Ideological Bias and Trust in Information Sources[†]

By Matthew Gentzkow, Michael B. Wong, and Allen T. Zhang*

We study the role of endogenous trust in amplifying ideological bias. Agents in our model learn a sequence of states from sources whose accuracy is ex ante uncertain. Agents learn these accuracies by comparing their own reasoning about the states based on introspection or direct experience to the sources' reports. Small biases in this reasoning 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 reasoning. Disagreements can be similar in magnitude whether agents see only ideologically aligned sources or diverse sources. (*JEL* D42, D72, D83, L25, L82, Z13)

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 (McIntyre 2018; Jerit and Zhao 2020). 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 Research Center 2020). Such divisions have deepened even as new media technologies have made information more widely and cheaply available than ever before. Paradoxically, the information age has 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 toward 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 even small cognitive biases when agents are uncertain about 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 Rao 2016; Barberá 2020), 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.

Agents in our model wish 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, each agent i 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. This is a separate state variable that captures the way that realizations of that issue map to the political arena. In particular, r_t is the value of ω_t that would be most consistent with conservative ideology, and $-r_t$ is the value of ω_t that would be most consistent with liberal ideology. If ω_t is the effectiveness of masks, r_t would be negative (since the conservative position tends to be that masks are ineffective). If ω_t is the extent of fraud in the 2020 election, r_t would be positive (since the conservative position is that fraud was extensive). We allow the information sources j to have ideological *biases* in the sense that they may be distorted in the direction of r_t—i.e., the errors in the sources' reports 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 each agent's own *reasoning* about the state based on introspection or direct experience. We model this as an observable random variable x_{it}. In the case of global warming, x_{it} might be a realization combining

information about local temperatures or weather events the agent has experienced directly and/or her understanding of the mechanisms of climate science. In the case of an election, x_{it} might be a realization combining information about the procedures in place at polling places in her neighborhood and her evaluation of the difficulty of committing fraud. In the case of masks, x_{it} might be a realization based on the agent's understanding of the way the disease spreads and the mechanisms by which masks could stop it.

Our key assumption is that agents believe their own reasoning to be unbiased, even though this may not be true. In other words, x_{it} may be systematically correlated with r_t conditional on ω_t, but each agent's prior assigns probability one to the case where this correlation is zero.[1] This prior belief is what distinguishes x_{it} from the other information sources s_{jt} in the model. It means that even if an agent knows that her own reasoning is noisy, and so has limited information about ω_t on its own, she is able to use it as a yardstick to determine the accuracy and biases of the s_{jt} and ultimately to learn which sources she can trust. If in fact her reasoning *is* biased, the result will be distorted learning about the accuracy of information sources and, consequently, the states ω_t. Our main results characterize the form that such distortions take and show that they can in some cases be large even when the magnitude of the agent's bias (i.e., the correlation between x_{it} and r_t conditional on ω_t) is small.

Both parts of our key assumption are strongly supported by evidence. Large literatures in psychology and behavioral economics establish mechanisms by which individuals' own reasoning about politically sensitive topics may be subject to ideological bias, including motivated reasoning, selective memory, and availability (Festinger 1957; Tversky and Kahneman 1973; Lord, Ross, and Lepper 1979; Kunda 1990; Eagly et al. 1999). The implication of such biases is that two agents who for exogenous reasons begin with different ideologies and who have access to the same direct experiences—say, observe the same observations of local weather events—might reach conclusions about ω_t that are systematically biased toward or away from r_t. Moreover, substantial evidence also suggests that individuals themselves are not aware of these biases or, at a minimum, significantly underestimate them (e.g., Pronin, Lin, and Ross 2002; Pronin 2007; Thaler 2024).

Our formal results characterize the limiting distribution of each agent's beliefs about accuracies and states as $t \to \infty$. We show that their 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 an agent has no ideological bias and there is no uncertainty about the accuracy of her reasoning x_{it}, her trust is a distribution degenerate at the true accuracies 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 an agent's reasoning changes the results of the benchmark case dramatically. An agent with a small conservative bias may come to trust

---

[1] In an extension, we show that our main results extend to the case where agents entertain the possibility that their reasoning is biased, provided the magnitude of bias in the support of their priors is sufficiently small.

right-leaning sources more than is warranted by their true accuracy and trust unbiased sources less than is warranted. She may become overconfident in the accuracy of her own reasoning (i.e., come to believe that x_t is more strongly correlated with ω_t 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 even if the bias is small, provided that the true accuracy of an agent's reasoning is sufficiently low.

To see the intuition for the way in which 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 reasoning x_{it}. This correlation will be zero when s_{jt} is perfectly uncorrelated with the state and positive when the source is perfectly accurate. But when the reasoning x_{it} is noisy, the magnitude of the correlation will be small even in the perfectly accurate case. Small differences in observed correlation—such as those that might be induced by a small ideological bias—thus imply large differences in perceived accuracy. An agent who knows that her own ability to discern the truth of global warming is limited will view even a weak correspondence between her own conclusions and the reports of an information source as highly diagnostic.

Distortions in the way that agents learn about the informational environment can translate into substantial 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 the disagreement between different agents depends on the accuracy and bias of their reasoning, as well as on those of the observed sources. When agents all observe a common unbiased source, the 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 disagreements.

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 bias by media outlets in a sequential positioning game. We find that media competition can lead to greater media bias as well as intensified disagreements among viewers. Second, we show that 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 work contributes to a small but growing set of models in which agents must simultaneously learn about the states of the world and the accuracy of sources that provide information about those states. Acemoglu, Chernozhukov, and Yildiz (2016) show that arbitrarily small differences in priors about the signal distribution as well as the state can generate large and persistent disagreements among fully Bayesian agents. Sethi and Yildiz (2016) study the trade-off between learning from well-informed sources, whose signals are precise, and well-understood

sources, whose perspectives are well known, and show that long-run behavior is history independent. Cheng and Hsiaw (2022a) study a model in which agents do not integrate information about states and credibility correctly and show that this can lead to persistent polarization. Cheng and Hsiaw (2022b) study a cheap talk game where receivers are uncertain of a state and the sender's type and show that persistent disagreement arises. Liang and Mu (2020) study social learning about information sources over time and focus on failures of learning rather than persistent polarization or disagreement.

Relative to existing models, we highlight and characterize a novel mechanism by which endogenous trust can amplify small behavioral biases. This is related to but distinct from the mechanism in Acemoglu, Chernozhukov, and Yildiz (2016), which focuses on differences in prior beliefs among fully Bayesian agents. It is also distinct from the mechanism in Cheng and Hsiaw (2022a), which is based on a different behavioral bias and does not produce amplification. Our findings—that (i) agents end up trusting unreliable but ideologically aligned sources more than accurate neutral sources, (ii) agents become overconfident in their own reasoning, (iii) agents believe that either the conservative or the liberal point of view is closer to the truth on average, (iv) small biases in reasoning are amplified into persistent disagreements about facts, (v) 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, and (vi) competition among information sources can deepen rather than mitigate ideological disagreement—are all novel relative to the existing literature on simultaneous learning about states and signals.

Other related models attribute belief polarization to behavioral biases but do not consider how endogenous learning about the accuracy of information sources may amplify and alter these biases.[2] Examples include models of confirmation bias (Rabin and Schrag 1999), ambiguity aversion (Baliga, Hanany, and Klibanoff 2013), limited memory (Fryer, Harms, and Jackson 2019), and limited attention (Che and Mierendorff 2019). Bowen, Dmitriev, and Galperti (2023) show that belief polarization can arise in a social network where agents have small misperceptions about the sharing behavior of their neighbors.[3] In contrast to these theories, we build a model in which endogenous trust is the central mechanism, and we assume that agents only have small biases in their reasoning but otherwise have Bayesian learning rules and can process information from an arbitrarily large set of high-quality sources.

Our model contributes to the literature on media bias and competition (Mullainathan and Shleifer 2005; Gentzkow, Shapiro, and Stone 2015). 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

---

[2] There are many models wherein initial differences in the interpretation of signals can generate belief polarization (e.g., Dixit and Weibull 2007; Andreoni and Mylovanov 2012; Kondor 2012; Glaeser and Sunstein 2014; Benoît and Dubra 2019). These papers are motivated by a large number of experimental studies (e.g., Lord, Ross, and Lepper 1979) showing that the beliefs of subjects polarized after the presentation of new evidence.

[3] Relatedly, studies have shown that disagreements can persist in models of opinion dynamics in social networks with non-Bayesian learning rules (e.g., DeGroot 1974; DeMarzo, Vayanos, and Zwiebel 2003; Golub and Jackson 2010, 2012). Bayesian individuals that only observe posterior beliefs or actions of other individuals may also fail to learn underlying states because they are able only to recall or communicate coarse information (e.g., Banerjee 1992; Bikhchandani, Hirshleifer, and Welch 1992).

static, however, and does not provide a mechanism by which diverging beliefs or trust can persist over time.[4] Our model also relates to a broader literature on misspecified learning that has recently blossomed, but these models typically do not feature agents who learn underlying states and signal distributions at the same time.[5]

The paper proceeds as follows. Section I describes the model. Section II presents results on trust and polarization. Section III analyzes ideology and perceived bias. Section IV allows for overconfidence. Section V considers the multi-homing case. Section VI presents extensions. Section VII concludes.

I. Model

A. *Setup*

Each agent i learns about a sequence of unobservable states $\omega_t \sim N(0, 1)$ over time periods $t = 1, 2, \ldots, T$. Agents may observe signals (i.e., random variables) s_{jt} from information sources $j = 1, \ldots, J$, such as media outlets or talkative neighbors. There is also a random variable x_{it}, which represents information arising from agent i's independent *reasoning* about ω_t based on logic, introspection, direct experience, or other factors unrelated to s_{jt}.

In each period, each agent observes her own reasoning x_{it} followed by at least one signal s_{jt} and possibly also the ideological valence r_t. Our baseline analysis focuses on the *single-homing* case, where agents choose one source to observe in each period to minimize their cumulative expected loss over all periods. We later study the *multi-homing* case, where all agents observe all available sources in every period. After observing the signal(s) and her own reasoning, an agent chooses an action $d_{it} \in \mathbb{R}$. At the end of the game, the agent receives a loss equal to $\sum_{t=1}^{T} (d_{it} - \omega_t)^2$. The true value of ω_t is never observed.

The ideological valence r_t is the conservative position on ω_t. We interpret this as the value of ω_t most consistent with conservative ideology—e.g., the level of mask effectiveness, election fraud, or global warming severity that a representative conservative ideologue would argue in favor of. In some cases, one could also think of it as the action d_{it} that a conservative ideologue would want to persuade an agent to take. The liberal position on ω_t is $-r_t$. In many cases, we expect agents to be aware of how issues like masks and election fraud map to ideology, so we assume that r_t is observed (formally, at the same time as signals s_{jt}). We will also consider the case where r_t is unobserved. This simplifies the analysis in some ways and may be

---

[4] A large empirical literature studies 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 experimental literature studies the link between trust in information sources and political beliefs (Levendusky 2013; Nisbet, Cooper, and Garrett 2015; de Benedictis-Kessner et al. 2019; Thaler 2024; Jo 2020).

[5] Berk (1966) provides a general statement that beliefs need not converge in the long run under misspecified learning. Heidhues, Kőszegi, and Strack (2018) study an overconfident agent who becomes misdirected away from the optimal action as she learns about a fundamental. Heidhues, Kőszegi, and Strack (2019) study a model in which agents are overconfident in their own abilities and form disfavorable beliefs about the abilities of the members of out-groups from observing signals that are influenced by both individual ability and group-level discrimination. Frick, Iijima, and Ishii (2020) study social learning with misspecification and show that small misspecification can lead to extreme failures of learning. Bohren and Hauser (2021) provide a general framework for the analysis of learning in misspecified models.

appropriate for issues like interest rate policy, General Agreement on Tariffs and Trade negotiations, or foreign policy, where less politically engaged agents may not know political parties' current stances.

Together, ω_t, r_t, x_{it}, and the J-vector s_t of s_{jt} are jointly normal and are drawn independently over time:

$$\begin{bmatrix} \omega_t \\ r_t \\ x_{it} \\ s_t \end{bmatrix} \sim N(0, \Omega_i),$$

where $\Omega_{i11} = 1$ and the other elements of Ω_i are free parameters.[6]

We define the *accuracy* α_j of signal j to be the correlation of s_{jt} with ω_t. Note that this correlation is a sufficient statistic for the value of observing s_{jt} to an agent who knows Ω_i and seeks to learn about ω_t because we can rescale all of the variables to have variance one without changing the agent's posterior beliefs. Accuracy will be high when the covariance of s_{jt} and ω_t is large and/or when the variance of the noise in s_{jt} orthogonal to ω_t is low. A signal with higher accuracy $\alpha_j \geq 0$ has higher precision conditional on ω_t in the usual statistical sense, and is more informative according to the Blackwell (1953) ordering.[7]

We define the accuracy a_i of each agent's reasoning analogously to be the correlation of x_{it} with ω_t. We are primarily interested in the case where the value of a_i as small, so that the agents' ability to learn ω_t through independent reasoning alone is limited and is substantially less than what she could potentially learn through media or other information sources s_{jt}.

To define bias, let $\tilde{r}_t$ denote the standardized residual from a regression of r_t on ω_t. We will focus throughout on the case where, in fact, r_t is independent of ω_t, so $\tilde{r}_t = r_t / \sqrt{\operatorname{var}[r_t]}$. However, for reasons that will become clear below, we will want to allow agents to entertain the possibility that ω_t and r_t are correlated. We let γ denote this correlation. In the case where they are correlated, we would not want to say that a signal is biased if it is correlated with r_t solely through its correlation with ω_t. Thus, we define the *bias* β_j of signal j to be the correlation of s_{jt} with $\tilde{r}_t$. We define the bias b_i of x_{it} to be its correlation with $\tilde{r}_t$. For ease of exposition, we say that a source j is *perfectly right-biased* if $\beta_j = 1$, *perfectly left-biased* if $\beta_j = -1$, and *perfectly accurate* if $\alpha_j = 1$. When b_i and β_j have the same sign, we say that agent i and source j are *like-minded*. When they have the opposite sign, we say agent i and source j are *opposite-minded*.

---

[6] The assumption that r_t, x_t, and s_t have zero mean is purely to simplify exposition. If the means of these variables were free parameters and agents' priors on these parameters had full support, agents would learn the true value of these parameters in the limit as $t \to \infty$. Agents' beliefs about other parameters—and, thus, the inferences that are relevant for our results—would be unchanged. See Appendix B for details.

[7] To see the first point, note that, conditional on ω_t, the random variable $s_{jt} / \left[\alpha_j \sqrt{\operatorname{var}[s_{jt}]} \right]$ is distributed normally with mean ω_t and variance $1/\alpha_j^2 - 1$. Thus, for a decision-maker who knows the values of α_j and $\operatorname{var}[s_{jt}]$, observing s_{jt} is equivalent to observing a normal variable with mean ω_t and precision (defined in the usual statistical sense of the inverse variance) increasing in α_j. To see the second point, note that $\alpha_j > \alpha_k \geq 0$ implies that s_{kt} is a garbling of s_{jt}, since we can produce a variable s_{kt}^* with the same conditional distribution as s_{kt} by rescaling s_{jt} and adding random noise orthogonal to ω_t.

We further assume that $a_i > 0$, so that x_{it} is always positively correlated with ω_t, that x_{it} is independent of s_t conditional on ω_t and r_t, and that both of these restrictions are known by the agents. As shown in Remark 1 below, the conditional independence of x_t and s_t implies that the correlations between the observable variables—namely r_t, x_{it}, and s_t—is fully determined by $(a_i, b_i, \alpha, \beta, \gamma, \Sigma)$, where γ is the correlation of r_t with ω_t, $\Sigma = \mathrm{corr}(s_t)$, and α and β denote the J-vectors of α_j and β_j, respectively. The independence of ω_t and $\tilde{r}_t$ implies that $a_i^2 + b_i^2 \le 1$ and $\alpha_j^2 + \beta_j^2 \le 1$.

Note that in the simple case where the variances of x_{it} and the elements of s_t are all one, our setup implies that we can write

$$(1) \qquad\qquad s_{jt} = \alpha_j \omega_t + \beta_j \tilde{r}_t + \varepsilon_{jt},$$

$$(2) \qquad\qquad x_{it} = a_i \omega_t + b_i \tilde{r}_t + \eta_{it}$$

for $\varepsilon_{jt} \sim N(0, 1 - \alpha_j^2 - \beta_j^2)$ and $\eta_{it} \sim N(0, 1 - a_i^2 - b_i^2)$ residuals orthogonal to both ω_t and $\tilde{r}_t$.[8]

In the single-homing case, each agent chooses a source j to observe in each period to minimize the expected loss. This choice is akin to a multi-arm bandit problem. However, unlike the standard bandit problem, the payoffs from each period are not observed immediately, and the agent observes auxiliary information in the form of x_{it}. This renders standard solutions to the bandit problem inapplicable. To tractably capture the trade-off between exploration and exploitation, we assume that all agents follow an ε-first (sometimes called an "explore-first") strategy (Slivkins 2019).[9] That is, the agents each observe a random source with uniform probability during the first εT "exploration" periods, where $\varepsilon \in (0, 1)$. In the remaining periods, the agent chooses the source that minimizes expected loss, subject to the restriction that she chooses a source for which the posterior mean on α_j is weakly positive, provided such a source is available.[10]

We are interested in agent i's beliefs about $\theta_i = (a_i, b_i, \alpha, \beta, \gamma, \Sigma)$ and ω_t during the exploitation phase in the limit where the number of periods is large, so $T \to \infty$. The true value of θ_i is denoted θ_{0i}, and where it adds clarity, we use a_{0i}, b_{0i}, α_{0j}, β_{0j}, γ_0, and Σ_0 to refer to the true values of the individual components. The set of all possible θ_i is denoted as $\Theta_i = \Theta$. In our main case of interest, the value of a_{0i} is small and $\gamma_0 = 0$.

---

[8] When the variances are not one, these same expressions hold with s_{jt} replaced by $s_{jt}/\sqrt{\mathrm{var}[s_{jt}]}$ and x_{it} replaced by $x_{it}/\sqrt{\mathrm{var}[x_{it}]}$.

[9] This strategy is not optimal but it dramatically increases the tractability of our model by separating the characterization of long-run beliefs from the endogenous choice of which source to observe. A more optimal strategy would incorporate information from sequentially observed signals into the dynamic choices of what sources to observe by trading off the gains from exploration and exploitation. However, in the limit as $T \to \infty$, it will be in the agent's interest to learn the accuracies as precisely as possible, so we expect long-run beliefs under the epsilon-first strategy to be the same as under a more optimal strategy.

[10] The restriction rules out an agent choosing to observe sources she believes are negatively correlated with the truth over one she believes is positively correlated with the truth; such a strategy could improve the accuracy of beliefs in principle, but we see it as unrealistic from a behavioral point of view. We relax this assumption in Appendix E and show that most of our results are robust to dropping this restriction.

TABLE 1—VARIABLES AND PARAMETERS

ω_t	True state
r_t	Ideological valence
x_{it}	Agent i's reasoning
s_t	Vector of signals s_{jt} from information sources
a_i	Correlation of x_{it} with ω_t
b_i	Correlation of x_{it} with the residual of a regression of r_t on ω_t
α	Correlation of s_t with ω_t
β	Correlation of s_t with the residual of a regression of r_t on ω_t
γ	Correlation of r_t with ω_t
Σ	Correlation matrix of s_t

The correlation matrix for $(\omega_t, r_t, x_{it}, s_t)$, which we denote as $\tilde{\Omega}_i$, is fully parametrized by θ_i.[11] The covariance matrix is thus $\Omega_i = V_i^{1/2} \tilde{\Omega}_i V_i^{1/2}$, where V_i is a diagonal matrix containing the variances of $(\omega_t, r_t, x_{it}, s_t)$. We assume that the set of all possible variances V_i is a compact set $\mathcal{V}$. The Lebesgue space on $(\Theta, \mathcal{V})$ is denoted $((\Theta, \mathcal{V}), \mathcal{L}, \nu)$.

Example (*Masks*).—In period t, the state ω_t indexes the effectiveness of masks at preventing the transmission of disease during a pandemic.[12] The ideological valence r_t is negative: conservatives believe that masks are relatively ineffective, while liberals believe that 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 true accuracy of these sources is captured by α_j. Right-biased sources (with $\beta_j > 0$) might distort information in a negative direction by downplaying evidence that masks are effective (causing the value of s_{jt} to be pulled closer to $r_t < 0$; left-biased news reports might do the opposite (causing the value of s_{jt} to be pulled closer to $-r_t > 0$). Reasoning x_{it} could be information from logical inference (e.g., about the likelihood that masks fully block disease-carrying droplets) and direct observation (e.g., of people in the agents' own social network who either did or did not wear masks and

---

[11] Specifically, given any θ_i, the correlation matrix for $(\omega_t, r_t, x_{it}, s_t)$ is given by

$$\tilde{\Omega}_i(\theta_i) = \begin{bmatrix} 1 & \gamma & a_i & \alpha' \\ \gamma & 1 & a_i\gamma + b_i\sqrt{1-\gamma^2} & \alpha'\gamma + \beta'\sqrt{1-\gamma^2} \\ a_i & a_i\gamma + b_i\sqrt{1-\gamma^2} & 1 & a_i\alpha' + b_i\beta' \\ \alpha & \alpha\gamma + \beta\sqrt{1-\gamma^2} & a_i\alpha + b_i\beta & \Sigma \end{bmatrix}.$$

Note then that Θ is the set of all $\theta \in (0,1] \times [-1,1]^{2J+2} \times [-1,1]^{J^2}$ such that $\tilde{\Omega}(\theta)$ is positive semidefinite and the diagonal entries of Σ are equal to one.

[12] It might be more intuitive to think of the effectiveness of masks as a bounded random variable. Note that we can accommodate this by interpreting effectiveness as $\exp(c_0 + c_1\omega_t)/[1 + \exp(c_0 + c_1\omega_t)]$ so that it has support on $[0,1]$ and the arbitrary constants c_0 and c_1 determine the location and scale of the distribution of effectiveness within the $[0,1]$ interval.

may or may not have caught the disease). The parameter a_i governs the accuracy of this reasoning, and it might depend on factors like the agent's knowledge of science. The parameter b_i governs the extent to which the agent's reasoning is systematically distorted toward either the liberal or the conservative position on mask effectiveness.

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 that fraud was high, while liberals believe that 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 true accuracy of these sources is captured by α_j. Right-biased sources (with $\beta_j > 0$) might distort information in a positive direction by emphasizing anecdotes suggesting fraud (causing the value of s_{jt} to be pulled closer to $r_t > 0$); left-biased news reports might do the opposite by emphasizing evidence that fraud is rare (causing the value of s_{jt} to be pulled closer to $-r_t < 0$). Reasoning x_{it} could be information from logical inference (e.g., about the likelihood that widespread election fraud would be detected) and direct observation (e.g., of the way that polling places, mail-in ballots, etc. were administered in the agents' own neighborhood). The parameter a_i governs the accuracy of this reasoning, and it might depend on factors like the extent to which the agent knows details of the vote-counting process. The parameter b_i governs the extent to which the agent's reasoning is systematically distorted toward either the liberal or the conservative position on election fraud.

B. *Learning*

At the beginning of the first period, agents have an absolutely continuous prior belief μ_0^i over $(\Theta, \mathcal{V})$ with a continuous density with respect to ν. For tractability, we assume that θ_i is independent of V_i under each agent's prior, and we let $\mu_{0,\theta}^i$ and $\mu_{0,V}^i$ respectively denote agent i's prior distribution on θ_i and V_i. Our central assumption is that all θ_i in the support of $\mu_{0,\theta}^i$ have $b_i = 0$. In other words, all agents a priori believe that their own reasoning x_{it} is unbiased, even though their true bias may be nonzero. In Appendix C, we show that our main results extend to the case where agents believe their own bias to be $b_i = b^*$ for some small $b^* \neq 0$.

All θ_i in the support of $\mu_{0,\theta}^i$ also have $a_i \in \mathcal{A}_i$ for some set $\mathcal{A}_i$. In our baseline analysis, we assume that agents know the accuracy of their own reasoning, so $\mathcal{A}_i = \{a_0\}$. We then extend the model to allow for the possibility that agents could become overconfident. In this case, $\mathcal{A}_i = (0, a_i^{max}]$ for some $a_i^{max} \geq a_0$. We will assume that $\mu_{0,\theta}^i$ has full support on the subset Θ_i^{prior} of Θ consistent with the restrictions that $b_i = 0$ and $a_i \in \mathcal{A}_i$.

In our baseline case, agent i observes both x_{it} and s_{jt} for each j in many periods, and so, in the limit as $t \to \infty$, she learns the J-vector ρ_{is} of correlations between x_{it} and the elements of s_t. In the case where r_t is observed, agents also learn the J-vector ρ_{rs} of correlations between r_t and the elements of s_t and the correlation ρ_{ir}

of x_{it} and r_t. Finally, in the case of multi-homing, agents observe all sources in each period, so they additionally learn the correlation matrix Σ_0 of the elements of s_t. We abuse notation and let R_{0i} denote the true value of the vector of correlations that can be learned by agent i—(ρ_{is}), $(\rho_{is}, \rho_{rs}, \rho_{ir})$, or $(\rho_{is}, \rho_{rs}, \rho_{ir}, \Sigma_0)$, depending on the case considered—and we let $R_i(\theta_i)$ denote the observable subset of the correlations $\tilde{\Omega}_i$ that would occur under parameters θ_i.

Our first result defines agent i's limiting beliefs about θ_i. We define the *identified set* of parameters in the support of agent i's prior that is consistent with correlations R_{0i} as $I_i(R_{0i}) = \{\theta_i \in \Theta_i^{prior} : R_i(\theta_i) = R_{0i}\}$. The following proposition shows that agent i's beliefs converge asymptotically to this identified set. Because all θ_i in the identified set imply the same distribution of observed data, beliefs within the set remain proportional to the prior.

PROPOSITION 1: *Suppose that the true correlations of the observed data are R_{0i}. As $T \to \infty$, agent i's posterior distribution in any period $t > \varepsilon T$ converges to a limit $\mu_{\infty,\theta}^i$ such that for all measurable $\vartheta \subseteq \mathcal{L}_\Theta$,*

$$(3) \qquad \mu_{\infty,\theta}^i(\vartheta) = \frac{\mu_{0,\theta}^i(\vartheta \cap I_i(R_{0i}))}{\mu_{0,\theta}^i(I_i(R_{0i}))}.$$

PROOF:

All proofs are in Appendix A.

The second result derives the agents' posterior expectations of ω_t given their limiting beliefs about θ_i. For tractability, we analyze the belief that an agent would hold after observing information source(s), but not her reasoning x_{it} or the ideological valence r_t. This is similar to analyzing the case where the accuracy a_i of the agent's reasoning is close to zero and the agents are aware of this. We denote the mean of the agent's posterior at this ex interim stage by $\bar{\omega}_t^i$. We define agent i's expectation of α_j under $\mu_{\infty,\theta}^i$, which we denoted as $\bar{\alpha}_j^i$, as agent i's *trust* in source j. The following characterization follows from standard conjugate prior results for the normal distribution.

PROPOSITION 2: *Suppose that agent i single-homes and her posterior belief on θ_i is $\mu_{\infty,\theta}^i$. In any period t where she observes source j, the mean of her posterior on ω_t given s_{jt} (but not x_{it} or r_t) is*

$$(4) \qquad \bar{\omega}_t^i(s_{jt}) = \bar{\alpha}_j^i \tilde{s}_{jt},$$

where $\tilde{s}_{jt} = s_{jt}/\sqrt{\text{var}[s_{jt}]}$ is the standardized version of s_{jt}.

This proposition shows that the mean of agent i's posterior on ω_t will be proportional to the normalized value of the signal s_{jt} that she observes, with the constant of proportionality equal to her trust $\bar{\alpha}_j^i$ in source j. When the agent believes that source j is unreliable ($\bar{\alpha}_j^i \approx 0$), her beliefs about the true state do not vary much with the

content of source j, and they stay concentrated close to the prior value of zero. When the agent's trust in source j is high ($\bar{\alpha}_j^i \approx 1$), her beliefs track the content of the source closely, and their variance approaches 1.

Our third result shows that, given her limiting belief about θ_i, a single-homing agent observes a source that she believes to be most accurate. To be precise, let $\mathcal{J}_+^i$ denote the set of sources j such that $\bar{\alpha}_j^i \geq 0$, which we assume here to be nonempty. The agent chooses to observe the source from this set that minimizes her expected mean squared error after observing both the source and her own reasoning:[13]

$$(5) \qquad \min_{j \in \mathcal{J}_+^i} \mathbb{E}\left[\left(d_{ijt}^*(s_{jt}, x_{it}) - \omega_t\right)^2\right],$$

where the expectation is taken under the marginal distribution of $(s_{jt}, x_{it}, \omega_t)$ under μ_∞^i and $d_{ijt}^*(s_{jt}, x_{it})$ is the optimal decision after seeing (s_{jt}, x_{it}).[14]

PROPOSITION 3: *Suppose that agent i single-homes, her posterior belief on θ_i is $\mu_{\infty,\theta}^i$, and the expected accuracy a_i of her own reasoning under $\mu_{\infty,\theta}^i$ is less than one. Then she chooses to observe in each period a source j for whom her trust $\bar{\alpha}_j^i$ is maximal (provided some source has $\bar{\alpha}_j^i \geq 0$).*

Proposition 3 shows that, setting aside the degenerate case where an agent believes her own reasoning to be perfectly accurate, she converges to always getting information from the source she trusts most. This result is not entirely trivial. The agent's choice of j is equivalent to minimizing posterior variance integrating over both the uncertainty in $(s_{jt}, x_{it}, \omega_t)$ given Ω_i and the uncertainty in Ω_i. If we were only integrating over the former, this would be a standard case of choosing among signals with known distributions, and it would be optimal to choose the one with the highest α_j. Once there is uncertainty over Ω_i, it may not be obvious that the agent would choose the one with the highest expected α_j, since the whole distribution of the agent's belief about α_j (and other parameters) could in principle matter for the expected loss. Our proof exploits the fact that the optimal decision rule depends linearly on α_j.

C. *Benchmark Cases*

As shown above, information about source j's accuracy α_j allows the agents to learn about the underlying state ω_t from observing s_{jt}. As shown below in Section II, agents rely on the observed correlation of reasoning x_{it} and signal s_{jt} to learn α_j. In this section, we discuss benchmark cases where either no signal or no reasoning is available. In these cases, agents do not endogenously learn α_j, and small ideological biases in the agents' reasoning are not amplified into large disagreements. For simplicity, we

---

[13] Here it is unimportant whether the agent observes x_{it}. The same source j is chosen whether the agent minimizes her mean squared error after observing both x_{it} and s_{jt} or after observing s_{jt} but not x_{it}.

[14] If the restriction that $j \in \mathcal{J}_+^i$ were dropped, the agent instead chooses to observe a source j for whom $|\bar{\alpha}_j^i|$ is maximal. In Appendix E, we show that all of our results in subsequent sections are robust to dropping this restriction.

assume here that there is a single source ($J = 1$), agents know the accuracy of their own reasoning ($\mathcal{A}_i = \{a_0\}$), and r_t is never observed.

No Signals.—First, suppose that agent i observes only her reasoning x_{it} in every period, and never observes any source j or the ideological valence r_t. In the main case of interest where the accuracy of her reasoning a_{0i} is small, her posterior beliefs will always be noisy. To see this, assume for the moment that x_{it} is known to have variance equal to one, and that agent i knows the accuracy of their own reasoning, so $\mathcal{A}_i = \{a_{0i}\}$. Then her posterior mean on ω_t after observing x_{it} is $\bar{\omega}_t^i = a_{0i} x_{it} = a_{0i}(a_{0i}\omega_t + \eta_{it}) = a_{0i}^2 \omega_t + a_{0i}\eta_{it}$, where η_t is an error term that is normally distributed with mean zero and variance $1 - a_{0i}^2$. When a_{0i} is small, $\bar{\omega}_t^i$ is a very noisy estimate of ω_t.

No Reasoning; Exogenous Trust.—Next, suppose that agent i observes a single source j every period but never observes her reasoning x_{it} or the ideological valence r_t. In this case, her beliefs about the source's accuracy α_j never change. The mean of her posterior belief about ω_t after observing s_{jt} in any period t is $\bar{\omega}_t^i = \bar{\alpha}_j^0 s_{jt}$, where $\bar{\alpha}_j^0$ is the expectation of the source's accuracy under the agent's prior μ_0.

If $\bar{\alpha}_j^0$ is zero—roughly speaking, agents think that the source is as likely to be a perverse "false news" source as an accurate "true news" source—then agent i's posterior belief about ω_t would be zero in all periods, no matter what value of s_{jt} is realized. She would therefore learn nothing from observing s_{jt}, even if s_{jt} is in fact highly accurate.

If instead $\bar{\alpha}_j^0 = \alpha_{0j}$, so the agents exogenously know the true accuracy of signals, then their posterior belief about ω_t will be the same as that of a Bayesian. Furthermore, if source j is highly accurate, $\bar{\omega}_t^i$ will be close to the true value ω_t. This will be true regardless of the agent's bias b_{0i} or the source's bias β_{0j}.

A closely related benchmark is one where agent i does observe x_{it} but trust is exogenous in the sense that the agent's prior on α_j is degenerate at $\bar{\alpha}_j^0$. Updating from s_{jt} in this case will be identical to that when x_{it} is not observed. The only difference is that agent i's beliefs about ω_t would also be affected by the realization of x_{it}. In the main case of interest where the accuracy a_{0i} of x_{it} is low and the agents know that it is low, the effect of x_{it} on beliefs will be small, and it will remain true that beliefs do not depend on bias. In this sense, endogenous trust is necessary in our model for bias to produce polarization.

Toward the Full Model.—Now consider an agent i who observes both x_{it} and s_{jt} in every period and whose prior mean is $\bar{\alpha}_j^0 = 0$. In this case, her reasoning x_{it} functions as a reference point that allows her to learn the accuracy of source j. To see this, note that the agent believes that (i) x_{it} is positively correlated with ω_t (since $a_i > 0$ with probability one under μ_0^i) and (ii) x_{it} is uncorrelated with $\tilde{r}_t$ (since $b_i = 0$ with probability one under μ_0^i). This means that she expects the correlation between x_{it} and s_{jt} to be increasing in the source's accuracy α_j. If she observed a more positive correlation between x_{it} and s_{jt} than she would expect under her prior, then she would endogenously revise her estimate of the source's accuracy upward. Having learned about source j's accuracy, the agent can then learn about ω_t from observing s_{jt}.

If agent i's reasoning is unbiased, her posterior belief on α_j converges to the true value α_{0j} as $t \to \infty$. If the source was highly accurate (i.e., α_{0j} was equal to one), then, in the limit as $t \to \infty$, her posterior mean $\bar{\omega}_t^i$ approaches the true value ω_t. In other words, having access to noisy but unbiased reasoning allows an agent to accurately learn which sources to trust by aggregating over many periods. Unbiased agents can thereby learn the underlying states of the world as if they exogenously knew the accuracy of the information sources—and, therefore, much more precisely than if they had access to only her reasoning. However, as we show in the following sections, endogenous trust can lead small biases in reasoning to become amplified into large disagreements.

D. *Discussion: Reasoning and Bias*

What does x_{it} represent in the real world? As a baseline case, we can think of x_{it} as information that results from the agent's reasoning and introspection about the likely value of ω_t.[15] In the examples above, this could involve thinking about the likelihood that masks fully block disease-carrying droplets or that widespread election fraud is detected. If ω_t relates to economic stimulus policy, agents might reason from first principles about how large the plausible costs and benefits could be. (They might even write down and solve a model!) In writing the result of this reasoning and introspection as a random variable, we think of engaging in reasoning and introspection as analogous to observing the results of a noisy experiment, just as a student trying to solve a math problem in their head will produce a result that is imperfectly correlated with the true value.

There are a large number of well-studied psychological phenomena that could provide a microfoundation for bias in agents' reasoning and introspection (i.e., $b_{0i} \neq 0$). Consider an agent who has grown up in a liberal family, benefitted from liberal policies, and taken actions (like voting) consistent with liberal ideology. Such an agent may engage in motivated reasoning (Kunda 1990), distorting her inferences to reach conclusions that support a liberal point of view. She may underweight arguments or evidence 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 (Eagly et al. 1999). 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 had 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). Note that these latter two mechanisms could produce bias even in the case where the agent does not know the value of r_t—e.g., she might be more likely to reach conclusions favoring teachers' unions even if she is unaware that this is a position typically supported by the Democratic Party.

---

[15] Note that in this interpretation we do not think of x_{it} as the agent's posterior belief about the state but rather as the information arising from reasoning and introspection that produces that posterior belief.

A large body of evidence suggests that individuals themselves are not aware of these kinds of biases or, at a minimum, significantly underestimate them (e.g., Pronin, Lin, and Ross 2002; Pronin 2007; Thaler 2024).

Another possibility is that x_{it} captures 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 agents' locality (when ω_t relates to global warming), the agents' own experiences with public schools (when ω_t relates to education policy), or the agent's personal economic situation (when ω_t relates to economic policy). Such information could be observed either before or after seeing s_{jt}. Bias in these observations could arise through many of the same mechanisms as bias in reasoning. For example, an agent from a liberal background might be more likely to remember unusually hot days or unusually severe storms that suggest that global warming is severe (Eagly et al. 1999). It could also be that certain kinds of direct observation tend to favor one side or the other inherently, and agents fail to adjust for this selection. For example, observing the effort, wages, and impacts of public school teachers may naturally provide more evidence favorable to teachers' unions, and observing polling places where voting fraud is unlikely to be visible might naturally provide more evidence favorable to the view that elections are conducted fairly.

A final possibility is that x_{it} is the signal of a particular information source that agents believe a priori to be unbiased. This might be what their mothers say, what the Bible says, what scientists say, or even the report of a particular news source that they begin with extraordinary faith in.

What is crucial for our model is that x_{it} satisfies two conditions. First, agents believe that it is free from ideological bias ($b = 0$) with probability one. Second, contrary to the agents' beliefs, x_{it} may in fact be subject to ideological bias ($b_{0i} \neq 0$).

The ultimate source of the agents' biases is left unmodeled in our framework—b_{0i} is an exogenous parameter that we think of as determined by experiences and motivations prior to $t = 0$. We show below that nonzero bias leads to distortion in agents' trust in sources and their beliefs about states of the world, and also shapes what we label "ideology"—agent i's belief $\bar{\gamma}^i$ about the correlation between the truth ω_t and the ideological valence r_t. This point might lead to some confusion, as we show how bias can shape ideology, but psychological phenomena like motivated reasoning also provide reasons why ideology would shape bias. A richer model that included this two-way feedback between ideology and bias might well lead to even more dramatic divergence of trust and beliefs. We isolate one direction of causality for analytic tractability, taking bias in reasoning to be a fixed characteristic that may depend on baseline ideology and other traits but does not change as ideology evolves.

It may seem strong to assume that the agent puts such dogmatic ex ante faith in her own reasoning or observation. However, if there is *no* source of truth in which she would place significant faith in this sense, the agent would never learn what sources she can trust. More precisely, she would never be able to reject that any particular source is either perfectly positively correlated, uncorrelated, or perfectly negatively correlated with the state. We demonstrate this result in Appendix F. We also show in Appendix C that our results are robust to allowing agents to entertain the possibility of *small* biases in their own reasoning.

II. Trust and Polarization

In this section, we assume that there are multiple available sources, so $J \geq 2$, but each agent single-homes and so observes only one source in each period. For simplicity, we focus on the case where the ideological valence r_t is never observed. This means that ideological valence is a latent source of bias in signals and reasoning but is not something that agents make inferences from directly. While in reality r_t is often observed, it is plausible that a politically disengaged person may not know the conservative or liberal positions on a wide range of issues. We also focus on the case where agents know the accuracy a_{0i} of their own reasoning, so that $\mathcal{A}_i = \{a_0\}$. We consider the case where r_t is observed in Section III, the case where agents are uncertain about the accuracy of their own reasoning in Section IV, and the case of multi-homing in Section V.

For ease of exposition, we assume that there are three agents $i \in \{U, R, L\}$. All agents have accuracy $a_{0i} = a_0$, where $a_0 > 0$. Agent U's reasoning has no bias, so $b_{0U} = 0$. Agent R's reasoning has positive bias, while agent L's reasoning has negative bias. We set $b_{0R} = b_0$ and $b_{0L} = -b_0$, where $b_0 > 0$.

ASSUMPTION 1: *The support of $\mu_{0,\theta}^i$ is the set $\Theta_i^{prior} \subset \Theta$ for which $b_i = 0$ and $a_i = a_{0i}$.*

For agent U, this is a simple case of Bayesian learning with a correctly specified model. For agents R and L, the true parameters of the model lie outside the support of their priors (since $b_{0i} \neq 0$). Our model is thus an example of Bayesian learning under misspecification (Lian 2009).

Because r_t is not observed, agent i learns only the correlation ρ_{is} between s_t and x_{it}. This restricts agent i's limiting beliefs, since the model requires that $\rho_{is} = a_i \alpha + b_i \beta$. The identified set is then $I_i(R_{0i}) = \{\theta_i \in \Theta_i^{prior} : a_i \alpha + b_i \beta = \rho_{is}\}$.

Misspecified learning can lead to instability or lack of convergence (Berk 1966), but we focus our baseline analysis on parameter values such that the data agents observed do not violate their model of the world—i.e., the identified set is nonempty. A sufficient condition in the single-homing case is that $b_0 \leq a_0$ and $|\alpha_j| + |\beta_j| \leq 1$ for all j. We relax this assumption when we extend the model to allow overconfidence in Section IV below.

ASSUMPTION 2: $b_0 \leq a_0$ *and* $|\alpha_j| + |\beta_j| \leq 1$ *for all j.*

We can then show that the identified set $I_i(R_{0i})$ is nonempty and contains a single value of α. It does not restrict the values of γ or β, because when r_t is not observed, the observed correlations do not contain information about these parameters.

PROPOSITION 4: *Suppose that r_t is never observed and agents single-home. Under Assumptions 1 and 2, agent i's identified set, $I_i(R_{0i})$, is nonempty and consists of all $\theta_i \in \Theta$ such that $a_i = a_{0i}$, $b_i = 0$, and $\alpha = \rho_{is}/a_{0i}$.*

Proposition 4 shows that when r_t is never observed, each agent's beliefs about each source's accuracy is concentrated on $\alpha_j = \rho_{ij}/a_{0i}$. The agent believes that source j is more accurate the more correlated that source j's reports are with i's reasoning. For a given level of correlation, the update is larger the less accurate i believes her own reasoning to be, since when the correlation of x_{it} with the true state is small, x_{it} and s_{jt} can only be substantially correlated if the correlation of s_{jt} with the true state is large. This result follows from the agent's assumption that her reasoning to be unbiased. Its implications are discussed further below.

Trust.—A large body of evidence shows divergence between the sources trusted by conservatives and the sources trusted by liberals (e.g., Pew Research Center 2014a, 2020), and many have pointed to this as a key factor undermining the media's role in democracy (Gallup and Knight Foundation 2018, 2020). Consistent with this, agents in our model may come to trust biased sources more than unbiased sources. Divergence in trust may be large even when biases are small, provided that the accuracy of the agents' own reasoning is low.

To see this, notice that under Assumptions 1 and 2, Propositions 1 and 4 imply that agent i's trust in information source j is

$$(6) \qquad \bar{\alpha}_j^i = \frac{\rho_{ij}}{a_{0i}} = \alpha_{0j} + \frac{b_{0i}\beta_{0j}}{a_{0i}}.$$

Therefore, when the accuracy a_{0i} of the agents' reasoning is low, small differences in biases b_{0i} and β_{0j} translate into large differences in trust. Agents come to trust source j more when their biases are more aligned (i.e., $b_{0i}\beta_{0j}$ is positive), and that divergent trust can be extreme even when b_{0i} is close to zero.

COROLLARY 1: *Suppose that r_t is never observed, agents single-home, and Assumptions 1 and 2 hold. Then agent R's (L's) trust in source j is increasing (decreasing) in the source's bias β_{0j} holding constant the source's accuracy α_{0j}. In the limit as $a_0 \downarrow b_0$, she will come to believe that a perfectly right-biased (left-biased) source is perfectly accurate and will trust it more than any unbiased source with $\alpha_{0j} < 1$.*

Polarization.—Substantial literatures document large and growing disagreement between Democrats and Republicans on both policy issues (Pew Research Center 2014b; Boxell, Gentzkow, and Shapiro 2017) and questions of fact (Marietta and Barker 2019). Consistent with these findings, we show that disagreement in our model can be large even when the underlying bias is small and accurate information is widely available.

We focus on the exploitation periods $t > \varepsilon T$ in the limit as $T \to \infty$, where each agent observes the source j for which trust is maximal by Proposition 3. We define agents' *expected disagreement* when they observe these sources in such period to be $\pi = \mathbb{E}\left[\frac{1}{4}\left(\bar{\omega}_t^R - \bar{\omega}_t^L\right)^2\right]$. Scaling by one-fourth here ensures that $\pi \in [0,1]$.

The expected disagreement of single-homing agents is sensitive to the set of available sources. If all sources are unbiased, expected disagreement will be zero

because all agents will come to trust the most accurate source, and the trust of R and L agents for this source will be identical. However, if highly biased sources are available, disagreement can grow large.

To see this, suppose that all sources have accuracy $\alpha_{0j} < 1$ and there is one perfectly right-biased source and one perfectly left-biased source. By Corollary 1, agent R's trust in the perfectly right-biased source and agent L's trust in the perfectly left-biased source are both maximal in the limit as $a_0 \to b_0$. Proposition 3 then implies that each agent observes their like-minded, perfectly biased source during the exploitation periods. Expected disagreement then reaches the maximum possible value $\pi = 1$.

COROLLARY 2: *Suppose that r_t is never observed, agents single-home, and Assumptions 1 and 2 hold. Further suppose that all sources have accuracy $\alpha_{0j} < 1$ and there is at least one perfectly right-biased source and at least one perfectly left-biased source. In the limit as $a_0 \downarrow b_0$, expected disagreement is one.*

III. Ideology and Perceived Bias

We now consider the case where the ideological valence r_t is observed in every period. Agents are able to learn the vector of correlations $(\rho_{is}, \rho_{ir}, \rho_{rs})$. In this case, the results from the previous section are broadly unchanged, but we can additionally derive each agent's beliefs about the correlation of the true states with the ideological valence and the bias of sources.

Remark 1: Given parameters $\theta_i = (a_i, b_i, \alpha, \beta, \gamma, \Sigma)$, the elements of $R(\theta_i)$ are given by

$$\rho_{is} = a_i \alpha + b_i \beta, \tag{7}$$

$$\rho_{ir} = a_i \gamma + b_i \sqrt{1 - \gamma^2}, \tag{8}$$

$$\rho_{rs} = \alpha \gamma + \beta \sqrt{1 - \gamma^2}. \tag{9}$$

For simplicity, we continue to focus on parameter values such that the data an agent observes do not violate her model of the world—i.e., the identified set is nonempty. In the current case, this will require $b_{i0}^2 / a_{i0}^2 \leq 1 - \alpha_{0j}^2 / (1 - \beta_{0j}^2)$ for all i and j. This assumption is more stringent than Assumption 2 because learning the values of ρ_{ir} and ρ_{rs} allows agents to reject their models of the world in more cases. We relax this assumption when we extend the model to allow overconfidence in Section IV below.

ASSUMPTION 2′: $b_0^2 / a_0^2 \leq 1 - \alpha_{0j}^2 / (1 - \beta_{0j}^2)$ *for all j.*

Under this assumption, the identified set $I_i(R_{0i})$ is nonempty.

PROPOSITION 5: *Suppose that r_t is observed and agents single-home. Under Assumptions 1 and 2′, agent i's identified set, $I_i(R_{0i})$, is nonempty and consists of all $\theta_i \in \Theta$ such that $a_i = a_{0i}$, $b_i = 0$, $\alpha = \rho_{is}/a_i$, $\gamma = \rho_{ir}/a_i$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$.*

Comparing Proposition 5 to Proposition 4 reveals three implications of allowing r_t to be observed: First, the agent's identified set is concentrated on the value $\gamma = \rho_{ir}/a_i$. Second, the agent's beliefs about a source's bias β_j is a known function of the observed correlations ρ_{ij}, ρ_{ir}, and ρ_{rj}. Third, the agent's identified set is still concentrated on the value $\alpha = \rho_{is}/a_i$ even when r_t is observed. The implications of these facts are explored below.

Ideology.—A striking feature of political polarization in the United States is the significant correlation between citizens' liberal-conservative ideologies and their views across a range of diverse issues (Gentzkow 2016). Consistent with this fact, agents in our model when r_t is observed may 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 agents' limiting beliefs about the correlation γ between ω_t and r_t.

We define an agent's *ideology* $\bar{\gamma}^i$ to be her limiting posterior mean on γ. We say that agent i's ideology is *right leaning* if $\bar{\gamma}^i > 0$. Under Assumptions 1 and 2′, and focusing on the case of interest where $\gamma_0 = 0$, Proposition 5 implies that agent i's ideology is

$$(10) \qquad\qquad \bar{\gamma}^i = \frac{b_{0i}}{a_0}.$$

Equation (10) implies that if agent R simply observes the ideological valence r_t, then her beliefs about ω_t become biased toward r_t, while agent L comes to believe the opposite. This is because agent R sees positive correlation between r_t and x_{it}, so she comes to believe that conservative views are closer to the truth on average than liberal views.

COROLLARY 3: *Suppose that r_t is observed, agents single-home, and Assumptions 1 and 2′ hold. Then agent R's (L's) ideology is increasing (decreasing) in her bias b_0.*

Perceived Bias.—Empirical evidence suggests that most Americans perceive media sources they distrust to be systematically biased (Gallup and Knight Foundation 2018, 2020). Consistent with this, agents in our model may perceive unbiased sources as biased and also perceive like-minded sources as less biased than they actually are.

We define an agent's *perceived bias* $\bar{\beta}_j^i$ of information source j as the agent i's limiting posterior mean on β_j.[16] We say that agent i perceives a source j to be

---

[16] Under Assumptions 1 and 2′, the agent's perceived bias of source j is $\bar{\beta}_j^i = \beta_{0j}\sqrt{1 - (b_{0i}/a_0)^2} - \frac{a_0 b_{0i} \alpha_{0j}}{a_0^2 \sqrt{1 - (b_{0i}/a_0)^2}}$.

oppositely biased if $\text{sign}\left(\bar{\beta}_j^i\right) \neq \text{sign}\left(\bar{\gamma}^i\right)$. We say that agent i perceives a source j as *less right-biased than it actually is* if $\bar{\beta}_j^i < \beta_{0j}$. Proposition 5 implies that:

COROLLARY 4: *Suppose r_t is observed, agents single-home, and Assumptions 1 and 2′ hold. Then agents R and L both perceive an unbiased source with $\alpha_{0j} > 0$ as oppositely biased. They also perceive a like-minded biased source with $\alpha_{0j} > 0$ as less biased than it actually is.*

Trust and Polarization When r_t Is Observed.—Our baseline results on trust (Corollary 1) and polarization (Corollary 2) continue to hold even when r_t is observed, by Proposition 5 and because observing r_t does not change either ρ_{is} or a_i. The one proviso is that the limit $a_0 \downarrow b_0$ can only be consistent with Assumption 2′ if $\alpha_{0j} = 0$ for all j. When a_0 is bounded away from b_0, however, it will continue to be the case that trust in perfectly biased sources diverges as a_0 approaches b_0, and polarization consequently grows large.

IV. Overconfidence

We now allow for the possibility that agents are uncertain about the accuracy a_i of their own reasoning. Because this will enable a biased agent to rationalize any signal she might observe via a higher value of a_i, it allows us to dispense with restrictions on the parameter space (Assumptions 2 and 2′). We continue to assume that the ideological valence r_t is observed in every period.

We now revise Assumption 1 to allow the agent to entertain any $a_i \in (0, a_i^{max}]$.

ASSUMPTION 1′: *The support of $\mu_{0,\theta}^i$ is the set $\Theta_i^{prior} \subset \Theta$ for which $b_i = 0$ and $a_i \in (0, a_i^{max}]$, where $a_i^{max} \geq \sqrt{a_0^2 + b_0^2}$.*

Agents' identified sets now may include multiple values of a_i, and this means that they may include multiple values of α, β, and γ as well.

PROPOSITION 6: *Suppose r_t is observed and agents single-home. Under Assumption 1′, agent i's identified set, $I_i(R_{0i})$, is nonempty and consists of all $\theta_i \in \Theta$ such that $a_i \in \left[\underline{a}_i, a_i^{max}\right]$, $b_i = 0$, $\alpha = \rho_{is}/a_i$, $\gamma = \rho_{ir}/a_i$, and $\beta = \frac{1}{\sqrt{1 - \gamma^2}}(\rho_{rs} - \gamma\alpha)$, where $\underline{a}_i = \max_j \sqrt{\zeta_{ij}}$ and ζ_{ij} is the population R^2 from a regression of x_{it} on r_t and s_{jt}.*

Proposition 6 is similar to Proposition 5, but now the identified set contains a range of values for a_i, the accuracy of the agent's reasoning. The lower bound $\underline{a}_i$ depends on the extent to which x_{it} is correlated with the other data the agent observes. Under the agent's maintained assumption that $b_i = 0$, the only way that the variation in x_{it} can be jointly explained by r_t and a particular s_{jt} is if the correlation of x_{it} with the true state is relatively large. Thus, increasing the maximum R^2 that can be achieved by regressing x_{it} on (r_t, s_{jt}) increases the lower bound on the possible values of a_i.

Proposition 6 also shows that an agent's own accuracy a_i and the accuracy α of the information sources are not separately identified. Since agents believe that $b_i = 0$, we have that $\rho_{is} = a_i\alpha$ for any $\theta_i \in I_i(R_{0i})$. A given value of ρ_{is} could result from a high value of a_i and low values of α_j or a low value of a_i and high values of α_j; these cannot be distinguished by the observed data. Beliefs within the identified set will therefore remain proportional to the prior.

Overconfidence.—A large literature in psychology, economics, and finance has documented overconfidence in many contexts, with early evidence including the pioneering study of Alpert and Raiffa (1982). Ortoleva and Snowberg (2015) explore in detail the implications of overconfidence for political behavior. While overconfidence is a primitive in their model, the results of this section show that it may arise endogenously as a consequence of other biases in reasoning.[17]

We refer to agent i's belief about the accuracy a_i of her own reasoning as her *confidence*. We say that she is *overconfident* if $a_i > a_0$ for all $\theta_i \in I_i(R_{0i})$ and *underconfident* if $a_i < a_0$ for all $\theta_i \in I(R_{0i})$.

Proposition 6 shows that agents are never underconfident and will be overconfident if and only if $\underline{a}_i > a_0$. The bound $\underline{a}_i$ is determined by the R^2 values ζ_{ij} from regressions of x_{it} on r_t and s_{jt}. Intuitively, agents are overconfident when they observe correlations between x_{it}, r_t, and at least one s_{jt} that are infeasible at the true value a_0 (under the agent's maintained assumption that $b_i = 0$). The proof of Proposition 5 shows that this R^2 is given by

$$(11) \qquad \zeta_{ij} = b_{0i}^2 + a_0^2\left(\frac{\alpha_{0j}^2}{1 - \beta_{0j}^2}\right).$$

Therefore, agents R and L are overconfident if and only if

$$(12) \qquad \frac{b_{0i}^2}{a_0^2} > 1 - \max_j\left\{\frac{\alpha_{0j}^2}{1 - \beta_{0j}^2}\right\}.$$

Noting that the fraction in curly braces approaches one as $\alpha_{0j}^2 + \beta_{0j}^2$ approaches one yields the following result.

COROLLARY 5: *Suppose that r_t is observed, agents single-home, and Assumption $1'$ holds. Then agents are never underconfident, Agent U is never overconfident, and Agents R and L are overconfident if (i) a_0 is sufficiently small or (ii) there is some source j with $\alpha_{0j}^2 + \beta_{0j}^2$ sufficiently close to one.*

---

[17] In Ortoleva 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 but also has a further effect on polarization via endogenous trust.

Overconfidence emerges in our model in order to reconcile the agents' dogmatic belief that their reasoning x_{it} is unbiased with the observed correlations of the sources. No agent is ever underconfident, since it is always possible to rationalize the observed correlations with a sufficiently high value of a_i. Overconfidence can arise if an agent's bias is large. However, even when an agent's bias is small, she is overconfident if the accuracy of her reasoning a_0 is sufficiently low. She is also overconfident if at least one source in the market is sufficiently accurate. When there is at least one source with $\alpha_{0j} \neq 0$ and $\alpha_{0j}^2 + \beta_{0j}^2 = 1$, agents with any bias are overconfident regardless of the value of a_0, since $\underline{a}_i = \sqrt{a_0^2 + b_{0i}^2}$.

Trust, Polarization, Ideology, and Perceived Bias in the Presence of Overconfidence.—Can divergence in trust, disagreement about ω_t, ideology, and perceived bias be large once we allow for overconfidence? The answer is yes if a_0 is small relative to b_0. To illustrate, we focus on the case where the upper bound of the agent's prior on a_i is $a_i^{max} = \sqrt{a_0^2 + b_{0i}^2}$. This is the smallest value for which an agent's identified set is always nonempty, and it approximates a situation where each agent's prior on her own accuracy a is concentrated on values that are close to the true value a_0.[18] We further focus on the case where the set of sources is sufficiently rich that it includes each agent's *trust-maximizing* source. These results are not knife-edge; they hold approximately if a_i^{max} is close to the assumed value and there is an available source that generates sufficiently high trust. More general results for trust and beliefs in the presence of overconfidence are derived in Appendix D.

Because we have relaxed Assumptions 1 and 2, agent i's trust-maximizing source is defined by the accuracy and bias (α_i, β_i) that maximizes the agent's trust over all pairs satisfying the feasibility condition $\alpha_i^2 + \beta_i^2 \leq 1$. We show in Appendix D that this is

$$(13) \qquad \left(\alpha_i^{max}, \beta_i^{max}\right) \equiv \left(\frac{a_0}{\sqrt{a_0^2 + b_{0i}^2}}, \frac{b_{0i}}{\sqrt{a_0^2 + b_{0i}^2}}\right).$$

Note that the correlation between x_{it} and a trust-maximizing source is $\rho_{isj} = a_0 \alpha_i^{max} + b_{0i} \beta_i^{max} = \sqrt{a_0^2 + b_{0i}^2}$.

Agent U's trust will be maximized by an unbiased source with accuracy $\alpha_j = 1$ and bias $\beta_j = 0$. For agents R and L, the trust-maximizing source will be biased. If b_0 is close to a_0, a biased agent will prefer a source with bias and accuracy close to $1/\sqrt{2}$. If the set of available sources includes agent i's trust-maximizing source, then in the limit as $T \to \infty$, the single-homing agent observes her trust-maximizing source in all periods $t > \varepsilon T$ by Proposition 3.

If the trust-maximizing source is available, the agent will be maximally overconfident, with confidence degenerate at the maximal value $a_i^{max} = \sqrt{a_0^2 + b_{0i}^2}$. This is

---

[18] In the alternative case where each agent's prior on a is concentrated on values much larger than a_0, their initial beliefs about a are severely mistaken. Furthermore, 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 of bias into divergent trust and disagreement.

because substituting $\left(\alpha_i^{max}, \beta_i^{max}\right)$ for α_j and β_j in equation (11) yields $\zeta_{ij} = a_0^2 + b_{0i}^2$, which in turn implies $\underline{a}_i = \sqrt{a_0^2 + b_{0i}^2}$ by Proposition 6. It follows that the agent's trust in the trust-maximizing source will be one, and her posterior will be degenerate at

$$(14) \qquad \bar{\omega}_t^i = \alpha_i^{max} \omega_t + \beta_i^{max} \tilde{r}_t.$$

We can now apply Proposition 6 to see that trust becomes highly divergent in the limit as $a_0 \to 0$. The fact that overconfidence is maximal means an agent R's trust $\bar{\alpha}_j^R$ for an arbitrary source j (not necessarily her trust-maximizing source) approaches the source's bias β_j in that limit. Her trust in a perfectly right-biased source approaches one. Her trust in a perfectly left-biased source approaches zero. The opposite is true of agent L. Agents may therefore underestimate the accuracy of unbiased sources and come to trust biased sources strictly more than even a perfectly accurate source.

The agents' ideologies and perceived biases also become highly polarized in the same limit. Agent R believes that the true state is almost perfectly correlated with r_t, i.e., $\bar{\gamma}^R \to 1$, and perceives a perfectly accurate source j to be almost perfectly left-biased. Agent L correspondingly believes that the true state is almost perfectly correlated with $-r_t$, i.e., $\bar{\gamma}^L \to -1$, and perceives a perfectly accurate source to be almost perfectly right-biased.

Expected disagreement increases with the agents' bias, since $\pi = b_0^2 / \left(a_0^2 + b_0^2\right)$. This will be at least $1/2$ if $a_0 \leq b_0$, and it will approach one in the limit as a_0 approaches zero. Thus, polarization is significant if a_0 is small relative to b_0.[19]

V. 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 this section, 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.

A. Trust under Multi-homing

We first consider the agent's limiting beliefs about θ_i under multi-homing. In the *multi-homing* case, $R_{0i} = \left(\rho_{is}, \rho_{ir}, \rho_{rs}, \Sigma_0\right)$, where Σ is the matrix of correlations among elements of s_t.[20]

PROPOSITION 7: *Under Assumption 1′, agent i's identified set under multi-homing, $I_i(R_{0i})$, when r_t is observed, is nonempty and consists of all $\theta_i \in \Theta$ with $a_i \in \left[\underline{a}_i, a_i^{max}\right], b_i = 0, \alpha = \rho_{is}/a_i, \gamma = \rho_{ir}/a_i, \beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha), and \Sigma = \Sigma_0,$*

where $\underline{a}_i = \sqrt{\zeta_i}$ and ζ_i is the population R^2 from a regression of x_{it} on r_t and the elements of s_t.

Proposition 7 is similar to Proposition 6 except that in the multi-homing case, the lower bound on a_i is $\underline{a}_i$ rather than $\underline{a}_{\,j}$. Rather than depending on the *maximum* R^2 value across regressions of x_{it} on r_t and various s_{jt}, the bound now depends on the R^2 from a regression of x_{it} on r_t and *all* the s_{jt}.[21] Since $\zeta_i \geq \zeta_{ij}$ for all j, the lower bound on a_i is tighter under multi-homing than under single-homing. This reflects the fact that observing the correlation among the different signals in s_t provides further constraints on the set of parameters that could be consistent with the data.

Since the multi-homing bound on confidence $\underline{a}_i$ is weakly greater than the single-homing bound $\underline{a}_{\,j}$, 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 $\left|\bar{\alpha}_j^i - \bar{\alpha}_k^i\right|$ between any two sources can be weakly smaller under multi-homing, and that ideology $\bar{\gamma}^i$ will tend to be less extreme.

Suppose for tractability 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 her own accuracy.

ASSUMPTION 3: a_i and (α, β, γ) are independently distributed under the agent's prior μ_0^i.

When Assumption 3 holds, each agent's limiting (marginal) posterior distribution on a_i, which we denote by $\mu_{\infty,a}^i$, is the agent's prior marginal distribution on a except truncated at the lower bound of $\underline{a}_{\,i}$. It follows that multi-homing may dampen divergent trust and ideology (while also increasing confidence).

COROLLARY 6: *Suppose that Assumptions 1$'$ and 3 hold. Then, for any agent i, the difference in trust $\left|\bar{\alpha}_j^i - \bar{\alpha}_k^i\right|$ between any two sources and the ideology $\left|\bar{\gamma}^i\right|$ 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_∞^i of a multi-homing agent to be exactly the same as a single-homing agent's. For example, when at least one source with $\alpha_{0j} \neq 0$ is located on the frontier, $\underline{a}_i$ already achieves its maximal value under single-homing, so the limiting posterior under multi-homing is unchanged.

B. *Polarization under Multi-homing*

How does expected disagreement about ω_t compare in the single- and multi-homing cases? As shown in this subsection, it is possible for multi-homing to reduce disagreement, but multi-homing may also increase disagreement in some situations.

---

[21] That is, $\zeta_i = \tilde{\rho}_i' \tilde{\Sigma}^{-1} \tilde{\rho}_i$ where $\tilde{\rho}_i = \begin{pmatrix} \rho_{ir} \\ \rho_{is} \end{pmatrix}$ and $\tilde{\Sigma} = \begin{pmatrix} 1 & \rho_{rs}' \\ \rho_{rs} & \Sigma \end{pmatrix}$.

To demonstrate this, we first characterize the agent's posterior expectation of ω_t when her beliefs about θ_i are given by the limiting posterior $\mu_{\infty,\theta}^i$. Under multi-homing, the average belief $\bar{\omega}_t^i$ is now a linear function of the observed signals, as shown below.

LEMMA 1: *Suppose that agent i multi-homes and Assumption 1' holds. As $T \to \infty$, in any period $t > \varepsilon T$, the mean of her posterior on ω_t given s_t (but not x_{it} or r_t) is*

$$(15) \qquad\qquad\qquad \bar{\omega}_t^i = A_i \rho_{is}' \Sigma^{-1} \tilde{s}_t,$$

where $\tilde{s}_t$ is the J-vector of standardized signals s_t and the amplification factor A_i is given by

$$(16) \qquad\qquad\qquad A_i = \int_{\underline{a}_i}^{a_i^{max}} \frac{1}{a} d\mu_{\infty,a}^i(a).$$

We consider a special case wherein each agent may observe three sources with respective biases of β, 0, and $-\beta$, where $\beta > 0$. All of these sources are on the *frontier*, so $\alpha_j^2 + \beta_j^2 = 1$ for all j. As shown in Lemma 2 below, a multi-homing agent's posterior mean $\bar{\omega}_t^i$ is the same as that of a single-homing agent who observes her trust-maximing source. The reason is that a multi-homing agent observing two distinct frontier sources can construct a linear combination of the sources' signals whose value will be equal to the signal of the agent's trust-maximizing source. Appendix G show that we obtain the same result in the limit of a sequence of random markets with nonfrontier sources.

LEMMA 2: *Suppose that Assumption 1' holds and there are at least two frontier sources with distinct biases. Then the mean of the multi-homing agent's posterior on ω_t given s_t is*

$$(17) \qquad\qquad\qquad \bar{\omega}_t^i = \alpha_i^{max} \omega_t + \beta_i^{max} \tilde{r}_t.$$

It follows, as formalized in Proposition 8 below, that multi-homing does not in general reduce expected disagreement. In fact, multi-homing may make it worse. To see this, let $\phi_b = \tan^{-1}(b_0/a_0)$ denote the angle between vectors $(1,0)$ and (a_0, b_0), and let $\phi_\beta = \sin^{-1}(\beta)$ denote the angle between $(1,0)$ and $\left(\sqrt{1-\beta^2}, \beta\right)$. If $\phi_b \in \left(0, \frac{1}{2}\phi_\beta\right)$, all single-homers observe the unbiased source, while the multi-homers' beliefs are the same as would occur if they observed their trust-maximizing source, so multi-homing results in greater disagreement. If instead $\phi_b \in \left(\frac{1}{2}\phi_\beta, \phi_\beta\right)$, biased single-homers observe a source with more bias than the trust-maximizing source, so multi-homing results in less disagreement.

PROPOSITION 8: *Suppose Assumption 1' holds and there are three frontier sources with respective biases, β, 0, and $-\beta$. Then expected disagreement π is greater under*

multi-homing than under single-homing if $\phi_b \in \left(0, \frac{1}{2}\phi_\beta\right)$, but smaller if $\phi_b \in \left(\frac{1}{2}\phi_\beta, \phi_\beta\right)$.

VI. Endogenous Media Bias

What does our model imply about media competition and political behavior? In this extension, we show that media competition can intensify disagreements in a population with ideological biases. In Appendix H, we show that ideological bias results in interpersonal mistrust and creates welfare losses in strategic games of collective decision-making.

To explore how media competition affects ideological disagreement, we endogenize the accuracies and biases of the information sources in a sequential positioning game. We consider a unit mass of agents. The agents are divided into three types $i \in \{U, R, L\}$, with accuracies a_i and biases b_i defined in Section IB. We assume that agents have the same priors as in our full model in Section IV, so Assumption 1$'$ holds. Each type i has mass $m_i > 0$ such that $m_U + m_R + m_L = 1$. For simplicity, we assume that $m_L = m_R$.

A set of E identical potential entrants sequentially choose whether or not to enter before the first period. If they enter, they may choose any accuracy α_{0j} and bias β_{0j} on the frontier (i.e., $\alpha_{0j}^2 + \beta_{0j}^2 = 1$). Prior to entry, each entering outlet observes all preceding entrants' choices of $(\alpha_{0j}, \beta_{0j})$. We use subgame perfect equilibrium as our solution concept.

All agents are single-homers who choose a single outlet j to observe in a given period t. We focus on media viewership choices during their exploitation period. We assume that agents have beliefs about the accuracies of the outlets during these periods corresponding to the limiting posterior $\mu_{\infty,\theta}^i$.[22] If Proposition 3 identifies multiple potential sources to observe, agents randomize between them with equal probability.

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

Let $\mathcal{J}_i$ be the set of outlets for which an i-type agent's trust $\bar{\alpha}_j^i$ is highest. Let $\xi(\bar{\alpha}_j^i)$ denote revenue per viewer of type i. 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 $\lambda > 0$. Each firm j thus has expected profit

$$(18) \qquad \Pi_j = \sum_{i \in \{U, R, L\}} \mathbf{1}\{j \in \mathcal{J}_i\} \frac{m_i}{|\mathcal{J}_i|} \xi(\bar{\alpha}_j^i) - \lambda,$$

---

[22] Strictly speaking, this is a behavioral assumption that the agents' inferences about media outlet accuracy condition not on the equilibrium strategies chosen by the outlets but only on the signals the outlets produce.

where $\mathbf{1}\{j \in \mathcal{J}_i\}$ is an indicator for whether outlet j is in the set of outlets that type-i agents observe, and $m_i/|\mathcal{J}_i|$ measures the probability of observing j within that set. Note that both $\mathcal{J}_i$ and $\bar{\alpha}_j^i$ are equilibrium outcomes that depend on the accuracy and bias choices of all media outlet entrants.

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

PROPOSITION 9: *Suppose that 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 but have expected disagreement $\pi = 0$.*

Proposition 9 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 bias, 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 $\alpha_{0j} = 1$ and $\beta_{0j} = 0$. 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 focus on the case where the set of potential entrants is sufficiently large that there are potential entrants who do not find it profitable to enter in equilibrium. We can see from Appendix D (Lemma 7) that sources gain maximum trust from biased agents by choosing those agents' trust-maximizing level of bias. It is then unsurprising that in the case of competition, some sources choose to be biased and successfully retain a large audience.

PROPOSITION 10: *For λ sufficiently low and a set of potential entrants $E(\lambda)$ sufficiently large, all entrant outlets locate at positions on the frontier with $\beta_{0j} \in \{\beta^L, 0, \beta^R\}$, where β^L and β^R are the trust-maximizing biases 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 $\pi = 1/2$ if $a_0 \leq b_0$, and will approach $\pi = 1$ in the limit as $a_0 \to 0$. Thus, the entry of partisan media leads to greater divergence in beliefs.*

In contrast to the monopoly case, there is now significant disagreement in the population. This stems 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 bias, and this bias approaches ± 1 as the ratio of b_0 to a_0 increases, competition can potentially give rise to maximal disagreement and perfectly negatively correlated beliefs. Note that these results are not at all dependent on our assumption that $m_R = m_L$.

VII. 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 these accuracies by comparing their own reasoning about the states to the sources' reports.

Our contribution is to characterize how endogenous inference about the accuracy of sources can cause small biases in reasoning to be amplified into significantly divergent beliefs about facts, even when accurate information is commonly observed. Our model generates a large set of novel and sharp predictions about the resulting beliefs: Partisans end up trusting unreliable but ideologically aligned sources more than accurate neutral sources, and become overconfident in their own reasoning. 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.

Taken together, these results highlight the outsized importance of trust in driving ideological differences in society. If individuals' reasoning has even a small amount of bias, then they may 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 reducing biases in the population's reasoning through scientific literacy, or increasing the prominence of commonly trusted sources—may yield larger gains.

APPENDIX A. PROOFS

A1. *Proof of Proposition 1*

Let D_{it} denote the data observed by agent i in period t in three cases.

Case 1: Agent i single-homes and does not observe r_t. We let $D_{it} = (s_{jt}, x_{it})$.

Case 2: Agent i single-homes and observes r_t. We let $D_{it} = (s_{jt}, x_{it}, r_t)$.

Case 3: Agent i multi-homes and observes r_t. We let $D_{it} = (s_t, x_{it}, r_t)$.

LEMMA 3: $D_{i1}, \ldots, D_{it}$ *is independent of* θ_{0i}, *conditional on* R_{0i} *and* V_{0i}.

PROOF:

First, consider Case 1. With slight abuse of notation, define $D_{i\tau}^{j} = (s_{j\tau}, x_{i\tau})$ to be the τ-th time that j's signal and reasoning are observed. The agent knows $D_{i\tau}^{j} \sim N(0, \Omega_{0ij})$ for some positive definite Ω_{0ij}, where $\Omega_{0ij} = V_{0ij}^{1/2} R_{0ij} V_{0ij}^{1/2}$, where $V_{0ij} = \text{diag}(\Omega_{0ij})$ and R_{0ij} is the correlation matrix for $D_{i\tau}^{j}$, which in Case 1 is the correlation ρ_{ij}. Independence across periods then generalizes the result to the entire sequence of data observations. Cases 2 and 3 can be proven in the same way. ∎

Let $P_{Y|X}^{i}$ denote the posterior distributions of Y given X and the prior μ_0^i. Then, by Lemma 3, we can see that for all $\vartheta \in \mathcal{L}_\Theta$, $P_{\theta_i|R_i, V_i, D_1, \ldots, D_t}^{i}(\vartheta) = P_{\theta_i|R_i, V_i}^{i}(\vartheta)$ and, hence, that

$$(A1) \qquad P_{\theta_i|D_{i1}, \ldots, D_{it}}^{i}(\vartheta) = \int_{\mathcal{R}, \mathcal{V}} P_{\theta_i|R_i, V_i, D_{i1}, \ldots, D_{it}}^{i}(\vartheta) \, dP_{R_i, V_i|D_{i1}, \ldots, D_{it}}^{i},$$

$$(A2) \qquad = \int_{\mathcal{R}, \mathcal{V}} P_{\theta_i|R_i, V_i}^{i}(\vartheta) \, dP_{R_i, V_i|D_{i1}, \ldots, D_{it}}^{i}.$$

We now characterize the limit of $P_{R_i, V_i|D_1, \ldots, D_t}^{i}$. First, consider Case 3, where agent i multi-homes and observes r_t. Let $P_{D_i|R_i, V_i}$ denote the (true) distribution of D_{it}, conditional on R_i and V_i. The experiment $(P_{D_i|R_i, V_i}: R_i \in \mathcal{R}; V_i \in \mathcal{V})$ is Gaussian with known mean zero and known variance, and its parameter space $(\mathcal{R}, \mathcal{V})$ is compact. It is straightforward to verify the following regularity conditions: (i) $P_{D_i|R_i, V_i} \neq P_{D_i|(R_i, V_i)'}$ for any $(R_i, V_i) \neq (R_i, V_i)'$, (ii) the mapping $(R_i, V_i) \mapsto P_{D_i|R_i, V_i}$ is continuous in total variation norm, (iii) $P_{D_i|R_i, V_i}$ has a nonsingular information matrix $I_{R_{0i}, V_{0i}}$ at (R_{0i}, V_{0i}) (recalling from Remark 1 that we focus on θ_0 such that $R_{0i} \in \text{int}(\mathcal{R})$), and (iv) $(P_{D_i|R_i, V_i}: R_i \in \mathcal{R}; V_i \in \mathcal{V})$ is differentiable in quadratic mean at (R_{0i}, V_{0i}). Then, by van der Vaart's (2000) lemma 10.6 and theorem 10.1 (the Bernstein–von Mises theorem), the limit of $P_{R_i, V_i|D_{i1}, \ldots, D_{it}}^{i}$ as $t \to \infty$ is a distribution degenerate at the true correlations (R_{0i}, V_{0i}).

Next, consider Case 1, where agent i single-homes and does not observe r_t. Reorder the experimentation periods so that those where the agent observes (s_{1t}, x_{it}) occur first, those where the agent observes (s_{2t}, x_{it}) occur second, through those where the agent observes (s_{Jt}, x_{it}). Denote these respective subsequences of data by $D_i^1, \ldots, D_i^J$. Since posterior beliefs are invariant to the order of data, this reordering does not affect the limit of $P_{R_i, V_i|D_{i1}, \ldots, D_{it}}^{i}$. The logic above implies that as $T \to \infty$, the agent's posterior belief $P_{R_i, V_i|D_{i1}, \ldots, D_{it}}^{i}$ at the end of the first set of periods converges to a limit whose marginal distribution on ρ_{i1} is degenerate at the true values of these correlations. Note that for every finite t, $P_{R_i, V_i|D_i^1}^{i}$ has continuous density on $(\mathcal{R}, \mathcal{V})$ and so is a valid prior under our model. Applying the same logic again then implies that the agent's posterior belief $P_{R_i, V_i|D_i^1, D_i^2}^{i}$ at the end of the second set of periods converges to a limit whose marginal distribution on (ρ_{i1}, ρ_{i2}) is degenerate at the true value of these correlations. Iterating this logic repeatedly shows that $P_{R_i, V_i|D_i^1, \ldots, D_i^J}^{i} = P_{R_i, V_i|D_{i1}, \ldots, D_{it}}^{i}$ converges to a limit whose marginal distribution $(\rho_{i1}, \ldots, \rho_{iJ})$ is degenerate at the full vector of true correlations R_{0i}. Case 2 can be proven in the same way.

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

(A3) $$\mu^i_{\infty,\theta}(\vartheta) = \lim_{T \to \infty} P_{\theta_i|D_{i1},\ldots,D_{it}}(\vartheta),$$

(A4) $$= \lim_{T \to \infty} \int_{\mathcal{R},\mathcal{V}} P^i_{\theta_i|R_i,V_i}(\vartheta) dP^i_{R_i,V_i|D_{i1},\ldots,D_{it}},$$

(A5) $$= P^i_{\theta_i|R_{0i},V_{0i}}(\vartheta),$$

(A6) $$= \mu^i_{0,\theta}(\vartheta|R_i = R_{0i}),$$

(A7) $$= \frac{\mu^i_{0,\theta}(\vartheta \cap I_i(R_{0i}))}{\mu^i_{0,\theta}(I_i(R_{0i}))},$$

where the third equality uses the convergence of $P^i_{R_i,V_i|D_{i1},\ldots,D_{it}}$.

A2. *Proof of Proposition* 2

Given any value of $\theta_i \in I_i(R_{0i})$, the state ω_t and the (normalized) signals $\tilde{s}_{jt}$ are jointly distributed:

$$\begin{pmatrix} \omega_t \\ \tilde{s}_{jt} \end{pmatrix} \sim N\left(0, \begin{bmatrix} 1 & \alpha_j \\ \alpha_j & 1 \end{bmatrix}\right).$$

The conditional expectation of ω_t given $\tilde{s}_{jt}$ (under single-homing) is then $\alpha_j \tilde{s}_{jt}$, by the properties of the multivariate normal distribution. The desired result follows from taking the expectation over the limiting posterior $\mu^i_{\infty,\theta}$.

A3. *Proof of Proposition* 3

The expected loss for a given j is the same for all t in the exploitation period, so we can focus on minimizing the single-period loss for a single agent. Thus, we can drop the i and t subscripts. For notational simplicity, we also focus on the simple case where all variances are known to be one. The agent solves

$$\min_{j \in \mathcal{J}_+} E_{s_j,x,\omega}\left[\left(d^*_j(s_j,x) - \omega\right)^2\right],$$

where the expectation is taken under the distribution of (s,x,ω) and $d^*_j(s_j,x)$ is the optimal decision after seeing (s_j,x). Note that $d^*_j(s_j,x) = E_{\omega|s_j,x}[\omega]$. The law of iterated expectations implies that

$$E_{s_j,x,\omega}\left[\left(d^*_j(s_j,x) - \omega\right)^2\right] = 1 - E_{s_j,x}\left[d^*_j(s_j,x)^2\right].$$

Note that in all cases here, expectations are taken under the joint distribution of (Ω,s,x,ω) given $\Omega \sim \mu_\infty$. Thus, an expression like $E_\omega[\cdot]$ refers to the expectation under the marginal distribution of ω in that distribution.

Define $d_j^*(s_j, x, \Omega)$ to be the optimal decision *conditional* on a particular Ω. Note that $d_j^*(s_j, x, \Omega) = E_{\omega|s_j, x, \Omega}[\omega]$, so we have

$$d_j^*(s_j, x, \Omega) = \binom{a}{\alpha_j}' \begin{bmatrix} 1 & \alpha_j a \\ \alpha_j a & 1 \end{bmatrix}^{-1} \binom{x}{s_j} = \binom{a}{\alpha_j}' \begin{bmatrix} 1 & \rho_j \\ \rho_j & 1 \end{bmatrix}^{-1} \binom{x}{s_j},$$

where the last line follows from observing that for all (a, α_j) in the support of μ_∞, we must have $\alpha_j a = \rho_j$, where ρ_j is a constant equal to the empirical correlation observed in the data. Since $E_{\omega|s_j, x}[\omega] = E_\Omega E_{\omega|s_j, x, \Omega}[\omega] = E_\Omega d^*(s_j, x, \Omega)$, we have

$$d_j^*(s_j, x) = \binom{\bar{a}}{\bar{\alpha}_j}' \begin{bmatrix} 1 & \rho_j \\ \rho_j & 1 \end{bmatrix}^{-1} \binom{x}{s_j}.$$

It follows from some algebraic manipulation that

$$E_{s_j, x}\left[d_j^*(s_j, x)^2\right] = \binom{\bar{a}}{\bar{\alpha}_j}' \begin{bmatrix} 1 & \rho_j \\ \rho_j & 1 \end{bmatrix}^{-1} \binom{\bar{a}}{\bar{\alpha}_j} = \frac{\bar{a}^2 - 2\rho_j \bar{a} \bar{\alpha}_j + \bar{\alpha}_j^2}{1 - \rho_j^2}.$$

By Jensen's inequality,

$$\bar{\alpha}_j = E_\Omega[\rho_j/a] \geq \frac{\rho_j}{\bar{a}},$$

so we have $\rho_j = c\bar{a}\bar{\alpha}_j$ for some $c \leq 1$, holding μ_∞ fixed. Therefore, if we compare the expected variance in decisions when the agent observes sources with different trust $\bar{\alpha}_j$, holding μ_∞ fixed, we have that

$$\frac{\partial E_{s_j, x}\left[d_j^*(s_j, x)^2\right]}{\partial \bar{\alpha}_j} = \frac{\partial}{\partial \bar{\alpha}_j}\left[\frac{\bar{a}^2 - 2c\bar{a}^2 \bar{\alpha}_j^2 + \bar{\alpha}_j^2}{1 - c^2\bar{a}^2 \bar{\alpha}_j^2}\right]$$

$$\geq \frac{\partial}{\partial \bar{\alpha}_j}\left[\frac{\bar{a}^2 - 2\bar{a}^2 \bar{\alpha}_j^2 + \bar{\alpha}_j^2}{1 - \bar{a}^2 \bar{\alpha}_j^2}\right]$$

$$> 0,$$

where the first inequality follows from $c \leq 1$ and the second inequality is strict under the maintained assumption that $\bar{a} < 1$. This then implies that for any j and k such that $\bar{\alpha}_j^2 > \bar{\alpha}_k^2$, we have

$$E_{s_j, x, \omega}\left[\left(d_j^*(s_j, x) - \omega\right)^2\right] < E_{s_k, x, \omega}\left[\left(d_k^*(s_k, x) - \omega\right)^2\right],$$

and so the problem is solved by choosing a j with the highest value of $\bar{\alpha}_j^2$.

A4. *Proof of Proposition 4*

Recall that $I_i(R_{0i}) = \{\theta_i \in \Theta_i^{prior} : a_i\alpha + b_i\beta = \rho_{is}\}$, where any $\theta_i \in \Theta_i^{prior}$ must give rise to a positive semidefinite correlation matrix $\tilde{\Omega}(\theta_i)$ for $(\omega_t, r_t, x_{it}, s_t)$.

Assumption 1 implies that $\theta_i \in I_i(R_{0i})$ if and only if $a_i = a_0, b_i = 0, \alpha = \rho_{is}/a_i$, and $\tilde{\Omega}(\theta_i)$ is positive semidefinite. Assumption 2 implies that $|\alpha_{0j}| = |\rho_{ij}/a_{0i}| = |\alpha_{0j} + (b_{0i}/a_{0i})\beta_{0j}| \leq \|\alpha_{0j}| + |\beta_{0j}\| \leq 1$. We can then pick $\beta = 0, \gamma = 0$, and assume that s_{jt} are mutually independent and independent of r_t and x_{it}, conditional on ω_t, so $\Sigma = \mathrm{corr}(s_t) = \alpha\alpha' + K$, where K is a diagonal matrix with entries equal to $1 - \alpha_j^2$. It follows that $\tilde{\Omega}(\theta_i)$ is positive semidefinite, so $I_i(R_{0i})$ is nonempty.

A5. *Proof of Proposition 5*

The proof is similar to the proof of Proposition 4, except r_t is now observed. By Assumption 1, Remark 1, and the definition of $I_i(R_{0i})$, $\theta_i \in I_i(R_{0i})$ if and only if $a_i = a_0, b_i = 0, \alpha = \rho_{is}/a_i, \gamma = \rho_{ir}/a_i, \beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$, and $\tilde{\Omega}(\theta_i)$ is positive semidefinite. We first show that the correlation matrix for (ω_t, r_t, s_{jt}), i.e., $\begin{bmatrix} 1 & \tilde{\alpha}_{ij} \\ \tilde{\alpha}_{ij} & \tilde{\Sigma}_j \end{bmatrix}$, where $\tilde{\alpha}_{ij} = \begin{pmatrix} \rho_{ir}/a_i \\ \rho_{ij}/a_i \end{pmatrix}$ and $\tilde{\Sigma}_j = \begin{bmatrix} 1 & \rho_{rj} \\ \rho_{rj} & 1 \end{bmatrix}$, is positive semidefinite for all j. This is equivalent to $1 - \tilde{\alpha}_{ij}'\tilde{\Sigma}_j^{-1}\tilde{\alpha}_{ij} \geq 0$ by standard matrix results (see Boyd and Vandenberghe 2004, appendix A.5.5). This in turn requires that for all j,

$$(A8) \qquad a_i^2 \geq \zeta_{ij} = \frac{\rho_{ir}^2 + \rho_{xj}^2 - 2\rho_{rj}\rho_{ij}\rho_{ir}}{1 - \rho_{rj}^2}.$$

Algebraic substitution shows that $\zeta_{ij} = b_0^2 + a_0^2\left(\frac{\alpha_{0j}^2}{1 - \beta_{0j}^2}\right)$. Since $a_i = a_0$, this condition is guaranteed by Assumption 2'. We can then assume that s_{jt} are mutually independent and independent of x_{it} conditional on ω_t and r_t, so $\Sigma = \alpha\alpha' + \beta\beta' + K$, where K is a diagonal matrix with entries equal to $1 - \alpha_j^2 - \beta_j^2$. It follows that $\tilde{\Omega}(\theta_i)$ is positive semidefinite, so $I_i(R_{0i})$ is nonempty.

A6. *Proof of Proposition 6*

The proof is similar to the proof of Proposition 5, except the agent may become overconfident. By Assumption 1', Remark 1, and the definition of $I_i(R_{0i})$, $\theta_i \in I_i(R_{0i})$ if and only if $a_i \leq a_i^{max}, b_i = 0, \alpha = \rho_{is}/a_i, \gamma = \rho_{ir}/a_i, \beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$, and $\tilde{\Omega}(\theta_i)$ is positive semidefinite. By the same logic as the proof of Proposition 5, the last condition is true if and only if $a_i \geq \underline{a}_i = \max_j \sqrt{\zeta_{ij}}$. Therefore, $\theta_i \in I_i(R_{0i})$ if and only if $a_i \in [\underline{a}_i, a_i^{max}], b_i = 0, \alpha = \rho_{is}/a_i, \gamma = \rho_{ir}/a_i$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$.

A7. *Proof of Proposition 7*

This proof is the same as the proof of Proposition 6, except that agent i multi-homes. By Assumption 1', Remark 1, and the definition of $I_i(R_{0i})$, $\theta_i \in I_i(R_{0i})$

if and only if $a_i \leq a_i^{max}, b_i = 0, \alpha = \rho_{is}/a_i, \gamma = \rho_{ir}/a_i, \beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$, and $\tilde{\Omega}(\theta_i)$ is positive semidefinite. Since the agent believes that $x_{it} = a_i\omega_t + \eta_{it}$ for some $a_i \in (0, a_i^{max}]$, $\tilde{\Omega}(\theta_i)$ is positive semidefinite if and only if the covariance matrix for (ω_t, r_t, s_t), given by $\begin{bmatrix} 1 & \tilde{\alpha}' \\ \tilde{\alpha} & \tilde{\Sigma} \end{bmatrix}$, where $\tilde{\alpha} = \begin{pmatrix} \rho_{ir}/a_0 \\ \rho_{is}/a_0 \end{pmatrix}$ and $\tilde{\Sigma} = \begin{bmatrix} 1 & \rho_{rs}' \\ \rho_{rs} & \Sigma \end{bmatrix}$, is postive semidefinite. This holds if and only if $a_i^2 \geq \zeta_i = \tilde{\rho}_i' \tilde{\Sigma}^{-1} \tilde{\rho}_i$, where $\tilde{\rho}_i = (\rho_{ir} \ \rho_{is}')'$, by Boyd and Vandenberghe (2004), appendix A.5.5. Therefore, $\theta_i \in I_i(R_{0i})$ if and only if $a_i \in [\underline{a}_i, a_i^{max}]$, $b_i = 0$, $\alpha = \rho_{is}/a_i$, $\gamma = \rho_{ir}/a_i$, and $\beta = \frac{1}{\sqrt{1-\gamma^2}}(\rho_{rs} - \gamma\alpha)$, where $\underline{a}_i = \sqrt{\zeta_i}$.

A8. *Proof of Lemma 1*

Given any value of $\theta_i \in I_i(R_{0i})$, the state ω_t and the (normalized) signals $\tilde{s}_{jt}$ are jointly distributed:

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

Under multi-homing, the conditional expectation of ω_t given s_t is then $\alpha'\Sigma^{-1}\tilde{s}_t$. By Proposition 7, $\alpha = \rho_{is}/a_i$. Furthermore, $I_i(R_{0i})$ includes all $a_i \in [\underline{a}_i, a_i^{max}]$. The desired result then follows from taking the expectation over the limiting posterior $\mu_{\infty,a}^i$.

A9. *Proof of Lemma 2*

We first prove the following Lemma.

LEMMA 4: *Suppose that the signals s_{jt} are mutually independent, conditional on ω_t and r_t. Then, in the multi-homing case, we have*

$$\rho_{is}'\Sigma^{-1}\tilde{s}_t = y'Z(Z'Z + K)^{-1}(Z'\varphi_t + \varepsilon_t)$$

and

$$\rho_{is}'\Sigma^{-1}\rho_{is} = y'Z(Z'Z + K)^{-1}Z'y,$$

where $y = [a_0 \ b_{0i}]'$, Z is the $2 \times J$ matrix where the j-th column is $[\alpha_{0j} \ \beta_{0j}]'$, K is a diagonal matrix such that the j-th diagonal is $\kappa_{0j}^2 = 1 - \alpha_{0j}^2 - \beta_{0j}^2$, $\varphi_t = [\omega_t \ \tilde{r}_t]'$, and ε_t is the J-vector of $\varepsilon_{jt} = s_{jt} - \alpha_{0j}\omega_t - \beta_{0j}\tilde{r}_t$.

PROOF:
The lemma follows from noting that $\rho_{is} = Z'y$, $\Sigma = Z'Z + K$, and $s_t = Z'\varphi_t + \varepsilon_t$. ∎

When all sources 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 4 then implies that $\rho'_{is}\Sigma^{-1}s_t = a_0\omega_t + b_{0i}r_t$. Note that the R^2 of the population regression of x_{it} on s_t is $\rho'_{is}\Sigma^{-1}\rho_{is}$, which is equal to $a_0^2 + b_{0i}^2$ by Lemma 4. Therefore, the R^2 of the population regression of x_t on s_t and r_t must be weakly greater than $a_0^2 + b_{0i}^2$. However, the R^2 from a regression of x_{it} on s_t and r_t cannot exceed the R^2 from a regression of x_{it} on r_t and ω_t, which is $a_0^2 + b_{0i}^2$. Therefore, $\underline{a}_i = \sqrt{\zeta_i} = \sqrt{a_0^2 + b_{0i}^2}$. Since $a_i^{max} = \sqrt{a_0^2 + b_{0i}^2}$ by assumption, we have that $A_i = 1/\sqrt{a_0^2 + b_{0i}^2}$.

Next, consider the case with additional nonfrontier sources. Let the vector of signals of the frontier sources be s_t^F. Note that the elements of $\rho'_{is}\Sigma^{-1}$ are the coefficients from a population regression of x_t on the elements of s_t. Further note that $x_{it} = a_0\omega_t + b_{0i}\tilde{r}_t + \eta_t$, where η_t is orthogonal to ε_{jt} for all j, while each element of s_t^F is a linearly independent linear combination of ω_t and r_t. Thus, x_t is orthogonal to s_{jt} conditional on s_t^F for all nonfrontier sources. By the Frisch–Waugh–Lovell theorem, the elements of $\rho'_{is}\Sigma^{-1}$ corresponding to nonfrontier sources must be equal to zero, and the elements of $\rho'_{is}\Sigma^{-1}$ corresponding to frontier sources are the same as in the case with all frontier sources. We conclude that $\rho'_{is}\Sigma^{-1}\tilde{s}_t = a_0\omega_t + b_{0i}r_t$. Furthermore, we can conclude that $\underline{a}_i = \sqrt{a_0^2 + b_{0i}^2}$ using the same argument as in the previous paragraph.

A10. *Proof of Proposition 8*

In the multi-homing case, there are two or more frontier sources, so $\underline{a}_i = \sqrt{a_0^2 + b_{0i}^2}$, $\bar{\alpha}_j^i = 1$, and $\bar{\omega}_t^i = \alpha_i^{max}\omega_t + \beta_i^{max}\tilde{r}_t$. Expected disagreement is $b_0^2/(a_0^2 + b_0^2) = \sin^2(\phi_b)$. In the single-homing case, r_t is observed along with a frontier source, so $\underline{a}_i = \sqrt{a_0^2 + b_{0i}^2}$ and $\bar{\alpha}_j^i = (a_0\alpha_{0j} + b_{0i}\beta_{0j})/\sqrt{a_0^2 + b_{0i}^2}$. Note that a biased agent observing a biased source results in trust equal to $\cos(\phi_b - \phi_\beta)$, while a biased agent observing an unbiased source results in trust equal to $\cos(\phi_b)$. If $\phi_b < \phi_\beta/2$, then the agent's trust in the unbiased source is higher, so all single-homers observe the unbiased source. This implies that $\pi = 0$ under single-homing, while disagreement is positive under multi-homing. If $\phi_b > \phi_\beta/2$, then biased single-homers observe the similarly biased source. In this case, disagreement under multi-homing is weakly greater than disagreement under single-homing if and only if $\sin^2(\phi_b) \geq \cos^2(\phi_b - \phi_\beta)\sin^2(\phi_\beta)$, which can be shown to be true if and only if $\phi_b \geq \phi_\beta$ using trigonometric identities and by noting that $0 < \phi_b, \phi_\beta \leq \pi$.

A11. *Proof of Proposition 9*

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

$$\Pi_j = \sum_{i \in \{L,U,R\}} m_i \xi(\bar{\alpha}_j^i) - \lambda,$$

where $\bar{\alpha}_j^i$ is type-i consumers' trust in the monopolist.

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

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

Letting $m = m_R = m_L$, the optimal frontier location must satisfy the first-order condition that

$$(A9) \qquad \left.\frac{\partial \Pi}{\partial \beta_{0j}}\right|_{\alpha_{0j}^2 + \beta_{0j}^2 = 1} \;=\; (1 - 2\mu)\xi'(\bar{\alpha}^U)\,\delta^U(\beta_{0j})$$

$$+ m\left[\xi'(\bar{\alpha}^R)\,\delta^R(\beta_{0j}) + \xi'(\bar{\alpha}^L)\,\delta^L(\beta_{0j})\right]$$

$$= 0.$$

This condition is satisfied at $\beta_{0j} = 0$ because $\delta^U(0) = 0$, $\bar{\alpha}_j^R = \bar{\alpha}_j^L$, and $\delta^R(\beta_{0j}) = -\delta^L(\beta_{0j})$. When $\beta_{0j} > 0$, $\xi'(\bar{\alpha}_j^R) \leq \xi'(\bar{\alpha}_j^L)$ because $\xi(\cdot)$ is assumed to be concave and $\bar{\alpha}_j^R \geq \bar{\alpha}_j^L$. 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 (A9) 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 revenues at this position are strictly positive, the monopolist will enter when λ is sufficiently low. Overconfidence for biased agents follows from noting that $\underline{a}_i = \sqrt{a_0^2 + b_{0i}^2}$, and the expected disagreement result follows from noting that $\bar{\alpha}_j^R = \bar{\alpha}_j^L$.

A12. *Proof of Proposition 10*

Suppose that 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 a trust of one and, hence, positive revenue $\xi(1)$ from the associated agent type, so entry will be profitable for sufficiently low λ. Let λ be any λ small enough to support positive profit for at least two trust-maximizing outlets in every position earning revenue exclusively from their corresponding agent type,[23] i.e.,

$$\left\{\lambda \;\middle|\; \min_{i \in \{U, R, L\}} \frac{1}{2}m_i\xi(1) - \lambda > 0\right\}.$$

---

[23] Requiring enough profit for two trust-maximizing outlets in every position eliminates edge cases where the market only supports a limited number of outlets. Such outlets could instead choose to pool in one or two of the locations rather than spread across all three.

Furthermore, when these positions are occupied, an outlet at any other position earns zero revenue and strictly negative profit since $\lambda > 0$. Thus, in any equilibrium, all entrants must locate at one of these positions.

It remains to show that an equilibrium exists with at least one outlet in each position. 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 in position β^L in a market with J_L, J_U, J_R firms in the three positions. We show two properties about Π_L that will be useful later. First, once the other positions (0 and β^R) have each been occupied by at least one outlet, the specific number of such outlets no longer affect Π_L: for any $J_L, J_U, J'_U, J_R, J'_R$ such that $\min\{J_U, J'_U, J_R, J'_R\} \geq 1$,

$$\Pi_L(J_L, J_U, J_R) = \Pi_L(J_L, J'_U, J'_R).$$

Moreover, by definition, Π_L is strictly decreasing in J_L and $\Pi_L(2, J_U, J_R) > 0$ for any J_U, J_R due to our earlier choice of λ, so there exists a number of potential entrants $E_L(\lambda) < \infty$ such that $\Pi_L(E_L, J_U, J_R) < 0$ for any $J_U, J_R \geq 1$. Combining this with our preceding result, we obtain the second property: there exists a unique $J^*_L \in [2, E_L)$ such that for any $J_U, J_R \geq 1$,

$$\Pi_L(J^*_L, J_U, J_R) \geq 0$$

$$\Pi_L(J^*_L + 1, J_U, J_R) < 0.$$

This unique J^*_L is the threshold beyond which entry into position β^L is no longer profitable. Let Π_U and Π_R denote similar objects for positions 0 and β^R, where similar arguments show that these two properties hold as well. Let J^*_U and J^*_R denote their counterparts to J^*_L.

The tuple (J^*_L, J^*_U, J^*_R) of outlets is an equilibrium if the following conditions hold for Π_L:

$$\Pi_L(J^*_L, J^*_U, J^*_R) \geq 0$$

$$\Pi_L(J^*_L + 1, J^*_U, J^*_R) < 0$$

$$\Pi_L(J^*_L, J^*_U, J^*_R) \geq \Pi_U(J^*_L - 1, J^*_U + 1, J^*_R)$$

$$\Pi_L(J^*_L, J^*_U, J^*_R) \geq \Pi_R(J^*_L - 1, J^*_U, J^*_R + 1),$$

and similar conditions hold for Π_U and Π_R. The first two conditions follow from the definition of J^*_L. Next, at (J^*_L, J^*_U, J^*_R), we have $\Pi_L(J^*_L, J^*_U, J^*_R) \geq 0$ by definition of J^*_L and $0 > \Pi_U(J^*_L - 1, J^*_U + 1, J^*_R)$ by definition of J^*_U (noting that $J^*_L \geq 2$ and so $J^*_L - 1 \geq 1$). Hence, $\Pi_L(J^*_L, J^*_U, J^*_R) > \Pi_U(J^*_L - 1, J^*_U + 1, J^*_R)$, giving us the third condition. The fourth condition follows similarly. The same arguments prove that the corresponding conditions hold for Π_U and Π_R. The tuple (J^*_L, J^*_U, J^*_R) is also the unique equilibrium: By uniqueness of (J^*_L, J^*_U, J^*_R), no other candidate tuple can

satisfy the first two conditions simultaneously. Thus, at least one outlet chooses each of the positions $\{\beta^L, 0, \beta^R\}$ in equilibrium.

Applying Corollary 5 shows that agents are overconfident. Expected disagreement follows from equation (D2), noting that all outlets have a trust of one and the positions $\{\beta^L, \beta^R\}$ approach ± 1 as $a_0 \rightarrow 0$.

Appendix B. Asymptotic Beliefs When x_{it}, s_{jt}, and r_t Have Nonzero Means

We show in this Appendix that allowing the means of x_{it}, s_{jt}, and r_t to be nonzero does not affect our main results.

Suppose that

$$
\begin{bmatrix} \omega_t \\ r_t \\ x_{it} \\ s_t \end{bmatrix} \sim N(M_i, \Omega_i), \quad \text{where} \quad M_i = \begin{bmatrix} 0 \\ \bar{r} \\ \bar{x}_i \\ \bar{s} \end{bmatrix}.
$$

We assume that the set of all possible means M_i is a compact set M. As before, the set of all possible variances V_i is a compact set $\mathcal{V}$. The correlation matrix for $(\omega_t, r_t, x_{it}, s_t)$, which we denote as $\tilde{\Omega}_i$, is fully parametrized by θ_i. With slight abuse of notation, the Lebesgue space on $(\Theta, \mathcal{M}, \mathcal{V})$ is now denoted $((\Theta, \mathcal{V}, \mathcal{M}), \mathcal{L}, \nu)$.

ASSUMPTION 4: *At the beginning of the first period, agents have an absolutely continuous prior belief μ_0^i over $(\Theta, \mathcal{V}, \mathcal{M})$ with a continuous density with respect to ν.*

It is straightforward to show that with a slight modification in definitions, our core results continue to hold.

PROPOSITION 11: *Under Assumption 4 and the modification that $\tilde{s}_{jt} = (s_{jt} - \bar{s}_j)/\sqrt{\text{var}[s_{jt}]}$, Propositions 1, 2, 3, 4, 5, 6, 7, 8, 9 and 10 hold.*

PROOF:
We first show the following:

LEMMA 5: *The sequence of data realizations $D_{i1}, \ldots, D_{it}$ is independent of θ_{0i}, conditional on M_{0i}, R_{0i} and V_{0i}.*

PROOF:
First, consider Case 1 (as defined in the proof of Proposition 1). With slight abuse of notation, define $D_{i\tau}^j = (s_{j\tau}, x_{i\tau})$ to be the τ-th time j's signal and reasoning are observed. The agent knows $D_{i\tau}^j \sim N(M_{0i}, \Omega_{0ij})$ for some positive definite Ω_{0ij}, where $\Omega_{0ij} = V_{0ij}^{1/2} R_{0ij} V_{0ij}^{1/2}$, where $V_{0ij} = \text{diag}(\Omega_{0ij})$ and R_{0ij} is the correlation matrix for $D_{i\tau}^j$, which in Case 1 is the correlation ρ_{ij}. Independence across periods then generalizes the result to the entire sequence of data observations. Cases 2 and 3 can be proven in the same way. ∎

By applying Lemma 5 to the proof of Proposition 1, we can prove that as $T \to \infty$, agent i's posterior distribution on θ in any period $t > \varepsilon T$ converges to a limit $\mu_{\infty,\theta}^i$ that depends only on her prior $\mu_{0,\theta}^i$ and the correlations R_{0i}. In other words, Proposition 1 continues to hold under Assumption 4. Then, by the properties of the multivariate normal distribution, in any period t where an agent observes source j, the mean of her posterior on ω_t given s_{jt} (but not x_{it} or r_t) under $\mu_{\infty,\theta}^i$ is

$$(B1) \qquad \bar{\omega}_t^i(s_{jt}) \;=\; \bar{\alpha}_j^i \tilde{s}_{jt},$$

where $\tilde{s}_{jt} \;=\; (s_{jt} - \bar{s}_j)/\sqrt{\mathrm{var}\left[s_{jt}\right]}$ is the new standardized version of s_{jt}. This proves that Proposition 2 holds under Assumption 4 and the modified $\tilde{s}_{jt}$. The remaining propositions can be shown to hold under Assumption 4 and the modified $\tilde{s}_{jt}$ by applying the modified Propositions 1 and 2 to the respective proofs of Propositions 3, 4, 5, 6, 7, 8, 9 and 10. ∎

APPENDIX C. TRUST AND POLARIZATION WHEN AGENTS BELIEVE THAT THEIR OWN BIAS IS NONZERO

We consider a case where agents place a dogmatic prior on $b_i = b^*$ for some $b^* \neq 0$. We follow the same setup as those in Section II, namely that there exist multiple available sources but agents are single-homing and ideological valence r_t is never observed. As in Section II, we make the following assumptions, with the first having been modified to incorporate the dogmatic prior on b_i:

ASSUMPTION 5: *The support of $\mu_{0,\theta}^i$ is the set $\Theta_i^{prior} \subset \Theta$ for which $b_i = b^*$ and $a_i = a_{0i}$.*

ASSUMPTION 6: $b_0 \leq a_0$ *and* $|\alpha_j| + |\beta_j| \leq 1$ *for all j.*

The following proposition (originally Proposition 4) only requires minor modification.[24,25]

PROPOSITION 12: *Suppose that r_t is never observed and agents single-home. Under Assumptions 5 and 6, agent i's identified set, $I_i(R_{0i})$, is nonempty and consists of all $\theta_i \in \Theta$ such that $a_i = a_{0i}$, $b_i = b^*$, $\alpha = \left(\rho_{is} - \beta b^*\right)/a_i$, and $\tilde{\Omega}(\theta_i)$ is positive semidefinite, where $\tilde{\Omega}(\theta_i)$ is the correlation matrix for $\left(\omega_t, r_t, x_{it}, s_t\right)$.*

PROOF:
The proof for Proposition 4 (reproduced with only minor modification below) applies the expressions from Assumption 5 and then shows that the identified set is

---

[24] The results from Section I do not rely on the agent's prior on b_i.

[25] Unlike in the original proposition, the presence of a nonzero b^* here means that the identified set may sometimes exclude extreme values of β. The specific bounds are not necessary for the results that follow, and we do not derive explicit expressions for them.

nonempty by picking (among other parameters) $\beta = 0$. The same argument can be used to prove the new setting, as the degenerate prior on $b_i = b^*$ has no effect in the case with $\beta = 0$.

Recall that $I_i(R_{0i}) = \{\theta_i \in \Theta_i^{prior} : a_i\alpha + b_i\beta = \rho_{is}\}$, where any $\theta_i \in \Theta_i^{prior}$ must give rise to a positive semidefinite correlation matrix $\tilde{\Omega}(\theta_i)$ for $(\omega_t, r_t, x_{it}, s_t)$. Assumption 5 implies that $\theta_i \in I_i(R_{0i})$ if and only if $a_i = a_0$, $b_i = b^*$, $\alpha = (\rho_{is} - \beta b^*)/a_i$, and $\tilde{\Omega}(\theta_i)$ is positive semidefinite. Assumption 6 implies that $|\alpha_{0j}| = |\rho_{ij}/a_{0i}| = |\alpha_{0j} + (b_{0i}/a_{0i})\beta_{0j}| \leq \|\alpha_{0j}| + |\beta_{0j}\| \leq 1$. We can then pick $\beta = 0$, $\gamma = 0$, and assume that s_{jt} are mutually independent and independent of r_t and x_{it} conditional on ω_t, so $\Sigma = \text{corr}(s_t) = \alpha\alpha' + K$, where K is a diagonal matrix with entries equal to $1 - \alpha_j^2$. It follows that $\tilde{\Omega}(\theta_i)$ is positive semidefinite, so $I_i(R_{0i})$ is nonempty. ∎

Under Assumptions 5 and 6, Propositions 1 and 12 imply that the agent's trust in information source j is[26]

$$\bar{\alpha}_j^i = \frac{\rho_{ij} - \bar{\beta}_j^i b^*}{a_{0i}} = \frac{a_{0i}\alpha_{0j} + b_{0i}\beta_{0j} - \bar{\beta}_j^i b^*}{a_{0i}} = \alpha_{0j} + \frac{b_{0i}\beta_{0j} - \bar{\beta}_j^i b^*}{a_{0i}},$$

where $\bar{\beta}_j^i$ denotes agent i's expectation of β_j under $\mu_{\infty,\theta}^i$.

PROPOSITION 13: *Suppose that r_t is never observed, agents single-home, and Assumptions 5 and 6 hold. Then, given any prior $\mu_{0,\theta}^i$, there exists $b_{max}^* > 0$ such that for all $b^* \in [-b_{max}^*, b_{max}^*]$, agent R's (L's) trust in source j is increasing (decreasing) in the source's bias β_{0j}, holding constant the source's accuracy α_{0j}. Additionally, in the limit as $b^* \to 0$ and then as $a_0 \downarrow b_0$, she will come to believe that a perfectly right-biased (left-biased) source is perfectly accurate, and trust it more than any unbiased source with $\alpha_{0j} < 1$.*

PROOF:

Without loss of generality, let i be an R-agent. For the first result, note that the derivative of her trust $\bar{\alpha}_j^i$ with respect to in a source j's bias β_{0j}, holding constant its accuracy α_{0j}, is

$$\frac{\partial}{\partial \beta_{0j}} \bar{\alpha}_j^i \propto b_{0i} - \left(\frac{\partial}{\partial \beta_{0j}} \bar{\beta}_j^i\right) b^*.$$

For any continuous prior $\mu_{0,\theta}^i$, $\partial\bar{\beta}_j^i/\partial\beta_{0j}$ is bounded above by some finite $\bar{M}_j$ and below by some finite $\underline{M}_j$ for the closed interval $\beta_{0j} \in [-1, 1]$. First, consider any $b^* > 0$. If $\bar{M}_j \leq 0$, then the second term above is always strictly negative. If $\bar{M}_j > 0$, then for any $b^* < b_{0i}/\bar{M}_j$, we have $b_{0i} > \bar{M}_j b^* \geq \left(\frac{\partial}{\partial\beta_{0j}}\bar{\beta}_j^i\right)b^*$. In both cases, we have the desired result of $\partial\bar{\alpha}_j^i/\partial\beta_{0j} > 0$ for all β_{0j}. A similar

---

[26]For any $\theta_i \in I_i(R_{0i})$, by Proposition 12, $\alpha_j^i = (\rho_{ij} - \beta_j^i b^*)/a_{0i}$. The expression for trust $\bar{\alpha}_j^i$ is then obtained by taking the expectation over the limiting posterior $\mu_{\infty,\theta}^i$.

argument with $\underline{M}_j$ proves the result for $b^* < 0$. Hence, we can set $b^*_{max} = \min\left\{\min_j \left|b_{0i}/\bar{M}_j\right|, \min_j \left|b_{0i}/\underline{M}_j\right|\right\}$ to guarantee the result for all $b^* \in \left[-b^*_{max}, b^*_{max}\right]$ and across all sources j.

For the second result, the R-agent's trust with respect to a perfectly right-biased source is $\bar{\alpha}^i_j = \left(b_{0i} - \bar{\beta}^i_j b^*\right)/a_{0i}$. Since $\bar{\beta}^i_j$ is bounded within $[-1,1]$ for any prior, taking the limit as $b^* \to 0$ and then $a_0 \downarrow b_0$ gives $\bar{\alpha}^i_j = 1$. For any unbiased source with $\alpha_{0j} < 1$, the agent's trust is $\bar{\alpha}^i_j = \alpha_{0j} - \bar{\beta}^i_j b^*/a_{0i}$, which approaches $\alpha_{0j} < 1$ in the limit as $b^* \to 0$. ∎

COROLLARY 7: *Suppose that r_t is never observed, agents single-home, and Assumptions 5 and 6 hold. Further suppose that all sources have accuracy $\alpha_{0j} < 1$ and there is at least one perfectly right-biased source and at least one perfectly left-biased source. In the limit as $b^* \to 0$ and then $a_0 \downarrow b_0$, expected disagreement is one.*

PROOF:
Follows from Proposition 13. ∎

APPENDIX D. TRUST AND DISAGREEMENT WITH OVERCONFIDENCE

In this section, we derive general results for trust and beliefs in the presence of overconfidence, as discussed in Section IV. Here we drop the simplifying assumption that $a^{max}_i = \sqrt{a_0^2 + b_{0i}^2}$ and instead assume that $a^{max}_i \geq \sqrt{a_0^2 + b_{0i}^2}$, as in Assumption 1′.

LEMMA 6: *Under Assumption 1′, agent i's trust in information source j is*

$$\bar{\alpha}^i_j = A_i \rho_{ij} = A_i\left(a_0 \alpha_{0j} + b_{0i} \beta_{0j}\right),$$

where the amplification factor A_i is given by

(D1) $$A_i = \int_{\underline{a}_i}^{a^{max}_i} \frac{1}{a} d\mu^i_{\infty,a}(a).$$

PROOF:
This follows from combining Propositions 1 and 6. ∎

Lemma 6 shows that, as before, small differences in biases b_0 and β_{0j} translate into large differences in trust, provided the amplification factor A_i is large. The amplification factor A_i is typically large whenever a_0 and b_{0i} are small. To see this, suppose that the posterior marginal probability density function on a_i is nonincreasing. It follows then that $A_i \to \infty$ as $\underline{a}_i \to 0$. In other words, A_i is large if $\underline{a}_i \leq \sqrt{a_0^2 + b_{0i}^2}$ is small and the agent's posterior places sufficient weight on values of a_i close to $\underline{a}_i$.
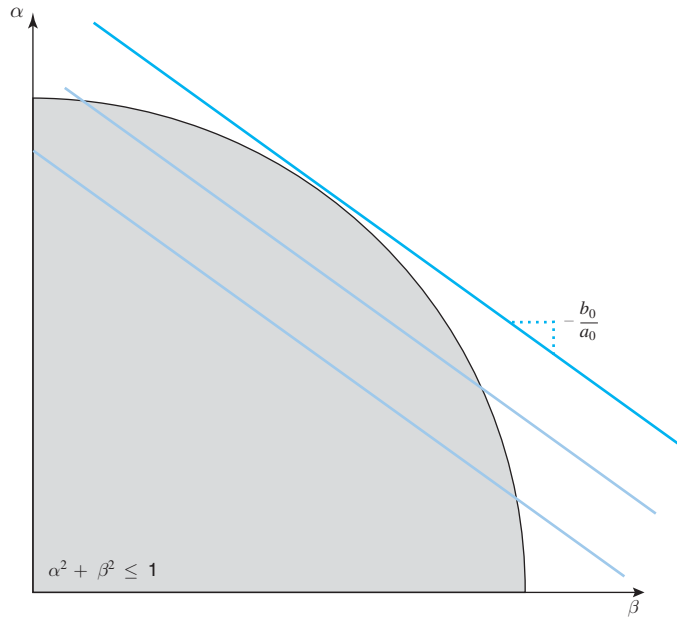
FIGURE D1. ISO-TRUST CURVES

LEMMA 7: *The accuracy and bias* (α_i, β_i) *that maximizes the agent i's trust over all pairs satisfying the feasibility condition* $\alpha_i^2 + \beta_i^2 \leq 1$ *is given by*

$$\left(\alpha_i^{max}, \beta_i^{max}\right) \equiv \left(\frac{a_0}{\sqrt{a_0^2 + b_{0i}^2}}, \frac{b_{0i}}{\sqrt{a_0^2 + b_{0i}^2}}\right).$$

PROOF:

By Lemma 6, we choose (α_i, β_i) to maximize $A_i(a_0\alpha_i + b_{0i}\beta_i)$, subject to $\alpha_i^2 + \beta_i^2 \leq 1$. Standard constrained optimization techniques yield the result. ∎

For intuition, Figure D1 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* that have maximum possible accuracy given their bias. The blue lines in the figure plot the set of *iso-trust curves*: combinations of α and β that yield the same trust. The slope of these lines is $-b_0/a_0$. Sources that fall on higher iso-trust curves are trusted more. From this graphical analysis, it is immediately apparent that the trust-maximizing source $\left(\alpha^{max}, \beta^{max}\right)$ will be the point on the frontier tangent to the iso-trust curves.

LEMMA 8: *In the limit as* $T \rightarrow \infty$, *whenever a single-homing agent i observes source j in any period t* $> \varepsilon T$, *the mean of the agent's posterior on* ω_t *given* s_{jt} *is*

(D2) $$\bar\omega_t^i(j) = \bar\alpha_j^i \tilde{s}_{jt} = A_i \rho_{ij} \tilde{s}_{jt},$$

where $\tilde{s}_{jt} = s_{jt}/\sqrt{\text{var}[s_{jt}]}$ is the standardized version of s_{jt}.

PROOF:

This follows from combining Lemma 6 with Proposition 2. ∎

Appendix E. Trust and Polarization When Agents May Observe Perverse Sources

This Appendix drops the restriction that a single-homing agent in any period $t > \varepsilon T$ chooses among sources that she believes to have *nonnegative* accuracy. We instead assume the following:

ASSUMPTION 7: *In any period $t > \varepsilon T$, a single-homing agent chooses to observe a source among all available sources to minimize her expected mean squared error after observing the source and her own reasoning.*

LEMMA 9: *Under Assumption 7, if all available sources are frontier sources (such that $\alpha_{0j}^2 + \beta_{0j}^2 = 1$) with positive accuracy (i.e., $\alpha_{0j} > 0$) and symmetric biases (i.e., for any source with $\beta_{0j} = \beta$ where $\beta \in (-1, 1)$, there exists a source k with $\beta_{0k} = -\beta$), then biased agents will observe a like-minded source in all periods $t > \varepsilon T$.*

PROOF:

This lemma follows from applying our assumption that $a_{0i} > 0$ to equation (6) and Lemma 6, and then applying the logic of the proof of Proposition 3. ∎

LEMMA 10: *Under Assumption 7, if a trust-maximizing source exists, the agent observes a source with $\bar{\alpha}_j^i = 1$ in all periods $t > \varepsilon T$.*

PROOF:

This lemma follows immediately from our assumption that $a_{0i} > 0$, which implies that if a trust-maximizing source exists, there exists a source such that $\bar{\alpha}_j^i = 1$ and there can be no source such that $\bar{\alpha}_j^i = -1$. ∎

The above two lemmas imply that all main propositions except Proposition 3 continue to hold.

PROPOSITION 14: *Suppose that Assumption 7 holds. Propositions 1, 2, 4, 5, 6, 7, 8, 9, and 10 hold. However, if agent i single-homes, her posterior belief on θ_i is $\mu_{\infty,\theta}^i$, and the expected accuracy a_i of her own reasoning under $\mu_{\infty,\theta}^i$ is less than one, then she chooses to observe in each period a source j for whom $|\bar{\alpha}_j^i|$ is maximal.*

PROOF:

Propositions 1, 2, 4, 5, 6, and 7 do not depend on the choice of sources to observe during the exploitation phase, so their proofs are unaffected by Assumption 7. Proposition 8 can be shown by applying Lemma 9, since sources are assumed to be

symmetric. For Proposition 9, there is only one source, so the agent's choice does not matter. Proposition 10 can be shown by applying Lemma 10, since outlets will still choose the trust-maximizing positions. The last sentence follows from the proof of Proposition 3. ∎

APPENDIX F. ASYMPTOTIC LEARNING WITHOUT REASONING

Can agent i learn about ω_t if the agent does not observe any information source that she ex ante believed to be unbiased? As discussed in Section ID and shown in this section, learning about ω_t from s_t is not possible if reasoning x_{it} were not available. Since ρ_{is} and ρ_{ir} are not observed in this case, the distribution of observable data is given by $R_{0i} = \rho_{rs}$ alone. The identified set $I_i(R_{0i})$, consistent with observed data ρ_{rs}, 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 are perfectly positively correlated, uncorrelated, or perfectly negatively correlated with the true state.

Without reasoning x_{it}, the agent's posterior mean on ω_t may always be zero regardless of what signals are available. This occurs whenever the agent's priors are (α, γ)-symmetric, meaning that $\mu_{0,\theta}^i(\vartheta) = \mu_{0,\theta}^i(\vartheta')$ for all measurable $\vartheta \subseteq \mathcal{L}_\Theta$, where $\vartheta' = \{(a, -\alpha, b, \beta, -\gamma, \Sigma) \mid (a, \alpha, b, \beta, \gamma, \Sigma) \in \vartheta\}$. 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.

PROPOSITION 15: *Suppose that agent i does not observe x_{it} in any period but still observes r_t. Under Assumption 1', agent i's identified set, $I_i(R_{0i})$, includes $\theta_i \in \Theta$ 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 in any period $t > \varepsilon T$ is zero in the limit as $T \to \infty$ if the agent's prior is (α, γ)-symmetric.*

PROOF:
Consider the multi-homing case first. Take any $z \in [-1, 1]$ and $j \in \{1, \ldots, J\}$. Define θ_i as follows. Set $a = a_0$, $b = 0$, $\alpha_j = z$, $\alpha_k = z\Sigma_{jk}$ for all $k \neq j$, $\gamma = z\rho_{rj}$, 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_i corresponds to a well-defined correlation matrix for the joint distribution of (ω_t, r_t, s_t).[27] Therefore, we have that $\theta_i \in I_i(R_{0i})$. The same θ_i works in the single-homing case, which only requires a

---

[27] This correlation matrix is given by

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

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 $\omega_t = z\tilde{s}_{jt} + (1-z)e_t$, where $e_t \sim N(0, 1)$ and is independent of (r_t, s_t). Since it follows that $\omega_t \sim N(0, 1)$, $\bar{\Omega}$ is well defined.

well-defined correlation matrix for the unit-normal joint distribution of (ω_t, r_t, s_{kt}) for each k.

By Proposition 1 and the properties of the multivariate normal distribution, the mean of the multi-homing agent's posterior on ω_t given s_t is

$$\frac{\int_{I_i(R_{0i})} \alpha' \Sigma^{-1} \tilde{s}_t \, d\mu^i_{0,\theta}(\theta_i)}{\mu^i_{0,\theta}\big(I_i(R_{0i})\big)},$$

where $I_i(R_{0i}) = \left\{\theta_i \in \Theta : \tilde{\alpha}' \tilde{\Sigma}^{-1} \tilde{\alpha} \leq 1; b = 0\right\}$. Under single-homing, the analogous mean is

$$\frac{\int_{I_i(R_{0i})} \alpha_j \tilde{s}_{jt} \, d\mu^i_{0,\theta}(\theta_i)}{\mu^i_{0,\theta}\big(I_i(R_{0i})\big)},$$

where $I_i(R_{0i}) = \left\{\theta \in \Theta : \tilde{\alpha}'_j \tilde{\Sigma}_j^{-1} \tilde{\alpha}_j \leq 1 \;\forall j; b = 0\right\}$.[28] Since $\mu^i_{0,\theta}$ is (α, γ)-symmetric, both integrals above are zero. ∎

APPENDIX G. LARGE MARKET OF NONFRONTIER SOURCES

This section considers a situation wherein there are many nonfrontier sources that together provide a relatively "large" quantity of information about both ω_t and r_t, as discussed in Section V. In particular, suppose that the number of sources is large, the signals s_{jt} are mutually independent conditional on ω_t and r_t, and there is at least a minimal amount of diversity in their biases. We formalize this notion of a "large and diverse" set of sources as follows.

DEFINITION 1: *A sequence of random markets is indexed by $J = 1, 2, \ldots, \infty$. Random market J has J sources, indexed by $j = 1, \ldots, J$, each with accuracy and bias $(\alpha_{0j}, \beta_{0j})$ drawn i.i.d. from some distribution F. The sources' signals s_{jt} are mutually independent, conditional on ω_t and r_t. Furthermore, under F,*

(i) *both $\alpha_{0j} \neq 0$ and $\beta_{0j} \neq 0$ have nonzero probability,*

(ii) *α_{0j} and β_{0j} are not perfectly correlated, and*

(iii) *$\alpha_{0j}^2 + \beta_{0j}^2 < 1$ with probability one.*

As shown in Proposition 16 below, a multi-homing agent's posterior mean $\bar{\omega}_t^i$ in a large random market is the same as a single-homing agent when she observes her trust-maximing source. The reason is that a multi-homing agent in a random market can construct a linear combination of the sources' signals whose value will approach the signal of the agent's trust-maximizing source, as in Lemma 2.

---

[28] See the proof of Propositions 6 and 7 for the definitions of $\tilde{\alpha}_j$, $\tilde{\Sigma}_j$, $\tilde{\alpha}$ and $\tilde{\Sigma}$.

PROPOSITION 16: *Suppose that Assumption $1'$ holds. Then the mean of the multi-homing agent's posterior on ω_t given s_t under $\mu^i_{\infty,\theta}$ in the probability limit of a sequence of random markets is*

$$\bar{\omega}^i_t = \alpha^{max}_i \omega_t + \beta^{max}_i \tilde{r}_t.$$

PROOF:

By Lemma 1, it suffices to show that $\lim_{J \to \infty} \rho'_{is} \Sigma^{-1} \tilde{s}_t = a_0 \omega_t + b_{0i} \tilde{r}_t$ and $\lim_{J \to \infty} A_i = 1/\sqrt{a_0^2 + b_{0i}^2}$. We use notation developed in Lemma 4 and let $Q = ZK^{-1}Z'$, $d_{\alpha\alpha} = \frac{1}{J}\sum_{j=1}^J \alpha^2_{0j}/\kappa^2_{0j}$, $d_{\beta\beta} = \frac{1}{J}\sum_{j=1}^J \beta^2_{0j}/\kappa^2_{0j}$, and $d_{\alpha\beta} = \frac{1}{J}\sum_{j=1}^J \alpha_{0j}\beta_{0j}/\kappa^2_{0j}$. It follows that $Q = J\begin{bmatrix} d_{\alpha\alpha} & d_{\alpha\beta} \\ d_{\alpha\beta} & d_{\beta\beta} \end{bmatrix}$. Next, let $W = Q(I+Q)^{-1}$. By Woodbury's matrix identity, we can write $Z(Z'Z + K)^{-1} = (I - W)ZK^{-1}$. It is easy to check that

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

It is also easy to check that $Z(Z'Z + K)^{-1}Z' = (I - W)Q = W$. By Definition 1, $\mathbb{E}[\alpha^2_{0j}/\kappa^2_{0j}]$, $\mathbb{E}[\beta^2_{0j}/\kappa^2_{0j}]$, and $\mathbb{E}[\alpha_{0j}\beta_{0j}/\kappa^2_{0j}]$ exist and are finite. Therefore, $d_{\alpha\alpha}$, $d_{\beta\beta}$, and $d_{\alpha\beta}$ converge in probability in the limit as $J \to \infty$ by the weak law of large numbers. Furthermore, Definition 1 implies that α_{0j} and β_{0j} are not linearly dependent, so neither are $\alpha_{0j}/\sqrt{\kappa^2_{0j}}$ and $\beta_{0j}/\sqrt{\kappa^2_{0j}}$. By Cauchy-Schwarz, we have that $\mathbb{E}[\alpha_{0j}\beta_{0j}/\kappa^2_{0j}]^2 < \mathbb{E}[\alpha^2_{0j}/\kappa^2_{0j}]\mathbb{E}[\beta^2_{0j}/\kappa^2_{0j}]$. This implies that $\text{plim}(d_{\alpha\alpha}d_{\beta\beta} - d^2_{\alpha\beta}) > 0$. Therefore, $Z(Z'Z + K)^{-1}Z' = W \to_p I$.

Further algebraic manipulation shows that

$$y'Z(Z'Z + K)^{-1}\varepsilon_t$$

$$= \frac{(a_0 d_{\beta\beta} + b_0 d_{\alpha\beta})\left(\frac{1}{J}\sum_{j=1}^J \frac{\alpha_{0j}}{\sqrt{\kappa^2_j}}\tilde{\varepsilon}_{jt}\right) + (a_0 d_{\alpha\beta} + b_0 d_{\alpha\alpha})\left(\frac{1}{J}\sum_{j=1}^J \frac{\beta_{0j}}{\sqrt{\kappa^2_j}}\tilde{\varepsilon}_{jt}\right)}{\frac{1}{J^2} + \frac{1}{J}d_{\alpha\alpha} + \frac{1}{J}d_{\beta\beta} + d_{\alpha\alpha}d_{\beta\beta} - d^2_{\alpha\beta}}$$

$$+ o_p(1),$$

where $\tilde{\varepsilon}_{jt} = \varepsilon_{jt}/\sqrt{\kappa^2_j} \sim N(0,1)$ are mutually independent across j as well as independent of α_{0j} and β_{0j}. Therefore, by the weak law of large numbers, $y'Z(Z'Z + K)^{-1}\varepsilon_t \to_p 0$. It immediately follows from Lemma 4 that $\rho'_{is}\Sigma^{-1}\tilde{s}_t = y'Z(Z'Z + K)^{-1}(Z'\varphi_t + \varepsilon_t) \to_p y'\varphi_t = a_0\omega_t + b_{0i}r_t$.

Next, note that the R^2 of the population regression of x_{it} on s_t is $\rho'_{is}\Sigma^{-1}\rho_{is}$. By Lemma 4 and the first paragraph in this proof, $\rho'_{is}\Sigma^{-1}\rho_{is} = y'Z(Z'Z + K)^{-1}Z'y \to_p y'y = a_0^2 + b_{0i}^2$. Therefore, the R^2 of the population regression of x_{it} on s_t and r_t also converges in probability to $a_0^2 + b_{0i}^2$. This implies that $\underline{a}_i \to_p \sqrt{a_0^2 + b_{0i}^2}$. By Assumption 3, $\mu^i_{\infty,a} = \mu^i_{0,a}$, so $A_i \to_p 1/\sqrt{a_0^2 + b_{0i}^2}$. ∎

Appendix H. Mistrust of Motives and Partisan Conflict

This Appendix shows how ideological bias can lead to mistrust of motives across ideological divides and results in intensified conflict in political settings. This exploration is motivated by studies that show that rising numbers of Americans hold negative views toward 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 2020).[29]

We augment our model by adding an observable policy decision d_t to be made by one of two agents, R or 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 $b = 0$ when in fact $b \neq 0$. This assumption is deliberately stark to illustrate that if people underestimate the extent to which others' reasoning is biased, they may attribute their behavior to biased motives instead. More precisely, we show that agents mistakenly learn that $B \neq 0$ even when in fact $B = 0$.

The setup is as follows. Suppose that Assumption $1'$ holds and all agents are single-homers observing their trust-maximizing source $(\alpha_i^{max}, \beta_i^{max})$ in some period $t > \varepsilon T$ in the limit as $T \to \infty$. 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 $-(\omega_t - d_t)^2$.

We assume that agents fail to appreciate both their own and others' ideological biases and believe that $b_i = 0$ for all agents i. 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. Under these assumptions, it is immediate that people who underestimate the extent to which others' *reasoning* is biased attribute observed behavior to biased *motives* instead.

If L observes R's decision d_t in any period when $r_t \neq 0$, then L infers that $B_R = 2b_0/\sqrt{a_0^2 + b_0^2} > 0$. Similarly, if R were to observe L's decisions, R would also conclude that L had an ulterior motive, $B_L = -2b_0/\sqrt{a_0^2 + b_0^2} < 0$. 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 the ideological bias, b_0, of the agents.

---

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

It follows that the political behavior of well-meaning agents with ideological bias can mimic 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, and inefficient expenditures in the battle for power.

Appendix I. Finite-Time Evolution of Beliefs

This Appendix studies the finite-time evolution of beliefs and speed of convergence by presenting numerical simulations. Recall that each period's state variable ω_t represents a different topic in the news—the efficacy of surgical masks in preventing the spread of COVID-19 in one period, the efficacy of vitamin D in reducing the severity of COVID-19 infections in another period, and so on. Here we characterize how quickly agents form beliefs about whether information sources provide accurate information about the states.

For simplicity, we focus on two agents, $i \in \{R, U\}$. Both agents have reasoning x_{it} with the accuracy $a_{0i} = a$, where $a > 0$, but different biases, $b_{0R} = b$ and $b_{0U} = 0$, where $b > 0$. We assume that both agents a priori believe that $b_i = 0$ and $a_i = a$. There are two sources $j \in \{R, U\}$, where $(\alpha_R, \beta_R) = (1/2, 1/2)$ and $(\alpha_U, \beta_U) = (1, 0)$. The sources and the agents' reasoning have unit-normal distributions.

Each agent i observes both x_{it} and s_{jt} for the same source j in each period t. Let the covariance matrix of the observed signals (x_{it}, s_{jt}) be denoted by $\Sigma_{ij} = \begin{bmatrix} \sigma_i^2 & \rho_{ij} \\ \rho_{ij} & \sigma_j^2 \end{bmatrix}$. Each agent has a prior $f(\Sigma_{ij})$, where

$$\Sigma_{ij} = \text{diag}(V_{ij}) \times \Omega_{ij} \times \text{diag}(V_{ij}),$$

Ω_{ij} is a correlation matrix, and V_{ij} is a vector of coefficient scales. We use a Lewandowski-Kurowicka-Joe prior with parameter $\eta = 1$ for the correlation matrix. Setting $\eta = 1$ implies that the density is uniform over all correlation
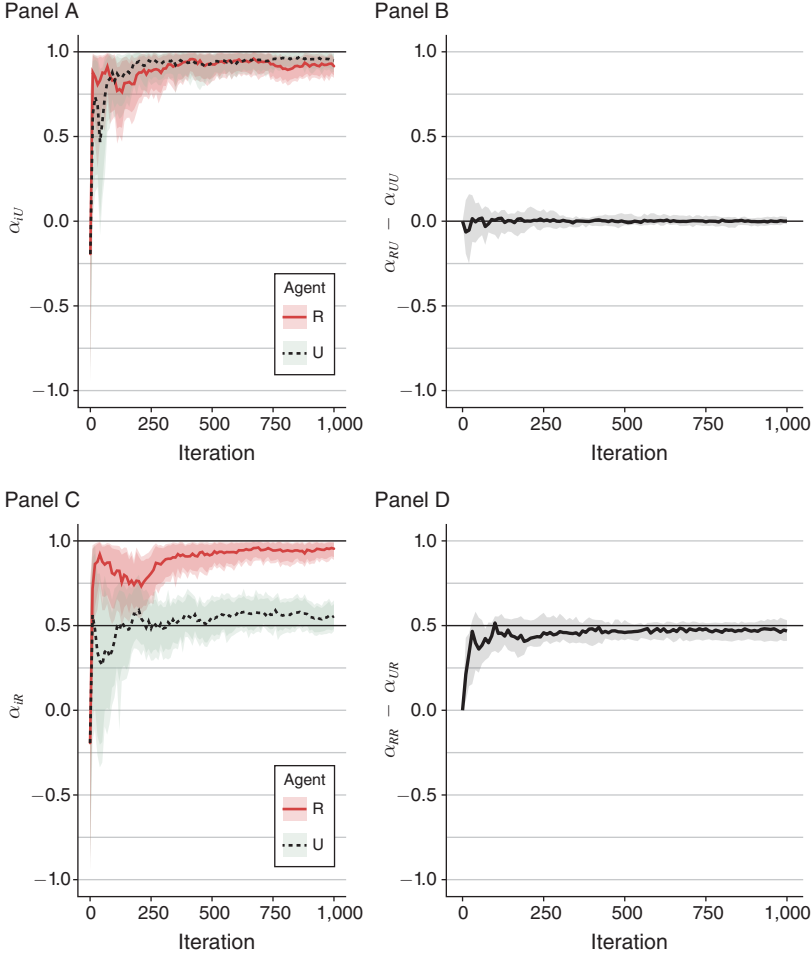
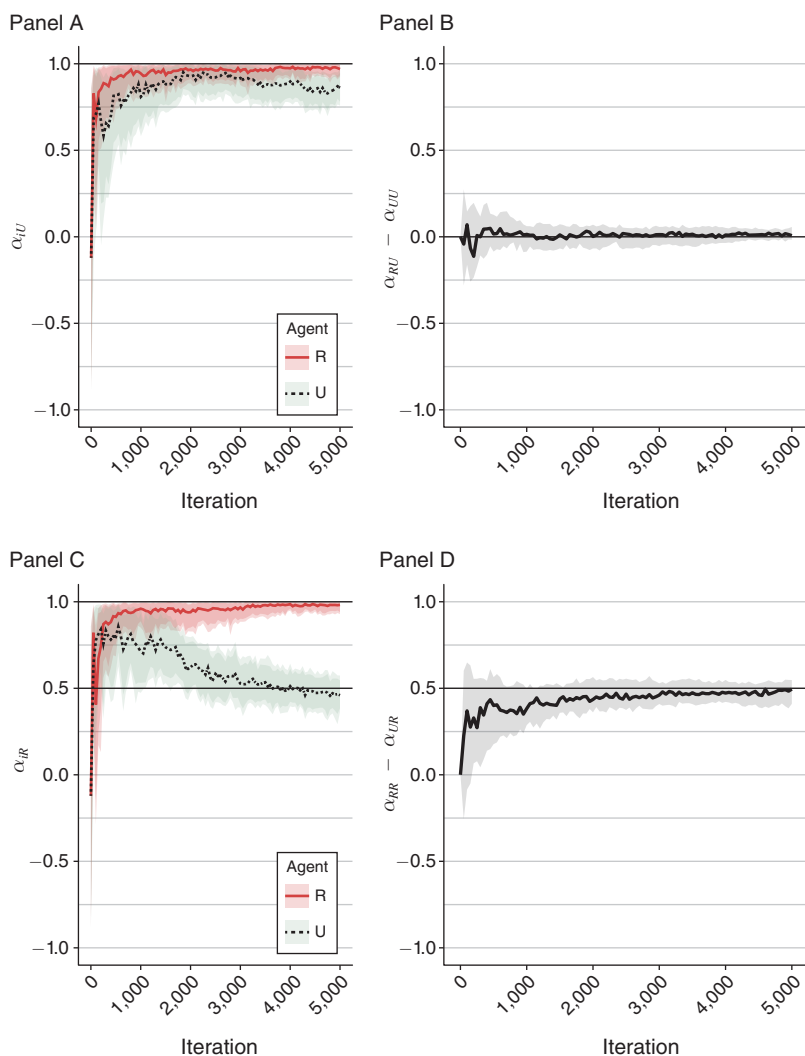FIGURE I1. FINITE-TIME BELIEF EVOLUTION WHEN $a = b = 1/2$

Notes: Panel A shows the simulated posterior distribution of α_{ij} for agents $i = R, U$ observing an unbiased source $j = U$ for a single sequence of signal realizations. We assume that $a = b = 1/2$. The thick lines indicate the posterior median. The lighter bands are for the $[0.05, 0.95]$ percentiles, and the darker bands are for $[0.10, 0.90]$ percentiles. The horizontal line denotes the asymptotic values. In panel B, the thick black line shows the median of the difference in posterior means between the two agents for 25 sequences of simulated signal realizations with different randomization seeds. The bands are for $[0.20, 0.80]$ percentiles. Panels C and D show the simulated posteriors when agents observe a biased source $j = R$ instead of an unbiased source.

matrices. We use a log normal prior with parameters $\mu = 0$ and $\nu = 0.0001$ for the scale. Furthermore, the prior is truncated for all values Σ_{ij} such that

$$\alpha_{ij} = \frac{1}{a} \frac{\rho_{ij}}{\sqrt{\sigma_i^2 \sigma_j^2}} \notin [-1, 1].$$

We use the software platform STAN to simulate 1,000 draws from the agent i's posterior distribution in each period t.

Figure I1 shows belief evolution when $a = b = 1/2$. On the left, we display the posterior distribution for α_{ij} for a single sequence of signal realizations. On the

Panel A

Panel B

Panel C

Panel D

FIGURE I2. FINITE-TIME BELIEF EVOLUTION WHEN $a = b = 1/5$

Notes: Panel A shows the simulated posterior distribution of α_{ij} for agents $i = R, U$ observing an unbiased source $j = U$ for a single sequence of signal realizations. We assume that $a = b = 1/5$. The thick lines indicate the posterior median. The lighter bands are for the $[0.05, 0.95]$ percentiles, and the darker bands are for $[0.10, 0.90]$ percentiles. The horizontal line denotes the asymptotic values. In panel B, the thick black line shows the median of the difference in posterior means between the two agents for 25 sequences of simulated signal realizations with different randomization seeds. The bands are for $[0.20, 0.80]$ percentiles. Panels C and D show the simulated posteriors when agents observe a biased source $j = R$ instead of an unbiased source.

right, we display the median difference in posterior mean between agents R and U across multiple sequences of signal realizations. The top two panels show that the trust of agents R and U in an unbiased source quickly converges toward the asymptotic value of one, while the difference in posterior mean converges to zero. The bottom two panels show that the trust of agents R and U in a R-biased source quickly diverges.

Figure I2 shows the case where $a = b = 1/5$. In this case, the agent's reasoning is noisier, so convergence is slower. However, in all of the above cases, the beliefs converge toward the asymptotic beliefs that we analytically derive.

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 *Judgement under Uncertainty: Heuristics and Biases*, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky, 294–305. 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 V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107 (3): 797–817.

Barberá, Pablo. 2020. "Social Media, Echo Chambers, and Political Polarization." In *Social Media and Democracy: The State of the Field, Prospects for Reform*, edited by Nathaniel Persily and Joshua A. Tucker, 34–55. Cambridge University Press.

Benoît, Jean-Pierre, and Juan Dubra. 2019. "Apparent Bias: What Does Attitude Polarization Show?" *International Economic Review* 60 (4): 1675–703.

Berk, Robert H. 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. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24 (2): 265–72.

Blackwell, David, and Lester Dubins. 1962. "Merging of Opinions with Increasing Information." *Annals of Mathematical Statistics* 33 (3): 882–86.

Bohren, J. Aislinn, and Daniel N. Hauser. 2021. "Learning with Heterogeneous Misspecified Models: Characterization and Robustness." *Econometrica* 89 (6): 3025–77.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *Quarterly Journal of Economics* 127 (3): 1243–85.

Bowen, T. Renee, Danil Dmitriev, and Simone Galperti. 2023. "Learning from Shared News: When Abundant Information Leads to Belief Polarization." *Quarterly Journal of Economics* 138 (2): 955–1000.

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." *PNAS* 114 (40): 10612–17.

Boyd, Stephen P., and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.

Campante, Filipe R., and Daniel A. Hojman. 2013. "Media and Polarization: Evidence from the Introduction of Broadcast TV in the United States." *Journal of Public Economics* 100: 79–92.

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

Cheng, Ing-Haw, and Alice Hsiaw. 2022a. "Bayesian Doublespeak." Preprint, SSRN. https://dx.doi.org/10.2139/ssrn.4082225.

Cheng, Ing-Haw, and Alice Hsiaw. 2022b. "Distrust in Experts and the Origins of Disagreement." *Journal of Economic Theory* 200: 105401.

Cotton, John L. 1985. "Cognitive Dissonance in Selective Exposure." In *Selective Exposure to Communication*, edited by Dolf Zillmann and Jennings Bryant, 11–33. Routledge.

de Benedictis-Kessner, Justin, 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): 902–16.

DeGroot, Morris H. 1974. "Reaching a Consensus." *Journal of the American Statistical Association* 69 (345): 118–21.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics* 118 (3): 909–68.

Dixit, Avinash K., and Jörgen W. Weibull. 2007. "Political Polarization." *PNAS* 104 (18): 7351–56.

Eagly, Alice H., Serena Chen, Shelly Chaiken, and Kelly Shaw-Barnes. 1999. "The Impact of Attitudes on Memory: An Affair to Remember." *Psychological Bulletin* 125 (1): 64–89.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*. Row, Peterson, and Company.

Flaxman, Seth, Sharad Goel, and Justin M. Rao. 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly* 80 (S1): 298–320.

Frick, Mira, Ryota Iijima, and Yuhta Ishii. 2020. "Misinterpreting Others and the Fragility of Social Learning." *Econometrica* 88 (6): 2281–2328.

Fryer, Roland G., 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): 1470–501.

Gallup and Knight Foundation. 2018. *American Views: Trust, Media and Democracy*. Gallup and Knight Foundation.

Gallup and Knight Foundation. 2020. *American Views 2020: Trust, Media and Democracy*. Gallup and Knight Foundation.

Gentzkow, Matthew. 2016. "Polarization in 2016." Toulouse Network for Information Technology Whitepaper.

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): 1799–839.

Gentzkow, Matthew, Jesse M. Shapiro, and Daniel F. Stone. 2015. "Media Bias in the Marketplace: Theory." In *Handbook of Media Economics*, Vol. 1, edited by Simon P. Anderson, Joel Waldfogel, and David Strömberg, 623–45. Elsevier.

Gentzkow, Matthew, Michael B. Wong, and Allen T. Zhang. 2025. *Data and Code for: "Ideological Bias and Trust in Information Sources."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E203381V1.

Gibbons, Robert, Niko Matouschek, and John Roberts. 2013. "Decisions in Organizations." In *Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 373–431. Princeton University Press.

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

Glaeser, Edward L., and Bryce A. Ward. 2006. "Myths and Realities of American Political Geography." *Journal of Economic Perspectives* 20 (2): 119–44.

Golub, Benjamin, and Matthew O. Jackson. 2010. "Naive Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics* 2 (1): 112–49.

Golub, Benjamin, and Matthew O. Jackson. 2012. "How Homophily Affects the Speed of Learning and Best-Response Dynamics." *Quarterly Journal of Economics* 127 (3): 1287–338.

Heidhues, Paul, Botond Kőszegi, and Philipp Strack. 2018. "Unrealistic Expectations and Misguided Learning." *Econometrica* 86 (4): 1159–214.

Heidhues, Paul, Botond Kőszegi, and Philipp Strack. 2019. "Overconfidence and Prejudice." PrePrint, arXiv. https://doi.org/10.48550/arXiv.1909.08497.

Hetherington, Marc J., and Thomas J. Rudolph. 2020. *Why Washington Won't Work: Polarization, Political Trust, and the Governing Crisis*. 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): 405–31.

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: 129–46.

Jerit, Jennifer, and Yangzi Zhao. 2020. "Political Misinformation." *Annual Review of Political Science* 23: 77–94.

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

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

Kondor, Péter. 2012. "The More We Know about the Fundamental, the Less We Agree on the Price." *Review of Economic Studies* 79 (3): 1175–207.

Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108 (3): 480–98.

Levendusky, Matthew S. 2013. "Why Do Partisan Media Polarize Viewers?" *American Journal of Political Science* 57 (3): 611–23.

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

Lian, Heng. 2009. "On Rates of Convergence for Posterior Distributions under Misspecification." *Communications in Statistics - Theory and Methods* 38 (11): 1893–900.

Liang, Annie, and Xiaosheng Mu. 2020. "Complementary Information and Learning Traps." *Quarterly Journal of Economics* 135 (1): 389–448.

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): 2098–109.

Marietta, Morgan, and David C. Barker. 2019. *One Nation, Two Realities: Dueling Facts in American Democracy*. Oxford University Press.

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

McIntyre, Lee. 2018. *Post-truth*. MIT Press.

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95 (4): 1031–53.

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." *Annals of the American Academy of Political and Social Science* 658 (1): 36–66.

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

Pew Research Center. 2014a. *Political Polarization and Media Habits*. Pew Research Center.

Pew Research Center. 2014b. *Political Polarization in the American Public*. Pew Research Center

Pew Research Center. 2020. *US Media Polarization and the 2020 Election: A Nation Divided*. Pew Research Center.

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

Pronin, Emily. 2007. "Perception and Misperception of Bias in Human Judgment." *Trends in Cognitive Sciences* 11 (1): 37–43.

Pronin, Emily, Daniel Y. Lin, and Lee Ross. 2002. "The Bias Blind Spot: Perceptions of Bias in Self versus Others." *Personality and Social Psychology Bulletin* 28 (3): 369–81.

Rabin, Matthew, and Joel L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *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. 2016. "Communication with Unknown Perspectives." *Econometrica* 84 (6): 2029–69.

Slivkins, Aleksandrs. 2019. "Introduction to Multi-armed Bandits." *Foundations and Trends in Machine Learning* 12 (1–2): 1–286.

Thaler, Michael. 2024. "The Fake News Effect: Experimentally Identifying Motivated Reasoning Using Trust in News." *American Economic Journal: Microeconomics* 16 (2): 1–38.

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. Texas A&M University Press.

Tversky, Amos, and Daniel Kahneman. 1973. "Availability: A Heuristic for Judging Frequency and Probability." *Cognitive Psychology* 5 (2): 207–32.

van der Vaart, A. W. 2000. *Asymptotic Statistics*. Cambridge University Press.