{xr}}{a^2} \right), \frac{\rho_{xr}}{a} \right)$ and let $\vartheta(v) = \{\theta(a) \mid a \in v\}$. Lemma 2 then implies that $\mu_\infty^a(v) = \frac{\mu_0(\vartheta(v))}{\mu_0(\vartheta([\underline{a}^R, a^{max}]))}$.

3.3 Trust

Consider next the agent's beliefs about the accuracy of the sources s_j . A large body of evidence shows divergence between the sources trusted by conservatives and the sources trusted by liberals (e.g., Pew 2014a; Pew 2020), and many have pointed to this as a key factor undermining the media's role in democracy (Gallup and Knight Foundation 2018, 2020). We show in this section that such divergent trust can be a natural implication of ideological bias in information processing, and that this divergence can be large even when the underlying biases are small. In Section 4.2 below, we show how divergent trust can feed back into polarization of beliefs.

We define an agent's *trust* in information source j as the agent's limiting posterior mean on α_j . We can combine Lemmas 1 and 2 with Remark 1 to derive the agent's trust.

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

$$\bar{\alpha}_j = A\rho_{xsj} = A(a_0\alpha_0j + b_0\beta_0j),$$

where the amplification factor A is given by

$$A = \int_{\underline{a}^R}^{a^{\max}} \frac{1}{a} d\mu_{\infty}^a(a), \quad (1)$$

and ρ_{xsj} denotes the j th element of ρ_{xs} .

Remark 2 shows the key amplification mechanism that will drive our main results. When the amplification factor A is large, small differences in biases b_0 and β_0j translate into large differences in trust. The amplification factor A is typically large whenever a_0 and b_0 are small. To see this, suppose that the posterior marginal probability density function on a is nonincreasing. It follows then that $A \rightarrow \infty$ as $\underline{a}^R \rightarrow 0$. In other words, A is large if $\underline{a}^R \leq \sqrt{a_0^2 + b_0^2}$ is small and the agent's posterior places sufficient weight on values of a close to \underline{a}^R .

Remark 2 also shows that the agent will come to trust source j more when their biases are more aligned (i.e., $b_0\beta_0j$ is greater). The model thus predicts an endogenous preference for like-minded information sources.


In fact, the comparison of the agent's trust for different sources generates a complete ordering among all possible information sources. Thus, 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 A fixed. Recall that a combination (α_j, β_j) 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.

Definition 1. 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 defined as the set of *frontier sources* which have

Figure 1: Iso-trust curves



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 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 valence r_j . 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.

To quantify the magnitude of divergence in trust, we suppose that the agent’s prior is tight, so $a^{max} = \sqrt{a_0^2 + b_0^2}$. This case approximates a situation where the agent’s prior on her own accuracy a is concentrated on values that are close to the true value a_0 .¹⁰ When the agent is unbiased (i.e.,

¹⁰We do not focus on the alternative case where the agent’s prior on a is concentrated on values much larger than a_0 for two reasons. First, it implies that the agent’s initial beliefs about a are severely mistaken. Second, if a_0 and b_0 are small, then the agent learns that all available sources are very noisy (i.e., $\bar{\alpha}_j$ is small for all j), and there is limited amplification.

$b_0 = 0$), it is easy to show that with a tight prior, the agent’s trust in any frontier source j is exactly correct (i.e., $\bar{\alpha}_j = \alpha_j$).

Trust becomes highly divergent even when bias $b_0 > 0$ is small, provided that a_0 is also relatively small. To formally show this, we take the limit of a sequence of prior distributions and true parameters such that the agent’s accuracy a_0 approaches zero and $a^{max} = \sqrt{a_0^2 + b_0^2}$ correspondingly approaches b_0 , but all remaining parameters are held constant. In this limit, the agent’s trust $\bar{\alpha}_j$ is equal to the source’s bias β_j . Her trust for an almost perfectly right-biased source (i.e., a source j where $\beta_{0j} \approx 1$) is maximal. By contrast, the same agent’s trust for an almost perfectly left-biased source (i.e., a source j where $\beta_{0j} \approx -1$) is minimal.

Proposition 2. *Suppose the agent has bias $b_0 > 0$. Her trust for a source with more right-bias is greater than her trust for a source with less right-bias but the same accuracy. Her trust for the trust-maximizing source is greater than her trust for any other source. If the agent’s prior is tight, her trust for a frontier source j with $\alpha_{0j} \neq 0$ converges to the source’s bias β_{0j} in the limit as $a_0 \rightarrow 0$.*

Proof. See appendix. □

3.4 Ideology and Perceived Bias

We next consider the agent’s beliefs about the correlation γ between ω_t and r_t . An important feature of our model is that agents may not only become polarized in their beliefs about individual issues (as we show in Section 4.2), but also form an *ex ante* conviction that either the conservative or the liberal point of view on *any* issue is likely to be closer to the truth on average. This conviction is captured in the agent’s limiting belief about γ . It will lead agents to hold ideologically divergent views on issues even in the complete absence of information. A striking feature of the recent history of polarization in the US is that the correlation between citizens’ liberal-conservative ideologies and their views across a range of diverse issues has increased substantially over time (Gentzkow 2016). Our model suggests that such correlated views may arise naturally as a result of diverging beliefs about γ that result from bias in information processing.

The agent’s *ideology* $\bar{\gamma}$ is the agent’s limiting posterior mean on γ , the correlation between ω_t and r_t . We say that the agent’s ideology is *right-leaning* if $\bar{\gamma} > 0$. A right-leaning agent 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 . A left-leaning agent comes to believe the opposite.

As a consequence of bias, ideology will be non-zero even though ω_t and r_t are in fact uncorrelated. To see this, note that $\bar{\gamma} = A\rho_{xr}$. 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 $\bar{\gamma} = Ab_0 > 0$, since $A > 0$.

We also consider agents' beliefs about the biases of the information sources s_j . Beliefs about bias are closely related to trust, but are distinct in our model. Trust may be low because an agent believes a source j to be highly biased (most of the variation in s_{jt} comes from \tilde{r}_t) or because the agent believes the source to be unbiased but noisy (most of the variation in s_{jt} comes from ε_{jt}). Empirical evidence suggests that most Americans in fact perceive media sources they distrust to be systematically biased (Gallup and Knight Foundation 2018, 2020).

We define an agent's *perceived bias* $\bar{\beta}_j$ of information source j as the agent's limiting posterior mean on β_j . We say that an agent perceives a source to be *oppositely biased* if $\text{sign}(\bar{\beta}_j) \neq \text{sign}(\bar{\gamma})$. We say that an agent perceives a source as *less right-biased than it actually is* if $\bar{\beta}_j < \beta_{0j}$. Having derived the posteriors on α and γ , we can apply Lemma 1 to derive the agent's limiting posterior on source j 's bias.

When the agent's prior is tight, her ideology and perceived bias may become highly polarized even when bias is small. If a right-biased agent's accuracy a_0 is sufficiently small relative to b_0 , then she believes that the true state is almost perfectly correlated with r_t (i.e., $\bar{\gamma} \rightarrow 1$) and perceives a perfectly accurate source j (i.e., $\alpha_{0j} = 1$) to be almost perfectly left-biased (i.e. $\bar{\beta}_j \rightarrow -1$). An analogous left-biased agent believes that the true state is almost perfectly correlated with $-r_t$ (i.e., $\bar{\gamma} \rightarrow -1$) and perceives a perfectly accurate source to be almost perfectly right-biased (i.e. $\bar{\beta}_j \rightarrow 1$).

Proposition 3. *Suppose the agent has bias $b_0 > 0$. Then her ideology is right-leaning. Furthermore, she perceives an unbiased source with $\alpha_{0j} > 0$ as oppositely biased. She also perceives a like-minded biased source with $\alpha_{0j} > 0$ as less right-biased than it actually is. If her prior is tight and she observes any frontier source with $\alpha_{0j} > 0$, then in the limit as $a_0 \rightarrow 0$, her ideology converges to one and she perceives a perfectly accurate source to be almost perfectly left-biased.*

Proof. See appendix. □

4 Learning about ω_t

We now turn to the agent’s posterior on ω_t . We focus on characterizing the agent’s posterior expectation of 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 direct feedback x_t nor ideological valence r_t , so we capture only the influence of the information source(s) she observes. We denote this expectation by $\bar{\omega}_t$. The following characterization follows from standard conjugate prior results for the normal distribution.

Lemma 3. *Suppose the agent’s beliefs about θ are the limiting posterior μ_∞ associated with $R \in \mathcal{R}$. In periods where the agent does not observe (x_t, r_t) , and observes source j , the mean of the agent’s posterior on ω_t given s_t is*

$$\bar{\omega}_t = Aw_t, \tag{2}$$

where A is the amplification factor defined in Remark 2, $w_t = \rho_{xsj}s_{jt}$ under single-homing, and $w_t = \rho'_{xs}\Sigma^{-1}s_t$ under multi-homing.

Proof. See appendix. □

Because the average belief $\bar{\omega}_t$ is a linear function of the observed signal(s), the asymptotic distribution of beliefs about ω_t follows immediately from the known distribution of the signal(s). Note that because a biased agent’s limiting posterior may be incorrect, her beliefs need not satisfy the martingale property. As a result, her expected posterior mean on ω_t given the ideological valence 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.

We next consider three special cases wherein the agent has access to sources that provide a relatively “large” quantity of information about both ω_t and r_t . In these cases, the agent’s posterior mean $\bar{\omega}_t$ can be written as an even simpler expression.

The first case considers single-homing agents who observe their trust-maximizing source.

The second case considers multi-homing agents who observe *any two* distinct frontier sources. Recall that frontier sources are defined as sources for which $\alpha_j^2 + \beta_j^2 = 1$. When there are at least two distinct frontier sources, there exists a linear combination of the two that is equal to the signal of a trust-maximizing source with probability one. This is easy to see in 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 $\alpha^{max} s_{jt} + \beta^{max} s_{kt}$. Note that this holds even if the biases of sources j and k are both *opposite* to that of the agent.

The third case considers a scenario 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, the number of sources is large, and there is at least a minimal amount of diversity in their slants. We formalize this notion of a “large and diverse” set of sources by considering a large random market as follows.

Definition 2. A sequence of *random markets* is indexed by $J = 1, 2, \dots, \infty$. Random market J has J sources, indexed by $j = 1, \dots, J$, 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., $\{\varepsilon_{jt}\}_{j=1, \dots, J}$, 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.

As we show in the proof of Lemma 4 below, a multi-homing agent in a random market can also 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.

We define Condition 1 to denote all of the above cases.

Condition 1. Information is *rich* if we consider:

1. Single-homing agents who always observe their trust-maximizing source;
2. Multi-homing agents who observe any two distinct frontier sources;

3. Multi-homing agents in the probability limit of a sequence of random markets.

Lemma 4. *Suppose Condition 1 holds. Then $\underline{a}^R = \sqrt{a_0^2 + b_0^2}$, and the mean of the agent's posterior on ω_t given s_t under μ_∞ is*

$$\bar{\omega}_t = A(a_0\omega_t + b_0r_t),$$

where A is the amplification factor defined in Remark 2.

Proof. See appendix. □

Three implications of Lemma 4 are immediate. First, the agent is overconfident if and only if $b_0 \neq 0$. Second, the agent's posterior mean $\bar{\omega}_t$ is the same in all three scenarios described by Condition 1. Third, the agent is able to learn about ω_t in periods even when there is no direct feedback x_t available. This is possible because the agent has used direct feedback in other periods to learn the accuracies of the available sources.

Recall that the amplification factor A is typically large whenever a_0 and b_0 are small. As previously shown in Remark 2, small differences in observed correlation lead a rational agent to infer large differences in the accuracy of s_{jt} . This in turn means that the agent's posterior mean $\bar{\omega}_t$ is a greatly “amplified” version of what would have been the agent's direct feedback x_t .

To make this intuition precise, we focus for much of the next subsections on the special case where the agent's prior is tight. In this case, the support of the agent's prior on her own accuracy a includes only values that are close to the true value a_0 , namely $(0, \sqrt{a_0^2 + b_0^2}]$. It immediately follows from Lemma 4 that if the agent's prior is tight and Condition 1 holds, then the support of μ_∞ contains a unique value θ^* whose elements are $a^* = \sqrt{a_0^2 + b_0^2}$, $\gamma^* = \frac{\rho_{xr}}{a^*}$, $\alpha_j^* = \frac{\rho_{xsj}}{a^*} \forall j$, and $\beta_j^* = \frac{1}{\sqrt{1-\gamma^{*2}}} (\rho_{rsj} - \gamma^* \alpha_j^*) \forall j$.

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 explore the properties of the agent's posterior mean $\bar{\omega}_t$ in the following three subsections.

4.1 Accuracy

We first show that the agent’s posterior beliefs about ω_t can be far away from the true states even when she has access to a rich set of information sources and the underlying bias is small.

To quantify the distance between the agent’s posterior and the true states, we will define the agent’s *mean squared error (MSE)* to be $\phi = \mathbb{E} \left[\frac{1}{4} (\bar{\omega}_t - \omega_t)^2 \right]$, where the expectation is taken over the joint distribution of $(\bar{\omega}_t, \omega_t)$ as described by equation (2). Scaling by one fourth here ensures that $\phi \in [0, 1]$.

We first ask how the agent’s MSE ϕ changes with the agent’s accuracy a_0 and bias b_0 , holding the amplification factor A fixed.¹¹ We suppose the agent has access to rich information, as formalized in the definition of Condition 1. For tractability, we also assume the following.

Condition 2. a and (α, β, γ) are independently distributed under the agent’s prior μ_0 .

Condition 2 says that the agent has an *a priori* belief that the accuracies or biases of observable external signals (i.e., s_t and r_t) are independent of the accuracy of her direct feedback x_t . In this case, the agent’s posterior on a , denoted by μ_∞^a , is simply the agent’s prior marginal distribution on a except truncated at the lower bound of \underline{a}^R .

When both Conditions 1 and 2 hold, holding A fixed is equivalent to holding \underline{a}^R fixed, which is in turn equivalent to holding $a_0^2 + b_0^2$ fixed. It immediately follows from Lemma 4 that when holding A fixed, a higher a_0 (and correspondingly lower b_0) implies a higher correlation between the agent’s posterior mean $\bar{\omega}_t$ and the true state ω_t , and hence a lower MSE ϕ .

If additionally the agent has a tight prior, then her posterior is degenerate at

$$\bar{\omega}_t = \alpha^{\max} \omega_t + \beta^{\max} r_t. \quad (3)$$

The intuition is as follows. In this case, the agent’s confidence is degenerate at $a^* = \sqrt{a_0^2 + b_0^2}$. The agent thus believes that the trust-maximizing source is perfectly correlated with ω_t . Since Condition 1 guarantees that the trust-maximizing source’s report can be written as a linear function

¹¹Note that μ_∞^a and \underline{a}^R are both functions of a_0 and b_0 . Therefore by equation (1), A is also a function of a_0 and b_0 . We emphasize that the effect of varying a_0 and b_0 can be considered in two parts: the effect of changing A and the effect of varying a_0 for a constant A .

of the elements of s_t , the agent's posterior mean after observing s_t is equal to what her trust-maximizing source would have reported.

The agent learns the true value of ω_t exactly in every period if she is unbiased. To see this, note that unbiasedness implies that $\alpha^{max} = 1$ and $\beta^{max} = 0$. This in turn implies that $\bar{\omega}_t = \omega_t$, as shown in equation (3), so $\phi = 0$.

However, the agent's posterior mean will not be equal to the true value if $b_0 \neq 0$. Equation (3) implies that the correlation of $\bar{\omega}_t$ with ω_t is given by $\alpha^{max} = a_0 / \sqrt{a_0^2 + b_0^2}$. As the magnitude of bias b_0 becomes large relative to a_0 , the agent's posterior mean becomes less correlated with the true value. In this limit where $b_0 \neq 0$ and a_0 approaches zero, there is no correlation between $\bar{\omega}_t$ and ω_t , so $\phi \rightarrow \frac{1}{2}$.

Proposition 4. *Suppose Conditions 1 and 2 hold. Holding A fixed, the agent's MSE ϕ decreases in a_0 and increases in $|b_0|$. Further suppose the agent's prior is tight. If $b_0 = 0$, then $\phi = 0$. If $b_0 \neq 0$, then $\phi \rightarrow \frac{1}{2}$ in the limit as $a_0 \rightarrow 0$.*

Proof. See appendix. □

4.2 Polarization

We now characterize the extent to which agents with opposite biases come to disagree about the value of ω_t . Substantial literatures document large and growing disagreement between Democrats and Republicans on both policy issues (Pew 2014b; Boxell, Gentzkow and Shapiro 2017) and questions of fact (Marietta and Barker 2019). Our model shows that such differences can arise as a result of information processing biases even when accurate information is widely available and agents' only motivation is to learn the truth, and that disagreement can be large even when the underlying bias is small.

We consider an R -biased agent with bias $b > 0$ and a corresponding posterior $\bar{\omega}_t^R$, and an L -biased agent with bias $-b$ and a corresponding posterior $\bar{\omega}_t^L$. We assume the agents' a_0 values are the same and they have the same prior μ_0 which satisfies Condition 2.

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

Lemma 4 implies that, if the agents have access to rich information, as formalized in the definition of Condition 1, then whenever \underline{a}^R is held fixed, a higher a_0 (and correspondingly lower b) results in a lower correlation between a right-biased's agent's posterior mean $\bar{\omega}_t$ and r_t . Thus, holding \underline{a}^R fixed, the agents' expected disagreement π decreases in a_0 and increases in b .

To derive the magnitude of expected disagreement, first consider the case where single-homing agents always observe the same unbiased source (with $\beta_{0j} = 0$). We know that $\bar{\omega}_t^R$ and $\bar{\omega}_t^L$ will be equal to the source's signal s_{jt} multiplied by the agents' trust $\bar{\alpha}_j^R$ and $\bar{\alpha}_j^L$ respectively. Expected disagreement could be positive in this case only if the agents differ in trust. If at least one source in the market is on the frontier and the agents' priors are tight, however, then the agent's confidence is degenerate at $\sqrt{a_0^2 + b^2}$. Thus 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. In particular, suppose that 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 $-\beta$. 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. Again suppose the agents' priors are tight. 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 $\bar{\alpha}_j = 1$. 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 small biases can lead agents to maximal polarization.¹² By Lemma 4, the same is true whenever Condition 1 holds.

Proposition 5. *Suppose Conditions 1 and 2 hold. Holding A fixed, expected disagreement π decreases in a_0 and increases in b . Further suppose the agents' priors are tight. If $b = 0$, then $\pi = 0$. If $b \neq 0$, then π approaches one in the limit as $a_0 \rightarrow 0$.*

Proof. See appendix. □

¹²Note 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.

4.3 Single- vs. 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 subsection, 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.

The intuition that multi-homing might reduce divergence in trust can be correct in our model. Since the multi-homing bound on confidence \underline{a}^R is weakly greater than the single-homing bound \underline{a}^R , a multi-homing agent's confidence will be weakly greater (in an FOSD sense) than a single-homing agent's confidence. This means that the difference in trust $|\bar{\alpha}_j - \bar{\alpha}_k|$ between any two sources will be weakly smaller under multi-homing, and that ideology $\bar{\gamma}$ will tend to be less extreme, provided Condition 2 holds. Thus multi-homing may dampen divergent trust and ideology (while also increasing confidence).

Remark 3. Suppose Condition 2 holds. Then the difference in trust $|\bar{\alpha}_j - \bar{\alpha}_k|$ between any two sources and the ideology $|\bar{\gamma}|$ 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 μ_∞ 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, \underline{a}^R already achieves its maximal value under single-homing, so the limiting posterior under multi-homing is unchanged.

We now turn to comparing expected disagreement about ω_t in the single- and multi-homing cases. 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.¹³ Then, because Lemma 3 implies that her posterior beliefs will be a weighted average of the signals she observes, her

¹³This 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 $\bar{\alpha}_j$).

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, Proposition 6 shows that multi-homing does not in general reduce expected disagreement. In fact, multi-homing may make it worse. The reasoning is as follows. Under single-homing, posterior beliefs depend crucially on the observed source in each period. When single-homing agents observe sources with less bias, there is less expected disagreement. In contrast, a multi-homing agent in a large random market can always construct a linear combination of the sources' signals whose value will be close to the signal of the agent's trust-maximizing source, so reducing the bias of the individual sources does not reduce expected disagreement.

Proposition 6. *Suppose Condition 2 holds. Then expected disagreement is the same when single-homing and multi-homing agents observe their trust-maximizing sources. However, if single-homing and multi-homing agents both observe a large random market, and the accuracies $|\alpha_j|$ are sufficiently large and the biases $|\beta_j|$ are sufficiently small for all sources j , then expected disagreement π is greater under multi-homing than under single-homing.*

Proof. See appendix. □

4.4 Learning without x_t

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 $R^\theta = \rho_{rs}$, since ρ_{xs} and ρ_{xr} are not observed. Consequently, the identified set $I(R^\theta)$ consistent with observed data R^θ contains a wide range of parameter values, including $\alpha_j = 1$ for some source j , or $\alpha_j = -1$ for the same source j . The agent thus cannot rule out the extreme possibilities that any of the sources is perfectly positively correlated, uncorrelated, or perfectly negatively correlated with the true state.

Without direct feedback x_t , the agent's posterior mean on ω_t may always be zero regardless of what signals are available. This occurs whenever the agent's prior are (α, γ) -symmetric, meaning that $\mu_0(\vartheta) = \mu_0(\vartheta')$ for all measurable $\vartheta \subseteq \mathcal{L}_\Theta$, where $\vartheta' = \{(a, -\alpha, b, \beta, -\gamma) \mid (a, \alpha, b, \beta, \gamma) \in \vartheta\}$. Intuitively, the agent's average belief about ω_t does not change after observing s_t if she believes a

priori that the correlation of ω_t with any observable source (i.e. any element of s_t or r_t) is zero in expectation. In this sense, learning about ω_t from s_t requires that the agent has *ex ante* dogmatic faith that the expected correlation of some information source with ω_t is nonzero.

Proposition 7. *If an agent does not observe x_t in any period, but still observes r_t , then the identified set $I(R^0)$ consistent with the observed data R^0 includes θ such that $\alpha_j = z$, for any source j and any $z \in [-1, 1]$. Furthermore, the mean of the agent's posterior on ω_t given s_t under μ_∞ is always zero if the agent's prior is (α, γ) -symmetric.*

Proof. See appendix. □

5 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.

5.1 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 and tight priors. 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 on the frontier (i.e., $\alpha_j^2 + \beta_j^2 = 1$). 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_∞ .¹⁴ 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 is highest, and randomize with equal probability among the sources that they trust most.¹⁵

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.

Since agents' priors are tight and the outlets are located on the frontier, a type- t agent's trust is equal to $\bar{\alpha}_j^t = (a_{0t}\alpha_{0j} + b_{0t}\beta_{0j}) / \sqrt{a_{0t}^2 + b_{0t}^2}$. Let \mathcal{J}_t be the set of outlets for which an t -type agent's trust $\bar{\alpha}_j^t$ is highest. Let $\xi(\bar{\alpha}_j^t)$ denote revenue per viewer of type t . 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. Firms also pay an entry cost λ . Each firm j thus has expected profit:

$$\Pi_j = \sum_{t \in \{L, U, R\}} \mathbf{1}\{j \in \mathcal{J}_t\} \frac{\mu_t}{|\mathcal{J}_t|} \xi(\bar{\alpha}_j^t) - \lambda,$$

where $\mathbf{1}\{j \in \mathcal{J}_t\}$ is an indicator for whether outlet j is in the set of outlets that type- t agents observe, and $\mu_t / |\mathcal{J}_t|$ measures the probability of observing j within that set. Note that both \mathcal{J}_t and $\bar{\alpha}_j^t$ 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 $\beta_j = 0$. Biased agents are overconfident,*

¹⁴This 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 μ_∞^a .

¹⁵Strictly 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.

but 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 Proposition 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. Biased agents are overconfident. 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 $a_0 \rightarrow 0$. Thus, the entry of partisan media leads to greater divergence in beliefs.*

Proof. See appendix. □

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.

5.2 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, Sood and Lelkes 2012; Iyengar et al. 2019), with potentially important consequences such as reducing the efficacy of government (Hetherington and Rudolph 2015).¹⁶

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 agents with tight priors and Condition 1 holds. We assume that the agents' beliefs about the sources' accuracies correspond to the limiting posteriors 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 agents with tight priors and Condition 1 holds. L assumes that she and the other agent both 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 $r_t \neq 0$, and her beliefs correspond to the limiting posterior μ_∞ , 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

¹⁶Relatedly, Ortoleva and Snowberg (2015) and Levy and Razin (2015) explore how correlation neglect and resulting overconfidence impact polarization and political behavior.

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 increasing in ideological bias b of the agents.

The political behavior of well-meaning agents with ideological bias 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” to obtain decision-making power, where the probability that R has power to decide d_τ depends on R ’s stock of arms relative to L ’s. 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.

6 Conclusion

We present a model to explain why individuals persistently disagree about both objective facts and the trustworthiness of information sources. In contrast to recent theories, we assume that agents have Bayesian learning rules and can process information from an arbitrarily large set of high-quality sources. Agents in our model learn about policy-relevant states by observing signals from information sources whose accuracy is *ex ante* uncertain. Agents learn about the accuracy of the sources using noisy feedback from direct experience that they assume to be unbiased.

We show that small biases in this direct feedback 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 direct information. They form a conviction that either the conservative or the liberal point of view is closer to the truth on average, and perceive unbiased sources to be oppositely biased. 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. Moving from a monopoly to a competitive market can deepen rather than mitigate ideological disagreement. Mistrust of motives results and leads to inefficient political outcomes.

Our theory highlights the outsized importance of trust in driving ideological differences in society. If individuals place their faith in inaccurate and biased direct feedback, then they learn to trust biased sources, and hence form biased beliefs about facts. For this reason, ideological disagreement can persist even among otherwise Bayesian agents who can process information about an arbitrarily large set of high-quality sources. Reducing selective exposure may therefore fail to redress political polarization. Targeting the underlying drivers of divergent trust – for example, by providing more accurate feedback about sources’ accuracy, or increasing the prominence of commonly trusted sources – may yield larger gains.

References

- Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz. 2016. “Fragility of asymptotic agreement under bayesian learning.” *Theoretical Economics* 11 (1): 187–225.
- Alpert, Marc and Howard Raiffa. 1982. “A progress report on the training of probability assessors.” In *Judgment under Uncertainty: Heuristics and Biases*, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky, 294-305. Cambridge: Cambridge University Press.
- 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 (4): 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.
- Bowen, T. Renee, Danil Dmitriev, and Simone Galperti. 2021. “Learning from Shared News: When Abundant Information Leads to Belief Polarization.” NBER Working Paper.
- Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. 2017. “Greater Internet Use is Not Associated with Faster Growth in Political Polarization among US Demographic Groups.” *Proceedings of the National Academy of Sciences* 114 (40): 10612-10617.
- Boyd, Stephen and Lieven Vandenbergh. 2004. *Convex Optimization*. New York: Cambridge University Press.
- Campante, Filipe and Hojman, Daniel. 2013. “Media and polarization. Evidence from the intro-

- duction 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 the American Statistical Association* 69 (345): 118-121.
- DeMarzo, Peter M., Dimitri Vayanos, Jeffrey Zwiebel. 2003. “Persuasion bias, social influence, and unidimensional opinions.” *Quarterly Journal of Economics* 118: 909-968.
- Dixit, Avinash K. and Jörgen W. Weibull. 2007. “Political polarization.” *Proceedings of the National Academy of Sciences* 104 (18): 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.
- Gallup and Knight Foundation. 2018. “American views: Trust, media and democracy.” Accessed at <<https://knightfoundation.org/reports/american-views-trust-media-and-democracy/>> on May 2, 2020.
- Gallup and Knight Foundation. 2020. “American Views 2020: Trust, Media and Democracy.” Accessed at <<https://knightfoundation.org/reports/american-views-2020-trust-media-and-democracy/>> on May 2, 2020.
- Gentzkow, Matthew. 2016. “Polarization in 2016.” Toulouse Network for Information Technology white paper.
- Gentzkow, Matthew and Jesse M. Shapiro. 2006. “Media bias and reputation.” *Journal of Political Economy* 114 (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 *Handbook of Media Economics*, edited by Simon Anderson, Joel Waldfogel, and David Strömberg, 623-645. Volume 2. Amsterdam and Oxford: Elsevier.
- Gibbons, Robert, Niko Matouschek and John Roberts. 2013. “Decisions in Organizations.” In

- R. Gibbons and J. Roberts, eds., *Handbook of Organizational Economics*. Princeton and Oxford: Princeton University Press.
- Glaeser, Edward L. and Cass R. Sunstein. 2014. “Does More Speech Correct Falsehoods?” *The 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 (2): 119–144.
- Golub, Benjamin and Matthew O. Jackson. 2010. “Naïve 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.
- Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. “Affect, not ideology: A social identity perspective 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: 129-146.
- 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 (3): 611-623.
- Levy, Gilat, and Ronny Razin. 2015. “Correlation Neglect, Voting Behavior, and Information Aggregation.” *American Economic Review* 105 (4): 1634-1645.
- 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 (11): 2098–2109.

- Marietta, Morgan and David C. Barker. 2019. *One Nation, Two Realities: Dueling Facts in American Democracy*. New York: Oxford University Press.
- McCarty, Nolan M., 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): 504–535.
- Pew Research Center. 2014a. "Political Polarization and Media Habits." Accessed at <www.journalism.org/2014/10/21/political-polarization-media-habits> on May 2, 2021.
- Pew Research Center. 2014b. "Political Polarization in the American Public." Accessed at <www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/> on May 2, 2021.
- 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: 101–127.
- Rabin, Matthew and Joel L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *The Quarterly Journal of Economics* 114 (1): 37–82.
- Rochet, Jean-Charles and Jean Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990-1029.
- 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.
- Tullock, Gordon. 1980. "Efficient Rent Seeking." In *Toward a Theory of the Rent-seeking Society*, edited by James M. Buchanan, Robert D. Tollison, and Gordon Tullock, 97-112. College Station, TX: Texas A&M University Press.
- Tversky, Amos and Daniel Kahneman. 1973. "Availability: a heuristic for judging frequency and

probability.” *Cognitive Psychology* 5 (2): 207-232.

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

Appendices

A Proofs

A.1 Proof of Lemma 1

First consider the multi-homing case. Suppose $\theta \in I(R)$. By definition, $b = 0$. Since $R = R(\theta)$ we also 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. Furthermore, θ must correspond to a well-defined covariance matrix for the unit-normal joint distribution of $(\omega_t, r_t, x_t, s_t)$. It follows that $|a| \leq 1$ (by unit-normality) and the covariance matrix for (ω_t, r_t, s_t) must be positive semi-definite.¹⁷ The latter holds if and only if $1 - \tilde{\alpha}'\tilde{\Sigma}^{-1}\tilde{\alpha} \geq 0$ by standard matrix results (see Boyd and Vandenberghe 2004 Appendix A.5.5). Since $\alpha = \frac{\rho_{xs}}{a}$ and $\gamma = \frac{\rho_{xr}}{a}$, an equivalent condition is that $a^2 \geq \zeta = \tilde{\rho}'\tilde{\Sigma}^{-1}\tilde{\rho}$, where $\tilde{\rho} = \begin{pmatrix} \rho_{xr} & \rho'_{xs} \end{pmatrix}'$. Thus, $a \geq \underline{a}^R = \sqrt{\zeta}$.

Now take θ such that $a \in [\underline{a}^R, a^{max}]$, $b = 0$, $\alpha = \frac{\rho_{xs}}{a}$, $\gamma = \frac{\rho_{xr}}{a}$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$. Note that $R(\theta) = (\rho_{xs}, \rho_{xr}, \rho_{rs}, \Sigma)$ by Remark 1. So it only remains to show that θ corresponds to a well-defined covariance matrix for a unit-normal joint distribution for $(\omega_t, r_t, x_t, s_t)$. Since $b = 0$ implies that x_t is independent of r_t and s_t conditional on ω_t , it suffices to show that $|a| \leq 1$ and the covariance matrix for (ω_t, r_t, s_t) must be positive semi-definite. The latter is guaranteed by $1 - \tilde{\alpha}'\tilde{\Sigma}^{-1}\tilde{\alpha} \geq 0$, which holds since $a^2 \geq \zeta = \tilde{\rho}'\tilde{\Sigma}^{-1}\tilde{\rho}$, $\alpha = \frac{\rho_{xs}}{a}$, and $\gamma = \frac{\rho_{xr}}{a}$. Together, these facts establish $\theta \in I(R)$.

The single-homing case proceeds in a similar manner, except that θ must correspond to a well-defined covariance matrix for the distribution of $(\omega_t, r_t, x_t, s_{jt})$ for each j , since the covariance of s_t is not observed. Therefore, by analogous logic, we must have $\tilde{\alpha}'_j\tilde{\Sigma}_j^{-1}\tilde{\alpha}_j \leq 1$ for all j .¹⁸ This in turn means that for all j ,

$$a^2 \geq \zeta_j = \frac{\rho_{xr}^2 + \rho_{xsj}^2 - 2\rho_{rsj}\rho_{xsj}\rho_{xr}}{1 - \rho_{rsj}^2}, \quad (4)$$

which is equivalent to $a \geq \underline{a}^R = \max_j \sqrt{\zeta_j}$.

¹⁷This covariance matrix is given by $\begin{bmatrix} 1 & \tilde{\alpha}' \\ \tilde{\alpha} & \tilde{\Sigma} \end{bmatrix}$, where $\tilde{\alpha} = \begin{pmatrix} \gamma \\ \alpha \end{pmatrix}$ and $\tilde{\Sigma} = \begin{bmatrix} 1 & \rho'_{rs} \\ \rho_{rs} & \Sigma \end{bmatrix}$.

¹⁸Here we define $\tilde{\alpha}_j = \begin{pmatrix} \gamma \\ \alpha_j \end{pmatrix}$ and $\tilde{\Sigma}_j = \begin{bmatrix} 1 & \rho'_{rsj} \\ \rho_{rsj} & 1 \end{bmatrix}$

A.2 Proof of Lemma 2

Let μ_0^R denote the distribution of $R(\theta)$ under the assumption that θ is distributed as μ_0 .¹⁹ Further, let $P_{\theta|D_1, \dots, D_t}$ and $P_{R|D_1, \dots, D_t}$ denote the posterior distributions of θ and R given the priors μ_0 and μ_0^R and the data D_1, \dots, 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 $\vartheta \in \mathcal{L}_\Theta$, $P_{\theta|R}(\vartheta) = \mu_0(\vartheta | R = R(\theta))$. Note that D_1, \dots, D_t is independent of θ conditional on R . Thus, we can see that for all $\vartheta \in \mathcal{L}_\Theta$, $P_{\theta|D_1, \dots, D_t}(\vartheta | R) = P_{\theta|R}(\vartheta)$ and hence that

$$\begin{aligned} P_{\theta|D_1, \dots, D_t}(\vartheta) &= \int_{\mathcal{R}} P_{\theta|D_1, \dots, D_t}(\vartheta | R) dP_{R|D_1, \dots, D_t} \\ &= \int_{\mathcal{R}} P_{\theta|R}(\vartheta) dP_{R|D_1, \dots, D_t}. \end{aligned}$$

We first characterize the limit of $P_{R|D_1, \dots, D_t}$. Begin with the case where all sources are observed in each period. (This includes the multi-homing case as well as the single-homing case where $J = 1$.) Consider only the subset of periods where the agent observes all components of $D_t = (s_t, x_t, r_t)$ and relabel t to only index such periods. (Note that the remaining periods, where the agent only observes s_t 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 $(P_{D|R} : R \in \mathcal{R})$ is Gaussian with known mean zero and known variance one, and its parameter space \mathcal{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 $R_0 \in \text{int}(\mathcal{R})$); (iv) $(P_{D|R} : R \in \mathcal{R})$ 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 $P_{R|D_1, \dots, D_t}$ as $t \rightarrow \infty$ is a distribution degenerate at the true correlations R_0 .

Now consider the single-homing case with $J > 1$. 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

¹⁹Formally, let $(\mathcal{R}, \mathcal{L}_{\mathcal{R}}, \nu_{\mathcal{R}})$ denote the Lebesgue space on \mathcal{R} , where $\mathcal{L}_{\mathcal{R}}$ is the σ -algebra of Lebesgue measurable sets, and $\nu_{\mathcal{R}}$ 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 $(\mathcal{R}, \mathcal{L}_{\mathcal{R}})$ -valued random variable. It is straightforward to see that, for all $\rho \in \mathcal{L}_{\mathcal{R}}$

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

those where the agent observes (s_{Jt}, x_t, r_t) . Denote these respective subsequences of data by D^1, \dots, D^J . Since posterior beliefs are invariant to the order of data, this reordering does not affect the limit of $P_{R|D_1, \dots, D_t}$. The logic above implies that as $t \rightarrow \infty$, 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 \mathcal{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 $P_{R|D^1, \dots, D^J} = P_{R|D_1, \dots, D_t}$ converges to a distribution degenerate at the full vector of true correlations R_0 .

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

$$\begin{aligned}
\mu_\infty(\vartheta) &= \lim_{t \rightarrow \infty} P_{\theta|D_1, \dots, D_t}(\vartheta) \\
&= \lim_{t \rightarrow \infty} \int_{\mathcal{R}} P_{\theta|R}(\vartheta) dP_{R|D_1, \dots, D_t} \\
&= P_{\theta|R_0}(\vartheta) \\
&= \mu_0(\vartheta | R = R_0) \\
&= \frac{\mu_0(\vartheta \cap I(R_0))}{\mu_0(I(R_0))}.
\end{aligned}$$

where the third equality uses the convergence of $P_{R|D_1, \dots, D_t}$, 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 1

It follows immediately from Lemma 1 that the agent is never underconfident. Now suppose $b_0 = 0$. Since η_t is independent of s_t and $b_0 = 0$, it must be that $\zeta \leq a_0^2$. Furthermore, $\underline{a}_R^2 \leq \zeta$ (with equality in the multi-homing case). Therefore, $\underline{a}_R \leq a_0$, so the agent is not overconfident by Lemma 1. Next consider example 1. Plugging in values of ρ_{xr} , ρ_{xsj} , and ρ_{rsj} from Remark 1 into equation (4) and setting $\gamma_0 = 0$ yields $\zeta_j = b_0^2 + a_0^2 \left(\frac{\alpha_{0j}^2}{1 - \beta_{0j}^2} \right)$. Under both single- and multi-homing, we must have that $\underline{a}_R^2 \geq \max_j \left\{ b_0^2 + a_0^2 \left(\frac{\alpha_{0j}^2}{1 - \beta_{0j}^2} \right) \right\} = b_0^2 + a_0^2 \max_j \left\{ \frac{\alpha_{0j}^2}{1 - \beta_{0j}^2} \right\}$, so the agent is overconfident if there

exists j such that $a_0^2 \leq b_0^2 + a_0^2 \left(\frac{\alpha_{0j}^2}{1 - \beta_{0j}^2} \right)$.

A.4 Proof of Proposition 2

Suppose there are two sources j and j' such that $\alpha_j = \alpha_{j'}$ but $\beta_j > \beta_{j'}$. Since $\rho_{xsj} > \rho_{xsj'}$, it follows immediately from Remark 2 that $\bar{\alpha}_j > \bar{\alpha}_{j'}$. The second result follows since the agent's trust-maximizing source has the highest ρ_{xsj} among all sources. The final sentence follows from observing that when the agent's prior is tight and she observes a frontier source with $\alpha_j \neq 0$, then her confidence is degenerate at $\sqrt{a_0^2 + b_0^2}$. This implies that the agent's limiting posterior on α_j is degenerate at $(a_0\alpha_{0j} + b_0\beta_0) / \sqrt{a_0^2 + b_0^2}$. We obtain the desired result by taking the limit as $a_0 \rightarrow 0$.

A.5 Proof of Proposition 3

By Lemma 1, we have that $a > 0$ and thus $\gamma = \rho_{xr}/a = b_0/a > 0$ for all $\theta \in I(R)$. Also by Lemma 1, we have that for all $\theta \in I(R)$

$$\begin{aligned} \beta_j &= \frac{1}{\sqrt{1 - \gamma^2}} (\rho_{rs} - \gamma\alpha_j) \\ &= \beta_{0j} \sqrt{1 - (b_0/a)^2} - \frac{a_0 b_0 \alpha_{0j}}{a^2 \sqrt{1 - (b_0/a)^2}}. \end{aligned}$$

Since $\sqrt{1 - (b_0/a)^2} < 1$ and $\frac{a_0 b_0 \alpha_{0j}}{a^2 \sqrt{1 - (b_0/a)^2}} > 0$, we have that $\beta_j < \beta_{0j}$ for any source with $\beta_{0j} \geq 0$, as desired. The final sentence follows from observing that when the agent's prior is tight and the agent observes some source j with $\alpha_{0j} > 0$, then the agent's confidence is degenerate at $\sqrt{a_0^2 + b_0^2}$. Thus when $\alpha_{0j} = 1$, β_j converges in probability to -1 by taking the limit as $a_0 \rightarrow 0$.

A.6 Proof of Lemma 3

Given any value of θ , the state ω_t and the signals s_{jt} are jointly distributed

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

Under single-homing, the conditional expectation of ω_t given s_{jt} is then $\alpha_j s_{jt}$, by the properties of the multivariate normal distribution. Similarly, the conditional expectation of ω_t given s_t is then $\alpha' \Sigma^{-1} s_t$ under multi-homing. By Lemma 1, we have that $\alpha = \rho_{xs}/a$ for all $\theta \in I(R_0)$. The desired result is then immediate from applying Lemma 2.

A.7 Proof of Lemma 4

It suffices to show that $w_t = a_0 \omega_t + b_0 r_t$ and $\underline{a}^R = \sqrt{a_0^2 + b_0^2}$ in each of the three cases in (*). We can then apply Lemma 3 to show the desired result.

Case 1: Single-homing agent observing their trust maximizing source. This case is straightforward. Example 1 shows that $\underline{a}^R = \sqrt{a_0^2 + b_0^2}$. Furthermore, $\rho_{xsj} = \sqrt{a_0^2 + b_0^2}$, so $w_t = \rho_{xsj} s_{jt} = \sqrt{a_0^2 + b_0^2} (\alpha^{max} \omega_t + \beta^{max} r_t) = a_0 \omega_t + b_0 r_t$.

Case 2: Multi-homing agent observing at least two distinct frontier sources. To study the multi-homing cases, we first prove the following Lemma.

Lemma A1. *In the multi-homing case, $\rho'_{xs} \Sigma^{-1} s_t = y' Z (Z' Z + K)^{-1} (Z' \varphi_t + \varepsilon_t)$ and $\rho'_{xs} \Sigma^{-1} \rho_{xs} = y' Z (Z' Z + K)^{-1} Z' y$, where $y = \begin{bmatrix} a_0 & b_0 \end{bmatrix}'$, Z is the $2 \times J$ matrix where the j th column is $\begin{bmatrix} \alpha_{0j} & \beta_{0j} \end{bmatrix}'$, K is a diagonal matrix such that the j th diagonal is $\kappa_{0j}^2 = 1 - \alpha_{0j}^2 - \beta_{0j}^2$, $\varphi_t = \begin{bmatrix} \omega_t & r_t \end{bmatrix}'$, and ε_t is the J -vector of ε_{jt} .*

Proof. The lemma follows from noting that $\rho_{xs} = Z' y$, $\Sigma = Z' Z + K$, and $s_t = Z' \varphi_t + \varepsilon_t$. \square

Now consider the case where there are exactly two frontier sources in the market. Since both are on the frontier and have distinct biases, $K = 0$ and Z spans \mathbb{R}^2 , so $Z (Z' Z + K)^{-1} Z' = I$. Lemma A1 then implies that $w_t = \rho'_{xs} \Sigma^{-1} s_t = a_0 \omega_t + b_0 r_t$. Note that the R^2 of the population regression of x_t on s_t is $\rho'_{xs} \Sigma^{-1} \rho_{xs}$, which is equal to $a_0^2 + b_0^2$ by Lemma A1. Therefore, the R^2 of the population regression of x_t on s_t and r_t must be weakly greater than $a_0^2 + b_0^2$. However, the R^2 from a regression of x_t on s_t and r_t cannot exceed the R^2 from a regression of x_t on r_t and ω_t , which is $a_0^2 + b_0^2$. Therefore, $\underline{a}^R = \sqrt{\zeta} = \sqrt{a_0^2 + b_0^2}$.

Next 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 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 $w_t = a_0\omega_t + b_0r_t$. Furthermore, we can conclude that $\underline{a}^R = \sqrt{a_0^2 + b_0^2}$ using the same argument as in the previous paragraph.

Case 3: Multi-homing agent observing large random market. 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 $Z(Z'Z + K)^{-1}Z' = (I - R)Q = R$. By Definition 2, $E[\alpha_j^2/\kappa_j^2]$, $E[\beta_j^2/\kappa_j^2]$, and $E[\alpha_j\beta_j/\kappa_j^2]$ exist and are finite. Therefore, $d_{\alpha\alpha}$, $d_{\beta\beta}$, and $d_{\alpha\beta}$ converge in probability in the limit as $J \rightarrow \infty$ by the weak law of large numbers. Furthermore, Definition 2 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 $E[\alpha_j\beta_j/\kappa_j^2]^2 < E[\alpha_j^2/\kappa_j^2]E[\beta_j^2/\kappa_j^2]$. This implies that $\text{plim}(d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2) > 0$. Therefore, $Z(Z'Z + K)^{-1}Z' = R \rightarrow_p I$.

Further algebraic manipulation shows that

$$y'Z(Z'Z + K)^{-1}\varepsilon_t = \frac{(a_0d_{\beta\beta} + b_0d_{\alpha\beta})\left(\frac{1}{J}\sum_{j=1}^J \frac{\alpha_j}{\sqrt{\kappa_j^2}}\tilde{\varepsilon}_{jt}\right) + (a_0d_{\alpha\beta} + b_0d_{\alpha\alpha})\left(\frac{1}{J}\sum_{j=1}^J \frac{\beta_j}{\sqrt{\kappa_j^2}}\tilde{\varepsilon}_{jt}\right)}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d_{\alpha\beta}^2} + o_p(1)$$

where $\tilde{\varepsilon}_{jt} = \varepsilon_{jt}/\sqrt{\kappa_j^2} \sim N(0, 1)$ are mutually independent across j as well as independent of α_j

and β_j . Therefore, by the weak law of large numbers, $y'Z(Z'Z + K)^{-1} \varepsilon_t \rightarrow_p 0$. It immediately follows that $w_t = y'Z(Z'Z + K)^{-1} (Z' \varphi_t + \varepsilon_t) \rightarrow_p y' \varphi_t = a_0 \omega_t + b_0 r_t$.

Next note that the R^2 of the population regression of x_t on s_t is $\rho'_{xs} \Sigma^{-1} \rho_{xs}$. Furthermore, $\rho'_{xs} \Sigma^{-1} \rho_{xs} = y'Z(Z'Z + K)^{-1} Z'y \rightarrow_p a_0^2 + b_0^2$. 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 $\underline{a}^R \rightarrow_p \sqrt{a_0^2 + b_0^2}$.

A.8 Proof of Proposition 4

Under Condition 1 and Condition 2, holding \underline{a}^R fixed is equivalent to holding A fixed, and is also equivalent to holding $a_0^2 + b_0^2$ fixed. By Lemma 4, $\bar{\omega}_t = A(a_0 \omega_t + b_0 r_t)$. Thus, $\phi = \mathbb{E} \left[\frac{1}{4} (\bar{\omega}_t - \omega_t)^2 \right] = \frac{1}{4} (1 - Aa_0)^2 + \frac{1}{4} (Ab_0)^2$. We have that $Aa_0 \leq 1$, because $\underline{a}^R = \sqrt{a_0^2 + b_0^2} \geq a_0$, so the first term decreases in a_0 . Increasing a_0 while holding $a_0^2 + b_0^2$ fixed decreases b_0 , so both terms of ϕ decrease in a_0 . By similar logic, ϕ increases in $|b_0|$.

By equation (3), $\bar{\omega}_t = \alpha^{max} \omega_t + \beta^{max} r_t$ if the agent's prior is tight. Therefore, $\phi = \frac{1}{4} (1 - \alpha^{max})^2 + \frac{1}{4} (\beta^{max})^2 = \frac{1}{4} \left(1 - \frac{a_0}{\sqrt{a_0^2 + b_0^2}} \right)^2 + \frac{1}{4} \left(\frac{b_0}{\sqrt{a_0^2 + b_0^2}} \right)^2$.

A.9 Proof of Proposition 5

Holding \underline{a}^R fixed is equivalent to holding $a_0^2 + b^2$ fixed and to holding A fixed. By Lemma 4, $\bar{\omega}_t^R = A(a_0 \omega_t + b r_t)$ and $\bar{\omega}_t^L = A(a_0 \omega_t - b r_t)$. Thus, $\pi = \mathbb{E} \left[\frac{1}{4} (\bar{\omega}_t^R - \bar{\omega}_t^L)^2 \right] = (Ab)^2$, which increases in b when A is held fixed. Holding $a_0^2 + b^2$ fixed while increasing a_0 implies a decrease in b , hence ϕ decreases in a_0 .

Further suppose the agents' priors are tight. By equation (3), $\bar{\omega}_t^R = \alpha^{max} \omega_t + \beta^{max} r_t$. Similarly, $\bar{\omega}_t^L = \alpha^{max} \omega_t - \beta^{max} r_t$. Thus, $\pi = \mathbb{E} \left[\frac{1}{4} (\bar{\omega}_t^R - \bar{\omega}_t^L)^2 \right] = \mathbb{E} \left[\frac{1}{4} (2\beta^{max} r_t)^2 \right] = \beta^{max} = b^2 / (a_0^2 + b^2)$. The final sentence follows from taking the limit as a_0 goes to 0.

A.10 Proof of Proposition 6

The first sentence of the Proposition is immediate from Lemma 4. We prove the remainder by noting that as $|\alpha_j| \rightarrow 1$ and as $|\beta_j| \rightarrow 0$ for all j , the single-homing agent has $\underline{a}_R \rightarrow \sqrt{a_0^2 + b_0^2}$ and posterior mean $\bar{\omega}_t = A(a_0 \alpha_{0j} + b_0 \beta_{0j}) s_{jt} \rightarrow Aa_0 \omega_t$ by Example 1 and Lemma 3. The multi-

homing agent has $\underline{a}_R = \sqrt{a_0^2 + b_0^2}$ and posterior mean $\bar{\omega}_t = A(a_0\omega_t + b_0r_t)$ by Lemma 4. Note that A is the same for the two agents provided Condition 2 holds. Single-homing agents then have expected disagreement 0, while multi-homing agents have positive expected disagreement when their biases b_0 are distinct.

A.11 Proof of Proposition 7

Consider the multi-homing case first. Take any $z \in [-1, 1]$ and $j \in \{1, \dots, J\}$. Define θ as follows. Set $a = 1, b = 0, \alpha_j = z, \alpha_k = z\Sigma_{jk}$ for all $k \neq j, \gamma = z\rho_{rsj}$ and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$. It is immediate that $\rho_{rs} = \alpha\gamma + \beta\sqrt{1-\gamma^2}$ (as required by Remark 1). Furthermore, note that θ corresponds to a well-defined covariance matrix for the unit-normal joint distribution of (ω_t, r_t, s_t) .²⁰ Therefore, we have that $\theta \in I(R^\theta)$. The same θ works in the single-homing case, which only requires a well-defined covariance matrix for the unit-normal joint distribution of (ω_t, r_t, s_{kt}) for each k .

The mean of the multi-homing agent's posterior on ω_t given s_t under μ_∞ is

$$\frac{\int_{I(R^\theta)} \alpha' \Sigma^{-1} s_t d\mu_0(\theta)}{\mu_0(I(R^\theta))},$$

where $I(R^\theta) = \{\theta \in \Theta : \tilde{\alpha}' \tilde{\Sigma}^{-1} \tilde{\alpha} \leq 1; b = 0\}$. Under single-homing, the analogous mean is

$$\frac{\int_{I(R^\theta)} \alpha_j s_{jt} d\mu_0(\theta)}{\mu_0(I(R^\theta))},$$

where $I(R^\theta) = \{\theta \in \Theta : \tilde{\alpha}'_j \tilde{\Sigma}_j^{-1} \tilde{\alpha}_j \leq 1 \forall j; b = 0\}$.²¹ Since μ_0 is (α, γ) -symmetric, both integrals above are zero.

²⁰This covariance matrix is given by

$$X = \begin{bmatrix} 1 & \gamma & \alpha' \\ \gamma & 1 & \rho'_{rs} \\ \alpha & \rho_{rs} & \Sigma \end{bmatrix}.$$

Setting $\gamma = z\rho_{rsj}, \alpha_j = z$ and $\alpha_k = z\Sigma_{jk}$ for all $k \neq j$, where $z \in [-1, 1]$, corresponds to supposing that $\omega_t = zs_{jt} + (1-z)e_t$, where $e_t \sim N(0, 1)$ and is independent of (r_t, s_t) . Since it follows that $\omega_t \sim N(0, 1)$, X is well-defined.

²¹See the proof of Lemma 1 for the definitions of $\tilde{\alpha}_j, \tilde{\Sigma}_j, \tilde{\alpha}$ and $\tilde{\Sigma}$.

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_{0j} and bias β_{0j} to maximize $\Pi = \sum_{t \in \{L, U, R\}} \mu_t \xi(\bar{\alpha}^t) - \lambda$, where $\bar{\alpha}^t$ is type- t consumers' trust in the monopolist.

The derivative of trust $\bar{\alpha}^t$ with respect to β_{0j} along the frontier is:

$$\delta^t(\beta_0) = \left. \frac{\partial \bar{\alpha}^t}{\partial \beta_{0j}} \right|_{\alpha_{0j}^2 + \beta_{0j}^2 = 1} = \frac{1}{\sqrt{a_{0t}^2 + b_{0t}^2}} \left(b_{0t} - a_{0t} \frac{\beta_{0j}}{\sqrt{1 - \beta_{0j}^2}} \right).$$

The optimal frontier location must satisfy the first order condition that

$$\left. \frac{\partial \Pi}{\partial \beta_{0j}} \right|_{\alpha_{0j}^2 + \beta_{0j}^2 = 1} = (1 - 2\mu) \xi'(\bar{\alpha}^U) \delta^U(\beta_{0j}) + \mu [\xi'(\bar{\alpha}^R) \delta^R(\beta_{0j}) + \xi'(\bar{\alpha}^L) \delta^L(\beta_{0j})] = 0. \quad (5)$$

This condition is satisfied at $\beta_{0j} = 0$ because $\delta^U(0) = 0$, $\bar{\alpha}^R = \bar{\alpha}^L$, and $\delta^R(\beta_{0j}) = -\delta^L(\beta_{0j})$. When $\beta_{0j} > 0$, $\xi'(\bar{\alpha}^R) \leq \xi'(\bar{\alpha}^L)$ because $\xi(\cdot)$ is assumed to be concave and $\bar{\alpha}^R \geq \bar{\alpha}^L$. Moreover it is straightforward to show that $\delta^R(\beta_{0j}) + \delta^L(\beta_{0j}) < 0$ and $\delta^L(\beta_{0j}), \delta^U(\beta_{0j}) < 0$. Thus the derivative in equation (5) is strictly negative. Symmetric reasoning shows that this derivative is strictly positive when $\beta_{0j} < 0$. Thus, the unique solution is for the monopolist to choose $\beta_{0j} = 0$ and $\alpha_{0j} = 1$.

Since profits at this position are strictly positive, the monopolist will enter when λ is sufficiently low. Overconfidence for biased agents follows from noting that $\underline{a}^R = \sqrt{a_{0t}^2 + b_{0t}^2}$, and the expected disagreement result follows from noting that $\bar{\alpha}^R = \bar{\alpha}^L$.

A.13 Proof of Proposition 9

First, suppose one of the positions $\{\beta^L, 0, \beta^R\}$ 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.

The remaining results follow from noting that $\underline{a}^R = \sqrt{a_{0i}^2 + b_{0i}^2}$ and applying Proposition 5.

A.14 Proof of Proposition 10

By equation (3), R has bias $b > 0$ and holds a posterior belief about ω_t in each period degenerate at the value $\bar{\omega}_t^R = \alpha^{max} \omega_t + \beta^{max} r_t$, while L has the opposite bias and a posterior belief degenerate at $\bar{\omega}_t^L = \alpha^{max} \omega_t - \beta^{max} r_t$. L observes R 's policy choice d_t in period t , which is equal to $\bar{\omega}_t^R$, while believing that the true state to be $\bar{\omega}_t^L$. 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 $B_R = 2\beta^{max}$.