Ideological Bias and Trust in Information Sources

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

October 2018

Abstract

Two key features describe ideological differences in society: (i) individuals persistently disagree about objective facts; (ii) individuals also disagree about which sources can be trusted to provide reliable information about those facts. We develop a model in which these patterns arise endogenously as the result of Bayesian learning with small biases. Individuals learn the accuracy of information sources by comparing the sources' reports about a sequence of states to noisy feedback the individuals observe about the truth. Arbitrarily small biases in this feedback can lead them to trust biased sources more than unbiased sources, create persistent divergence in beliefs, and cause them to become overconfident in their own judgment. Increasing the number of available sources can deepen rather than mitigate ideological divergence. All of these patterns can be similar whether agents see only ideologically aligned sources or see a diverse range of sources.

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

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

1 Introduction

Ideological divisions in society often seem intractable, with those on either side persistently disagreeing about objective facts. In recent years, for example, fervent debates over the validity of global warming, evolution, and vaccination have persisted long after the establishment of a scientific consensus. Partisans also disagree about which sources can be trusted to provide reliable information about such facts. In the United States, for instance, 88 percent of consistent conservatives say they trust news and information from Fox News, while 81 percent of consistent liberals say they distrust it; 56 percent of consistent liberals say they trust CNN, while 61 percent of consistent conservatives say they distrust it (Pew 2014). Such divisions have persisted or deepened even as new media technologies have made information more widely and cheaply available than ever before. The information age has, paradoxically, produced what has been dubbed a "post-truth" era (Keyes 2004).

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

In this paper, we explore a different possibility, which is that rational Bayesian inference may magnify the influence of even small biases when agents are uncertain which sources they can trust. Building on recent insights by Acemoglu et al. (2016) and Sethi and Yildiz (2014), among others, we show that arbitrarily small biases in information processing may lead to substantial and persistent divergence in both trust in information sources and beliefs about facts, with partisans on each side trusting unreliable ideologically aligned sources more than accurate neutral sources, and also becoming overconfident in their own judgment. We show that increasing the number of available information sources in such a setting may deepen rather than mitigate ideological differences. Consistent with recent evidence suggesting that the magnitude of selective exposure has generally been limited (Gentzkow and Shapiro 2011; Flaxman et al. 2016), we show that these patterns arise whether agents selectively view only ideologically aligned sources or are exposed to a diverse range of sources.

An agent in our model wishes to learn about a sequence of unobserved states $\omega_t \sim N(0,1)$, which are drawn independently in each period t. Each period she observes a normally distributed signal s_{jt} from one or more information sources j which are correlated with ω_t. This correlation α_j, which we refer to as the *accuracy* of source j, is *ex ante* uncertain. The agent never learns the true value of ω_t, but she observes some feedback x_t by direct observation that is correlated with ω_t. The key features of x_t are that it is noisy (in the sense that its correlation with ω_t is low) and that the agent believes it to be independent (in the sense that errors in x_t are independent of errors in the s_{jt}). Over time, therefore, comparing this feedback to the reports s_{jt} allows the agent to learn which information sources she can trust.

We think of the states ω_t as individual items discussed in the news: facts about global warming, the fitness of a presidential candidate, risks from vaccines, and so on. The information sources could be traditional or digital media outlets, as well as other sources such as word of mouth from particular friends. The feedback x_t captures whatever individuals can learn themselves through direct observation, such as information about global warming from observing local weather events, or information about vaccine risk from observing which children in one's social network have developed autism. The key properties we impose on x_t are natural in such examples, as direct observation of weather or children's outcomes is a noisy signal of the true state, but the agent may pay attention to it nevertheless because they believe it to be independent of the distortions that may plague other information sources.

Bias in the model takes the form of systematic ideological errors in the agent's feedback x_t. We introduce an unobserved *ideological state* r_t each period, where r_t is the belief about ω_t that would be held by a conservative ideologue, and $-r_t$ is the belief that would be held by a liberal ideologue. We allow the agent's feedback to be distorted either toward or away from r_t, and refer to the correlation of x_t and r_t as the agent's bias. When the reports of sources are also biased, this causes the assumed independence of x_t and the s_{jt} to be violated, and so implies that the true data generating process lies outside the support of the agent's priors. Our primary interest is in the case where the magnitude of this bias is small, but on a similar scale to the accuracy of x_t. Such bias would not significantly impact beliefs in a model where agents know the accuracies of the sources, but its effects are amplified when agents must learn which sources to trust.

Our main objects of interest are the agent's beliefs about the accuracies α_j of the sources and her beliefs about the states ω_t, both in the limit as $t \to \infty$. Assuming that the agent has full support priors on the parameters θ characterizing the model, we show that the limiting distribution of her posterior beliefs on θ has support in an identified set, and is distributed within this set according to the conditional distribution of the prior. The fact that this set is always non-empty shows that even when the agent is biased the data never invalidate her model of the world. To obtain a sharper characterization, we then focus on a special case in which the agent's priors on the accuracy of x_t are concentrated at the true value, and show that in this case her beliefs about the accuracies of information sources eventually place probability one on a single vector α^*, which we define as the agent's asymptotic *trust* in the respective sources. Under this limiting posterior on θ, the agent's belief about ω_t after observing signals s_t has a mean $\overline{\omega}_t$, and we refer to the correlation between $\overline{\omega}_t$ and r_t as the asymptotic *polarization* of the agent's beliefs.

We analyze two scenarios, one in which the observes exactly one source j in each period (i.e., she "single homes" in the language of Rochet and Tirole 2003), and another in which she observes all j in each period (i.e., she "multi homes"). We begin with the benchmark case in which the agent has no ideological bias. In this case, her trust α^* is equal to the true accuracies α_0 of the information sources. This is true under either single- or multi-homing. The agent is also not overconfident, in the sense that she has correct beliefs about the accuracy of her own feedback. Consequently, the agent's asymptotic beliefs are the same as if she knew the true data generating process. If the accuracy of the information sources is sufficiently high, there is no polarization of beliefs and agents with different biases will agree on the value of ω_t in each period.

The benchmark case allows us to isolate an amplification mechanism that will be the key reason small biases have outsized effects in this model. An agent's trust in source j will be higher the greater is the correlation between the source's signal s_{jt} and the agent's direct observations x_t. When the accuracy of x_t is low, this correlation will be small even if s_{jt} is perfectly informative. Thus, small differences in observed correlation (as will be produced by small ideological biases in both x_t and s_{jt}) are rationally interpreted to imply large differences in accuracy.

Our main results show that introducing arbitrarily small biases in the agent's feedback changes the results of the benchmark case dramatically. Single-homing agents come to trust biased sources more than unbiased sources, and partisans' beliefs about ω_t and trust in outlets become polarized. Biased agents believe like-minded sources are highly accurate, unbiased sources are less-accurate, and opposite-minded sources are perverse, in the sense that their reports are negatively correlated with the state. The entry of partisan media thus reduces trust in unbiased sources. Moreover, all biased agents become overconfident, in the sense that they believe the accuracy of their feedback is greater than it is. If the bias of the agent's feedback is on a similar scale to its accuracy, the polarization of the agent's beliefs will be large in magnitude whenever the sources also have significant bias. Thus, when two agents with opposite biases observe their respective trust-maximizing source, they have large disagreements about ω_t.

One might expect these results to be quite different under multi-homing, as an agent in that case can compare results from diverse sources and so moderate her beliefs. While this force is present in our model, we show that multi-homing does not in general reduce polarization and divergent trust, and may in fact make them worse. We focus on two main cases, one in which the market contains the agent's trust maximizing source, and the other in which it contains a sufficiently large set of sources that she can combine their signals to approximate her trust maximizing source. We show in both cases that our main results from the single-homing case continue to hold, and that polarization may actually be greater under multi-homing than under single-homing. The driving force is that the agent endogenously comes to believe that unbiased (or less biased) sources in the market are essentially noisier versions of her preferred biased source, and so ignores their reports in updating her beliefs.

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

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

Our work relates to a growing literature on models of opinion polarization. Dixit and Weibull (2007) show how the beliefs of Bayesians with different priors can polarize when new information arrives. Benoit and Dubra (2015) argue that findings of group attitude polarization in psychological studies (e.g., Lord, Ross, and Lepper 1979) can be rationalized using purely Bayesian models. Fryer, Harms, and Jackson (2015) show that opinion polarization can persist when Bayesian agents have limited memory. Ortoleva and Snowberg (2015) explore how overconfidence drives polarization and affects political behavior.

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

The paper proceeds as follows. Section 2 describes the model. Section 3 characterizes the asymptotic distribution of beliefs. Section 4 presents our main results for the single-homing case. Section 5 presents our main results for the multi-homing case. Section 6 concludes.

2 Model

2.1 Setup

An agent wishes to learn about a sequence of states $\omega_t \sim N(0,1)$, which are independently distributed over time periods $t = 1, 2, \ldots \infty$. The agent learns about ω_t by observing signals $s_{jt} \sim N(0,1)$ from some subset of the available information sources $j = 1, ..., J$. In each period, with some independent probability strictly between zero and one, she also observes direct feedback $x_t \sim N(0,1)$ about ω_t based on her personal experience. Our main focus will be the limit of the agent's beliefs as $t \to \infty$.

Signals and feedback are given by

$$\begin{aligned} s_{jt} &= \alpha_j \omega_t + \varepsilon_{jt} \\ x_t &= a\omega_t + \eta_t, \end{aligned} \tag{1}$$

where $\alpha_j \in [-1,1]$ and $a \in (0,1]$. We let s_t, ε_t, and α denote the J-vectors of s_{jt}, ε_{jt}, and α_j respectively. We assume that the vector of errors (ε_t, η_t) is jointly normal with marginal distributions $\varepsilon_{jt} \sim N\left(0, 1-\alpha_j^2\right)$ and $\eta_t \sim N\left(0, 1-a^2\right)$, and is drawn i.i.d. across periods. We denote the correlation matrix of s_t by $\Sigma = \alpha\alpha' + E\left(\varepsilon_t \varepsilon_t'\right)$ and the vector of correlations between the s_t and x_t by ρ, where the j-th element $\rho_j = \alpha_j a + E\left(\varepsilon_{jt}\eta_t\right)$.[1] Note that α_j is the correlation of s_{jt} with ω_t, and a is the correlation of x_t with ω_t. The joint distribution of signals s_t, feedback x_t, and states ω_t is fully parameterized by $\theta := (a, \alpha, \rho, \Sigma)$.[2]

We refer to α_j as the *accuracy* of information source j, and to a as the accuracy of the agent's feedback. We will refer to the agent's limiting belief about α_j as her *trust* in source j and her

[1] We assume that $\left|\rho_{0j}\right| < 1$ for all j and that $\left|\Sigma_{0jj'}\right| < 1$ for any $j' \neq j$, so none of the observed signals or feedback are perfectly correlated.

[2] As will become clear below, it is convenient to parameterize the covariances of the error terms by Σ and ρ rather than directly by $E\left(\varepsilon_t \varepsilon_t'\right)$ and $E\left(\varepsilon_{jt}\eta_t\right)$. The two formulations are equivalent, however, since given a and α there is a one-to-one correspondence between (Σ, ρ) and $\left(E\left(\varepsilon_t \varepsilon_t'\right), E\left(\varepsilon_{jt}\eta_t\right)\right)$.

limiting belief about a as her *confidence*. Both of these limits are defined precisely below. As discussed above, we think of feedback in most settings as noisy, so we focus on the case where a is close to zero. Note that we allow for the possibility that sources may be highly accurate ($\alpha_j \approx 1$), inaccurate (α_j positive but close to zero), or even perverse, in the sense that they tend to report the opposite of the truth ($\alpha_j < 0$).

We will focus on two scenarios. Under *single-homing*, the agent observes the signal s_{jt} of exactly one source j in each period, and each source is observed infinitely many times as $t \to \infty$. Under *multi-homing*, the agent observes the signals of all J sources in each period.

While we describe the model from the perspective of a representative agent, the model can also be interpreted as describing a population of heterogeneous agents each of which solves the problem we describe.

2.2 Prior Beliefs

Agents in the model are initially uncertain not only about the states ω_t, but also about the parameters θ that govern the distribution of feedback and signals. We let Θ denote the set of all such parameter values.[3] We let θ_0 denote the true value of θ, and where it adds clarity we will use a_0, α_{j0}, Σ_0, and ρ_0 to refer to the true values of the individual components.

At the beginning of the first period, the agent has a prior belief on Θ with a continuous density which we denote μ_0. The agent correctly believes that each ω_t is distributed i.i.d. $N(0,1)$.

The central assumption of our model is that the agent believes the direct feedback x_t she observes is independent evidence about ω_t, in the sense that the errors in x_t are uncorrelated with the errors in the s_{jt}. For θ in the set $\Theta^{prior} \subset \Theta$ consistent with this restriction, $\rho_j = a\alpha_j$. We assume that the agent's prior has full support on this set.

Assumption 1. *The support of μ_0 is the set $\Theta^{prior} \subset \Theta$ of all values of θ such that $E\left(\varepsilon_{jt}\eta_t\right) = 0$.*

We will be mainly interested in the case where the agent's beliefs about the accuracy of her own signal are concentrated at the true value a_0. However, we do not assume that the agent places

[3] Θ is the set of all $\theta \in \mathbb{R}^{J^2+2J+1}$ that satisfy the following conditions: (1) $a \in (0,1]$, (2) $\alpha \in [-1,1]^J$, (3) $\begin{bmatrix} \Sigma & \rho \\ \rho' & 1 \end{bmatrix}$ is nonsingular and is the sum of $\begin{pmatrix} \alpha \\ a \end{pmatrix}\begin{pmatrix} \alpha \\ a \end{pmatrix}'$ and a positive-semidefinite matrix with diagonal elements equal to $\left(1-\alpha_1^2, ..., 1-\alpha_J^2, 1-a^2\right)$.

prior probability one on this value because, as will become clear below, there will be cases where data the agent observes cannot be rationalized by an element of Θ^{prior} with $a = a_0$. In these cases, we want to allow the agent to adjust her belief about a to rationalize the data. We will thus focus on the limit of a sequence of priors that become increasingly concentrated at values close to a_0 while maintaining the property of full support on Θ^{prior}.

Definition 1. Let $\{\mu_{0,n}\}_{n=1}^{\infty}$ be a sequence of prior beliefs with full support on Θ^{prior}. We say that such a sequence *becomes concentrated at* a_0 if for any $\theta, \theta' \in \Theta^{prior}$ such that the associated a and a' satisfy $|a' - a_0| > |a - a_0|$, we have

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

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

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

2.3 Ideological Bias

The final component of the model allows for the possibility that the agent may be subject to ideological bias that leads her feedback x_t to be systematically distorted in either a conservative or a liberal direction. Underlying this could be a range of well-studied psychological phenomena including the availability heuristic (Tversky and Kahneman 1973), salience (Bordalo et al. 2012), confirmation bias (Lord Ross and Lepper 1979), and motivated reasoning (Kunda 1990). The result will be that the true θ_0 may fall outside Θ^{prior}.

To model such bias, we introduce an unobserved ideological state $r_t \sim N(0,1)$ that is independent of ω_t and drawn independently each period. We think of r_t as representing the conservative position on issue ω_t (that is, the belief that a conservative ideologue would hold) and $-r_t$ as representing the liberal position.

We define the agent's *bias* $b \in [-1,1]$ to be the correlation of her feedback x_t with r_t. That is, we assume the error term that enters x_t in equation (1) is given by:

$$\eta_t = br_t + \tilde{\eta}_t,$$

where $\tilde{\eta}_t$ is independent of both ω_t and r_t as well as all other errors in the model, and its marginal distribution is $N\left(0, 1-a^2-b^2\right)$. We restrict attention to biases b satisfying $b^2 \leq 1-a^2$. Positive values of b represent conservative bias, while negative values represent liberal bias.

We also allow for the possibility that information sources may be biased. We define the *slant* $\beta_j \in [-1,1]$ of outlet j to be the correlation of s_{jt} with r_t. That is, we assume the error term that enters s_{jt} in equation (1) is given by:

$$\varepsilon_{jt} = \beta_j r_t + \tilde{\varepsilon}_{jt},$$

where $\tilde{\varepsilon}_{jt}$ is independent of both ω_t and r_t as well as all other errors in the model, and its marginal distribution is $N\left(0, 1-\alpha_j^2-\beta_j^2\right)$. We restrict attention to slants β_j satisfying $\beta_j^2 \leq 1-\alpha_j^2$. Positive values of β_j represent conservative slant, while negative values represent liberal slant. We say the agent and source j are *like-minded* if b and β_j are non-zero and have the same sign and *opposite-minded* if they are non-zero and have opposite signs.

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

The form of bias and slant we have introduced here is fully consistent with the model of Section 2.1. It simply represents a particular parameterization of the covariance of the errors $(\varepsilon_{1t},...,\varepsilon_{Jt},\eta_t)$. Moreover, if either (i) $b=0$ or (ii) $\beta_j=0$ for all j, this parameterization is consistent with the agent's prior beliefs defined in Section 2.2. The main case of interest, however, is where both the agent and the outlets may be biased, and so neither (i) nor (ii) hold. In this case, the true correlation of the errors violates the agent's assumption that x_t is independent information, and so the true θ_0 lies in a region of the parameter space outside the support Θ^{prior} of her prior. It is in this sense that we refer to $b \neq 0$ as a bias.

3 Asymptotic Beliefs

3.1 Beliefs about θ

In the case where both the agent and at least some sources are biased, our model is an example of Bayesian learning under misspecification (Lian 2009). Characterizing the evolution of beliefs in such cases—when the observed data generating process lies outside the support of the agent's prior—is complicated in general, and can lead to instability or lack of convergence in the limit (Berk 1966). In the context of our model, however, we show that the data the agent observes can always be rationalized by some $\theta \in \Theta^{prior}$ that *does* fall within the support of her prior. Thus, the data will never violate her model of the world, and we show that her beliefs will be well behaved asymptotically as a result.

Let $\mathscr{F}$ be the set of all mean zero joint normal distributions on $\mathbb{R}^{J+1}$ with a non-degenerate covariance matrix all of whose diagonal elements are equal to one. For any θ (including those not in Θ^{prior}), the implied distribution of (s_t, x_t) is an element of $\mathscr{F}$ which we denote F_θ. Note that this mapping is many to one, so we may have $F_\theta = F_{\theta'}$ for $\theta \neq \theta'$. In this section, we characterize the agent's asymptotic beliefs when the data are generated by an arbitrary $F \in \mathscr{F}$. In the following section, we consider the implications for the specific F generated by the model of ideological bias in Section 2.3.

Given $F \in \mathscr{F}$, let the vector ρ_F denote the correlation of x_t with s_t under F, and let Σ_F denote the correlation matrix of s_t under F. Note that F is fully characterized by ρ_F and Σ_F. We say that a parameter θ is consistent with a statistic such as ρ_F or Σ_F if that statistic is the same under F_θ.

We first show that *any* element of $\mathscr{F}$ can be rationalized by some θ that falls within the support Θ^{prior} of the agent's prior. We then characterize, for any F, the subset of Θ^{prior} that is (i) consistent with the correlations ρ_F and (ii) consistent with both ρ_F and Σ_F. The former will define the set of parameters a single-homing agent will come to believe are possible, and the latter will define the set of parameters a multi-homing agent will come to believe are possible. With these building blocks in place, we then characterize the limit as $t \to \infty$ of both single-homing and multi-homing agents' beliefs under any $F \in \mathscr{F}$.

Lemma 1. *For any $F \in \mathscr{F}$, there exists $\theta \in \Theta^{prior}$ such that $F_\theta = F$.*

Proof. Choose θ with $a = 1$, $\alpha = \rho_F$, $\rho = \rho_F$, and $\Sigma = \Sigma_F$. □

In order for a value of the parameter $\theta \in \Theta^{prior}$ to be consistent with the observed correlations ρ_F, it must be the case that the a associated with θ satisfies $a \geq \left|\rho_{Fj}\right|$ for all j. This is the case because the correlation of x_t and s_{jt} under the independence assumption is $a\alpha_j$, and we have both $a \geq 0$ and $\left|\alpha_j\right| \leq 1$. We denote the resulting lower bound on the possible values of a for a given $F \in \mathscr{F}$ by $\underline{a}_F = \max_j \left\{\left|\rho_{Fj}\right|\right\}$.

For θ to also be consistent with the covariances Σ_F observed by a multi-homing agent, it must also be the case that the a associated with θ satisfies $a \geq \sqrt{(\rho_F)'(\Sigma_F)^{-1}\rho_F}$. This is the case because under the independence assumption the R^2 of a population regression of x_t on the full vector s_t cannot exceed a^2.[4] We denote this tighter lower bound by $\underline{\underline{a}}_F = max\left\{\underline{a}_F, \sqrt{(\rho_F)'(\Sigma_F)^{-1}\rho_F}\right\}$.

Lemma 2. *The set $\Theta^{prior}_{\rho_F}$ of $\theta \in \Theta^{prior}$ consistent with ρ_F is the set of all $(a, \alpha, \rho, \Sigma) \in \Theta$ with (i) $a \in [\underline{a}_F, 1]$, (ii) $\alpha_j = \frac{\rho_{Fj}}{a} \forall j$, and (iii) $\rho = \rho_F$.*

Proof. See appendix. □

Lemma 3. *The set $\Theta^{prior}_{\rho_F,\Sigma_F}$ of $\theta \in \Theta^{prior}$ consistent with both ρ_F and Σ_F is the set of all $(a, \alpha, \rho, \Sigma) \in \Theta$ with (i) $a \in \left[\underline{\underline{a}}_F, 1\right]$, (ii) $\alpha_j = \frac{\rho_{Fj}}{a} \forall j$, (iii) $\rho = \rho_F$, and (iv) $\Sigma = \Sigma_F$.*

Proof. See appendix. □

Note that for any F both $\Theta^{prior}_{\rho_F}$ and $\Theta^{prior}_{\rho_F,\Sigma_F}$ are non-empty by Lemma 1.

Proposition 1. *Suppose the distribution of (s_t, x_t) is $F = F_{\theta_0}$ for* some θ_0. *Then as $t \to \infty$, the agent's posterior beliefs on θ converge in distribution to a limit with density μ_∞. Under single-homing, μ_∞ is equal to zero for any $\theta \notin \Theta^{prior}_{\rho_F}$ and proportional to μ_0 for any $\theta \in \Theta^{prior}_{\rho_F}$. Under multi-homing, μ_∞ is equal to zero for any $\theta \notin \Theta^{prior}_{\rho_F,\Sigma_F}$ and proportional to μ_0 for any $\theta \in \Theta^{prior}_{\rho_F,\Sigma_F}$.*

---

[4]The R^2 from this population regression is $(\rho_F)'(\Sigma_F)^{-1}\rho_F$. Since the variance of x_t is one and the variance of the predictable component $a\omega_t$ of x_t is a^2, this R^2 cannot exceed a^2. Thus, $(\rho_F)'(\Sigma_F)^{-1}\rho_F \leq a^2$.

Proof. See appendix. □

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

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

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

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

Proposition 2. *Suppose the distribution of (s_t, x_t) is $F = F_{\theta_0}$ for* some θ_0. *Then the limiting posterior μ_∞^* exists.*

- *In the single-homing case, its support contains a unique value of a given by $a_F^* = \max\{a_0, \underline{a}_F\}$, a unique value of α with elements given by $\alpha_{Fj}^* = \frac{\rho_{Fj}}{a^*}\ \forall j$, a unique value of ρ given by $\rho_F^* = \rho_F$, and all values of Σ consistent with $(a_F^*, \alpha_F^*, \rho_F^*)$.*
- *In the multi-homing case, its support contains a unique value of a given by $a_F^* = \max\left\{a_0, \underline{\underline{a}}_F\right\}$, a unique value of α with elements given by $\alpha_{Fj}^* = \frac{\rho_{Fj}}{a^*}\ \forall j$, a unique value of ρ given by $\rho_F^* = \rho_F$, and a unique value of Σ given by $\Sigma_F^* = \Sigma_F$.*

Proof. See appendix. □

The values a^* and α^* will be one of the main focuses of our analysis. We refer to a^* as the agent's *confidence* and to α_j^* as the agent's *trust* in source j. It will turn out that for a biased agent both a^* and α^* may be different from the true values a_0 and α_0. If $a^* > a_0$, the agent overestimates the accuracy of her own feedback and we say she is *overconfident.*

3.2 Beliefs about ω_t

We define the agent's asymptotic *distribution of beliefs about* ω_t to be the distribution of her posterior mean $\overline{\omega}_t$ about ω_t given that her beliefs about θ are given by the limiting posterior μ_∞^*. For a single-homing agent this is defined separately for each source j she might observe. For a multi-homing agent it is defined for a period in which she observes all sources. For simplicity, we define it for all agents assuming we are in a period when she observes no feedback x_t, so we capture only the influence of the information source(s) she observes. The following characterization follows from standard conjugate prior results for the normal distribution.

Proposition 3. *Suppose the agent's beliefs about θ are the limiting posterior μ_∞^* associated with $F \in \mathscr{F}$. If a single-homing agent observes source j and no feedback in period t, her posterior belief about ω_t will be a normal distribution with mean $\overline{\omega}_t = \alpha_{Fj}^* s_{jt}$ and variance $v_t = 1 - \alpha_{Fj}^{*2}$. If a multi-homing agent observes the vector of realizations s_t and no feedback in period t, her posterior belief about ω_t will be a normal distribution with mean $\overline{\omega}_t = \alpha_F^{*'} \left(\Sigma_F^*\right)^{-1} s_t$ and variance $v_t = 1 - \alpha_F^{*'} \left(\Sigma_F^*\right)^{-1} \alpha_F^*$.*

Proof. See appendix. □

In both the single-homing and multi-homing cases, $\overline{\omega}_t$ is a linear function of the observed signals and so the asymptotic distribution of beliefs about ω_t follows immediately from the known distribution $N(0, \Sigma_F)$ of the signals. Note that because a biased agent's limiting posterior may be incorrect, her beliefs need not satisfy the martingale property and so her expected posterior mean of ω_t given the ideological state r_t may be significantly different from zero. When we consider multiple agents, this will also lead to systematic disagreement, and even the possibility that two agents may both be certain about the value of ω_t but differ in what they think is the true value.

When we focus below on the case where F is determined by our model of ideological bias, we will define the agent's *polarization* in a given period to be $\pi^* = |Corr(\overline{\omega}_t, r_t)|$. One way to motivate this definition is to imagine there are two agents who both have polarization π^*, one of whom has right-biased beliefs $\overline{\omega}_t^R$ with $Corr\left(\overline{\omega}_t^R, r_t\right) > 0$ and one of whom has left-biased beliefs $\overline{\omega}_t^L$ with $Corr\left(\overline{\omega}_t^L, r_t\right) < 0$. Define the *expected disagreement* between these agents to be $\mathrm{E}\left[\left(\overline{\omega}_t^R - \overline{\omega}_t^L\right)^2\right] \in \left[0, \frac{1}{4}\right]$. It is straightforward to show that holding constant the correlation of the agents' beliefs with ω_t, their expected disagreement will be proportional to π^* and will achieve the maximal value $\frac{1}{4}$ when $\pi^* = 1$.

4 Main Results: Single Homing

The previous section characterized the agent's asymptotic beliefs when the data are generated by an arbitrary distribution F. We now consider the implications in the case where that distribution is the one defined by our model of ideological bias in Section 2.3. We denote the distribution that results from that case by F_0, and we continue to use ρ_0 and Σ_0 to denote the associated ρ and Σ. We consider outcomes for a single-homing agent here, and turn to outcomes under multi-homing in the next section. All results in this section are straightforward implications of Proposition 2, and so we state them as corollaries. To simplify the exposition, we focus without loss of generality on the case of $b \geq 0$.

A useful starting point is to note that the observed correlation ρ_{0j} between the agent's feedback x_t and the signal of source j will be given by the value $a_0\alpha_{0j}$ that an unbiased agent would observe, plus the product of the agent's and the source's biases. The correlation will thus be higher than an unbiased agent would observe when the source is like-minded, and lower than an unbiased agent would observe when the source is opposite-minded.

Remark 1. The observed correlation with a source that has accuracy and bias (α, β) is:

$$\rho_{0j} = a_0\alpha + b\beta.$$

It will also be useful to define the source that would maximize this correlation for a given accuracy a_0 and bias b. Because we show below that a source with this (α, β) will also lead to the highest possible trust on the part of the agent, we refer to it as her trust-maximizing source. Recall that a combination (α,β) is feasible if $\beta^2 \le 1-\alpha^2$. It is straightforward to show that the values which maximize ρ_{0j} subject to this constraint are those given in the following remark.

Remark 2. The agent's *trust-maximizing source* is one with accuracy and slant

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

Finally, note that when the agent is unbiased, we recover the standard result that her beliefs converge to the truth (Blackwell and Dubins 1962). She will observe correlation $\rho_{0j} = a_0\alpha_{0j}$ with any source j. Because $a_0 \ge \underline{a}_F = \max_j \left\{\left|\rho_{0j}\right|\right\}$, she will not be overconfident ($a^* = a_0$). We then have $\alpha_j^* = \frac{\rho_{0j}}{a_0} = \alpha_{0j}$, so her trust in source j matches its true accuracy.

4.1 Overconfidence

We begin by considering the determinants of the agent's confidence a^*. It is immediate that when all information sources are unbiased, the agent will not be overconfident, since in this case all ρ_{0j} are weakly less than a_0, and so $a^* = a_0$.

Suppose, now, that we add at least one biased source to the market. The agent is overconfident if $\max_j \left\{\left|\rho_{Fj}\right|\right\} > a_0$. If the biases of all sources are sufficiently small, we will have $\left|\rho_{0j}\right| = \left|a_0\alpha_{0j} + b\beta_j\right| \le a_0$ for all j and so $a^* = a_0$. Otherwise, we will have $a^* > a_0$ and the agent will be overconfident. This will occur regardless of the size of her bias $b > 0$, provided that a_0 is sufficiently small. Note that this will be true even if the only biased source has $\beta_j < 0$—even if the source's bias is the opposite of the agent's bias, the *negative* correlation she observes with its signal could be enough to convince her that her own feedback is more accurate than she expected.

Finally, suppose that at least one source in the market is the agent's trust-maximizing source. The observed correlation with such a source will be $\rho_{0j} = \sqrt{a_0^2+b^2} > a_0$. The agent will therefore be overconfident for any $b > 0$ regardless of the value of a_0. By continuity of a^*, this will also be true if there is at least one source with (α,β) sufficiently close to the trust-maximizing values.

Corollary 1. *Suppose $b > 0$ and the agent is single-homing. If all sources are unbiased, the agent will not be overconfident. Adding at least one biased source to the market will lead the agent to* *be overconfident provided a_0 is sufficiently small. If at least one source is her trust-maximizing source, she will be overconfident regardless of the value of a_0, with $a^* = \sqrt{a_0^2 + b^2}$.*

4.2 Trust

We can apply Proposition 2 to derive the agent's trust as a function of her confidence a^*.

Remark 3. The agent's trust in information source j is

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

This expression shows two key forces that will drive our main results. First, small differences in bias b will have a large effect on trust whenever the agent believes her own feedback to be noisy ($a^* \approx 0$). To see the intuition for this, note that when x_t is noisy its correlation with even a perfectly accurate signal s_{jt} will be small. Its correlation with a completely inaccurate signal will be zero. Thus, small differences in observed correlation lead a rational agent to infer large differences in the accuracy of s_{jt}. This *amplification mechanism* is the reason why even small amounts of bias in our model can translate into large differences in trust and substantial polarization.

Second, as the agent becomes more overconfident (a^* increases), she will trust all like-minded or unbiased sources less. In particular, henever an agent is overconfident she will end up believing unbiased sources are less accurate than they really are ($\alpha_j^* < \alpha_{0j}$). This provides the mechanism by which the introduction of biased sources into a market can lead agents to trust unbiased sources less. Note that this would not be possible if the agent's confidence a^* were fixed exogenously; in this case, trust in source j is solely determined by the observed correlation ρ_{0j}, and so trust in an unbiased source cannot be affected by the set of other sources available in the market.

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

From this graphical analysis, it is immediately apparent that the trust-maximizing source $(\alpha^{max}, \beta^{max})$

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

We can then derive the way trust depends on the set of sources available. If all sources are unbiased ($\beta_j = 0$), we know from above that $a^* = a_0$ and so $\alpha_j^* = \alpha_{0j}$. Thus, the agent's trust in these sources is accurate. Suppose, now, that we add at least one biased source to the market. A key question of interest is how this changes her trust in unbiased sources. If all sources' biases are sufficiently small relative to a_0 and b, the agent will still not be overconfident and so her trust in unbiased sources will not change. It follows from Corollary 1 and Remark 3, however, that in our main case of interest where a_0 is small, her trust in such sources will fall. The agent will also tend to trust like-minded sources excessively (in the sense that $\alpha_j^* > \alpha_{0j}$). She may not do so if their biases β_j are very small, because the continuity of the problem implies her trust in such a source must be close to her trust in an unbiased source (for which $\alpha_j^* < \alpha_{0j}$). But it is easy to see that she will trust a source j excessively provided its bias β_j is sufficiently large.

Finally, suppose that at least one source in the market is the agent's trust-maximizing source. We know from Corollary 1 that the agent will be overconfident in this case, and so she will underestimate the accuracy of unbiased sources. Moreover, plugging her confidence $a^* = \sqrt{a_0^2 + b^2}$ and the values of $(\alpha^{max}, \beta^{max})$ from Remark 2 into Remark 3 shows that her confidence in a trust maximizing source will in fact be $\alpha_j^* = 1$. She will come to believe that a trust maximizing source is *perfectly* accurate, reporting $s_{jt} = \omega_t$ with probability one.

Corollary 2. *Suppose $b > 0$ and the agent is single-homing. If all sources are unbiased, the agent's trust in all sources will be equal to their true accuracies ($\alpha_j^* = \alpha_{0j}$). Provided a_0 is sufficiently small, adding at least one biased source to the market will lead the agent to underestimate the accuracy of unbiased sources ($\alpha_j^* < \alpha_{0j}$), and to overestimate the accuracy of sufficiently biased like-minded sources ($\alpha_j^* > \alpha_{0j}$). Regardless of the value of a_0, if at least one source is her trust-*

maximizing source, she will underestimate the accuracy of unbiased sources and believe the trust-maximizing source is perfectly accurate ($\alpha_j^* = 1$).

4.3 Polarization

In any period where an agent observes a source with accuracy and bias (α, β) for whom her trust is α_j^*, her posterior mean will be given by $\overline{\omega}_t = \alpha_j^* (\alpha\omega_t + \beta r_t)$. This implies that her polarization will be $\pi^* = \alpha_j^* \beta$.

Polarization in any period when she observes an unbiased source will therefore be zero. Since the agent's trust in a like-minded source must be strictly positive, her polarization in any period when she observes such a source must be strictly positive as well.

The main question of interest is whether polarization can be large even when the agent's bias b is small. It turns out that this will be the case whenever the agent observes a trust-maximizing source, provided that a_0 is also small. To see this, recall that the bias of a trust maximizing source is $\beta^{max} = b/\sqrt{a_0^2 + b^2}$ and the agent's trust in such a source is $\alpha_j^* = 1$ by Corollary 2. Thus, her polarization will be equal to the absolute value of the source's bias: $\pi^* = |b|/\sqrt{a_0^2 + b^2}$. This will be at least $1/\sqrt{2}$ provided that $a_0 \leq b$, and it will approach one in the limit as a_0 becomes small.

Corollary 3. *Suppose $b > 0$ and the agent is single-homing. In any period in which the agent observes an unbiased source, her polarization will be $\pi^* = 0$. In any period in which she observes a like-minded source, her polarization will be strictly positive. If at least one source is her trust-maximizing source, her polarization in periods where she observes that source will be at least $\pi^* = \frac{1}{\sqrt{2}}$ if $a_0 \leq b$, and will approach $\pi^* = 1$ in the limit as $a_0 \to 0$.*

5 Main Results: Multi-Homing

A common intuition is that polarization and divergent trust could be reduced or eliminated if agents were exposed to an ideologically diverse set of information sources. In any given period, agents might observe both biased and unbiased sources and so have less extreme beliefs than if they observed their preferred biased source alone. Over time, the ability to compare the reports of different outlets might help them more accurately identify trustworthy sources.

Both of these forces are present in our model. First, suppose we focus for simplicity on a case where the agent's trust is α^* and she has come to believe that all of the ε_{jt} of all the sources in the market are independent (i.e., her belief about Σ converged to a value $\Sigma^* = \alpha^* \alpha^{*'}$). Suppose all sources are on the frontier and that under single-homing she would only observe the most biased source. Then, because we can see from Proposition 3 that her posterior beliefs will be a weighted average of the signals she observes, it is correct that her beliefs under multi-homing will be more correlated with ω_t and less correlated with r_t than under single-homing.

Second, consider how asymptotic trust will vary with the correlation she observes among the signals over time. Intuitively, we might expect her to trust a set of sources more if their reports are positively correlated with each other, and less if they are uncorrelated or even negatively correlated with each other. This is indeed what Proposition 2 implies. Note that the agent's confidence a^* will be greater under multi-homing than under single-homing if and only if $\sqrt{(\rho_0)'(\Sigma_0)^{-1}\rho_0} > \max_j\left\{\left|\rho_{Fj}\right|\right\}$. It is straightforward to show that this condition will not hold if the elements of Σ_0 are sufficiently positive, but will hold if they are all sufficiently small or negative. This in turn means that an agent's trust in all sources will be unchanged under multi-homing when she observes sufficiently positive correlation among sources (because a^* is unchanged), but may fall under multi-homing if she sees sufficient disagreement among sources (because a^* increases and trust is decreasing in a^*).

In spite of these forces being present, however, we can show that multi-homing does not in general reduce polarization and divergent trust, and may in fact make them worse. We begin with four lemmas that together will provide the proof of our main result in this section.

The first lemma follows from observing that when at least one source is the agent's trust maximizing source, we must have $\sqrt{(\rho_0)'(\Sigma_0)^{-1}\rho_0} \leq \max_j\left\{\left|\rho_{Fj}\right|\right\}$, which implies the agent's confidence and trust under multi-homing is the same as under single-homing. In essence, a trust-maximizing source already makes the agent sufficiently over-confident under single-homing that also observing the correlations among the different s_{jt} does not force any change in her beliefs about her own accuracy.

Lemma 4. *Suppose at least one source is the agent's trust-maximizing source. Then the agent's asymptotic confidence a^* and asymptotic trust in all sources α^* is the same under single–homing and under multi-homing.*

Proof. See appendix. □

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

Lemma 5. *Suppose at least one source is the agent's trust-maximizing source. Then the distribution of the agent's beliefs about ω_t under multi-homing is invariant to what other sources are in the market, and in every period her posterior belief about ω_t is degenerate at the value equal to the signal of her trust-maximizing source.*

Proof. See appendix. □

The third lemma shows that observing *any two* distinct sources on the frontier is equivalent for a multi-homing agent to observing her trust-maximizing source. This perhaps surprising result follows from showing that for any two frontier signals s_{jt} and s_{kt}, there exists a linear combination of the two that is equal to the trust-maximizing signal with probability one. This is easy to see in the event where j is an unbiased source ($s_{jt} = \omega_t$) and k is a perfectly biased source ($s_{kt} = r_t$); in this case we can take the linear combination $\alpha^{max} s_{jt} + \beta^{max} s_{kt}$. The proof of the lemma shows that such a linear combination always exists. Note that this holds even if the biases of sources j and k are both *opposite* to that of the agent.

Lemma 6. *Suppose the market contains at least two frontier sources with distinct biases. Then the confidence, trust, and distribution of beliefs of a multi-homing agent are the same as in the case where the market contains the agent's trust-maximizing source.*

Proof. See appendix. □

The fourth lemma addresses the case in which sources are not necessarily located on the frontier. We show that provided the number of sources is large, and provided there is a minimal amount of diversity in their slants, this case is also equivalent to observing her trust-maximizing source. We formalize this notion of a "large and diverse" set of information as follows.

Definition 3. Let $d_{\alpha\alpha} := \sum_{j=1}^{J} \alpha_j^2/\kappa_j^2$, $d_{\beta\beta} := \sum_{j=1}^{J} \beta_j^2/\kappa_j^2$, and $d_{\alpha\beta} := \sum_{j=1}^{J} \alpha_j\beta_j/\kappa_j^2$ where $\kappa_j^2 = 1 - \alpha_j^2 - \beta_j^2$. The limit where there are *many and diverse information sources* is any situation in which (i) $J \to \infty$, (ii) $d_{\alpha\alpha} \to \infty$, $d_{\beta\beta} \to \infty$, and $d_{\alpha\beta}^2/d_{\alpha\alpha}d_{\beta\beta} \to 0$ as $J \to \infty$, and (iii) $0 < \alpha_j^2 + \beta_j^2 < 1$ for all j.

The proof of the lemma shows that in this case, too, the agent can construct a linear combination of the sources' signals whose value will be equal to the signal of the agent's trust-maximizing source in the limit as the number of sources grows large.

Lemma 7. *In the limit where there are many and diverse information sources, the confidence, trust, and distribution of beliefs of a multi-homing agent approach those in the case where the market contains the agent's trust-maximizing source.*

Proof. See appendix. □

The following proposition is then immediate.

Proposition 4. *Suppose one of the following is true: (i) one of the sources is the agent's trust-maximizing source; (ii) the market contains at least two frontier sources with distinct biases; (iii) we take the limit as there are many and diverse sources. Then a multi-homing agent with any* $b > 0$*:*

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

Under condition (i), the agent's confidence, trust, and beliefs about ω_t *are the same under multi-homing and under single-homing. If conditions (ii) or (iii) but not (i) hold, an agent under multi-homing, relative to the same agent under single-homing, is more overconfident, less trusting in source* j *if* $\rho_j > 0$*, and more trusting if* $\rho_j < 0$*. If additionally both the accuracies* $|\alpha_j|$ *are sufficiently large and the biases* $|\beta_j|$ *are sufficiently small for all sources j, then the agent is more polarized under multi-homing than under single-homing.*

6 Conclusion

We present a model in which agents learn about policy-relevant states through observing signals from multiple information sources over time. In the model, agents learn about the accuracies of the information sources based on feedback that they assume to be independent. However, arbitrarily small biases in the feedback result in persistent divergence in both trust in information sources and beliefs about facts. Partisans end up trusting unreliable ideologically aligned sources more than accurate neutral sources, and also becoming overconfident in their own feedback. Increasing the number of available information sources can deepen rather than mitigate ideological disagreement.

Perhaps most surprisingly, divergent trust and beliefs can arise to a similar extent whether agents selectively view only ideologically aligned sources or are exposed to a diverse range of sources. These predictions of the model are consistent with recent evidence suggesting that the magnitude of selective exposure has generally been limited, even while individuals persistently disagree about both objective facts and which sources can be trusted to provide reliable information about those facts.

References

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

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

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.

Benoit, Jean-Pierre and Juan Dubra. 2014. "A theory of rational attitude polarization." Working paper.

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

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

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

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

DeGroot, Morris. 1974. "Reaching a consensus." *Journal of American Statistical Association* 69: 118- 121.

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

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

Fryer, Roland G., Jr., Philipp Harms, and Matthew O. Jackson. 2015. "Updating beliefs when evidence is open to interpretation: Implications for bias and polarization." Working paper.

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

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

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

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

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

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

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

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

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

Lai, T. L. and Herbert Robbins. 1985. "Asymptotically Efficient Adaptive Allocation Rules." *Advances in Applied Mathematics* 6, 4-22.

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

Lord, Charles G., Lee Ross, and Mark R. Lepper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology* 37: 2098–2109.

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

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

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

Pew survey. 2013. "GOP Deeply Divided Over Climate Change." Accessed at <http://www.people-press.org/2013/11/01/gop-deeply-divided-over-climate-change/> on September 28, 2014.

Pew survey. 2013. "Republicans' views on evolution." Accessed at <http://www.pewresearch.org/fact-tank/2014/01/03/republican-views-on-evolution-tracking-how-its-changed/> on September 28, 2014.

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

Public Policy Polling. 2013. "National Survey Results." Accessed at <http://www.publicpolicypolling.com/pdf/20 on September 28, 2014.

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

Sethi, Rajiv and Muhamet Yildiz. 2016. "Communication with Unknown Perspectives." Working paper.

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

Figure 1: Iso-Trust Curves

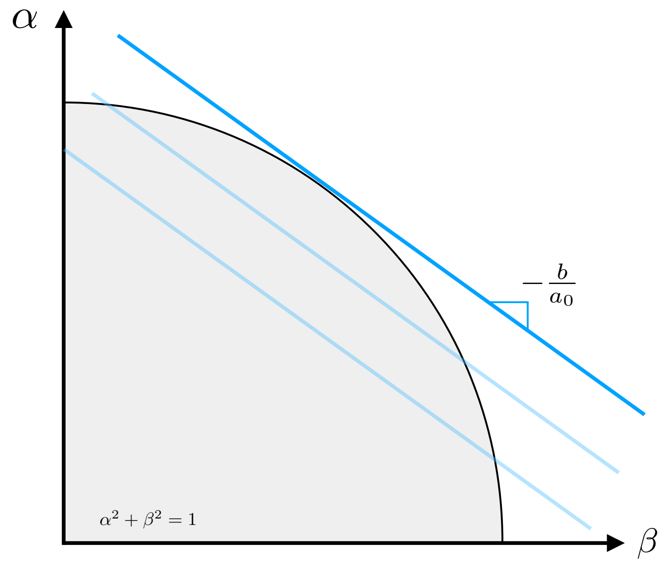

Appendices

A Proofs

A.1 Proof of Lemma 2

Suppose $\theta \in \Theta^{prior}_{\rho_F}$. Because θ is consistent with ρ_F, we must have $\rho_j = \rho_{Fj}$ for all j, which implies (iii). Because θ is in Θ^{prior}, it must satisfy $E\left(\varepsilon_{jt}\eta_t\right) = 0$ and therefore $\rho_j = \alpha_j a$ for all j, which implies (ii). Because $a \in (0,1]$ and $\left|\alpha_j\right| \leq 1$ by assumption, this in turn implies (i). Conversely, suppose that a tuple $(a,\alpha,\rho,\Sigma) \in \Theta$ satisfies (i)-(iii). Because $\alpha_j = \frac{\rho_j}{a}$ for all j, this tuple satisfies $E\left(\varepsilon_{jt}\eta_t\right) = 0$. Thus $\theta \in \Theta^{prior}$. By (iii), θ is consistent with ρ_F.

A.2 Proof of Lemma 3

Suppose $\theta \in \Theta^{prior}_{\rho_F,\Sigma_F}$. Because $\Theta^{prior}_{\rho_F,\Sigma_F} \subset \Theta^{prior}_{\rho_F}$, Lemma 2 implies that θ satisfies (ii)-(iii). Because θ is consistent with Σ_F, it must satisfy (iv). To see that it must satisfy (i), note that the R^2 from a population regression of ω_t on s_t when the data are generated by θ is $\alpha'\Sigma^{-1}\alpha \leq 1$. This implies that $\rho'\Sigma^{-1}\rho = (a \otimes \alpha)'\Sigma^{-1}(a \otimes \alpha) \leq a^2$, and so $a \geq \sqrt{\rho'\Sigma^{-1}\rho}$. This combined with the fact that Lemma 2 implies $a \geq \underline{a}_F$ implies that (i) must hold.

A.3 Proof of Proposition 1

Define D^t as the sequence of signal and feedback realizations that an agent has observed by the end of period t. Let $\mu(D^t \mid \theta)$ denote the probability distribution function over D^t under F_θ. We are interested in the limit of the agent's posterior $\mu(\theta \mid D^t)$ given prior $\mu_0(\theta)$ as $t \to \infty$.

We partition θ into two component vectors θ^{obs} and θ^{unobs}, which denote the parameters the agent can or cannot learn about from observing D^t. Under single-homing, $\theta^{obs} = \rho$ and $\theta^{unobs} = (a,\alpha,\Sigma)$. Furthermore, under single-homing, we define $\theta^{obs}_F = \rho_F$ and $\Theta^{prior}_F = \Theta^{prior}_{\rho_F}$. Under multi-homing, $\theta^{obs} = (\rho,\Sigma)$ and $\theta^{unobs} = (a,\alpha)$. Analogously, under multi-homing, we define $\theta^{obs}_F = (\rho_F,\Sigma_F)$ and $\Theta^{prior}_F = \Theta^{prior}_{\rho_F,\Sigma_F}$. Slightly abusing notating, let $\mu_0\left(\theta^{obs}\right)$ be the marginal prior distribution over θ^{obs} and $\mu\left(\theta^{obs} \mid D^t\right)$ be the conditional marginal distribution of θ^{obs} given D^t.[5]

[5] More precisely, let $\mu_0\left(\theta^{obs}\right) = \int_{\Theta^{unobs}} \mu\left(\theta^{obs},\theta^{unobs}\right) d\theta^{unobs}$ and so on.

Let $\tilde{\mu}_0\left(\theta^{unobs} \mid \theta^{obs}\right) = \mu_0(\theta)/\mu_0\left(\theta^{obs}\right)$.

Lemma A1. *$\mu\left(\theta \mid D^t\right) = \tilde{\mu}_0\left(\theta^{unobs} \mid \theta^{obs}\right)\mu\left(\theta^{obs} \mid D^t\right)$ for any realized D^t.*

Proof. We first show that under both single- and multi-homing, $\mu\left(D^t \mid \theta\right) = \mu\left(D^t \mid \theta^{obs}\right)$. Under single-homing, the agent observes $\left(x_t, s_{jt}\right)$ for some j in each period. Each element in $\left(x_t, s_{jt}\right)$ is i.i.d. normal with mean zero and variance one, and their covariance is ρ_j. Therefore, $\mu\left(D^t \mid \theta\right) = \mu\left(D^t \mid \rho\right)$ under single-homing. Under multi-homing, the agent observes $\left(x_t, s_t\right)$ in every period. The vector $\left(x_t, s_t\right)$ is i.i.d. $N\left(0, \bar{\Sigma}\right)$, where $\bar{\Sigma} = \begin{bmatrix} 1 & \rho' \\ \rho & \Sigma \end{bmatrix}$. This implies that $\mu\left(D^t \mid \theta\right) = \mu\left(D^t \mid \rho, \Sigma\right)$ under multi-homing.

We can then conclude that

$$\begin{aligned}
\mu\left(\theta \mid D^t\right) &= \mu\left(D^t \mid \theta\right)\mu_0(\theta)/\mu\left(D^t\right) \\
&= \mu\left(D^t \mid \theta^{obs}\right)\mu_0(\theta)/\mu\left(D^t\right) \\
&= \mu\left(D^t \mid \theta^{obs}\right)\tilde{\mu}_0\left(\theta^{unobs} \mid \theta^{obs}\right)\mu_0\left(\theta^{obs}\right)/\mu\left(D^t\right) \\
&= \tilde{\mu}_0\left(\theta^{unobs} \mid \theta^{obs}\right)\mu\left(\theta^{obs} \mid D^t\right)
\end{aligned}$$

where the first equality follows from Bayes rule, the second follows from the preceding paragraph, the third follows from the definition of $\tilde{\mu}$, and the fourth is another application of Bayes rule. □

Lemma A2. *In the limit as $t \to \infty$, $\mu\left(\theta^{obs} \mid D^t\right)$ converges uniformly to a point mass at $\theta^{obs} = \theta^{obs}_F$ where $F = F_{\theta_0}$.*

Proof. Let $P_{\theta^{obs}}$ denote the distribution of the signal(s) and feedback observed by an agent in a given period, indexed by parameter $\theta^{obs} \in \Theta^{obs}$. By assumption, as $t \to \infty$, the multi-homing agent observes realizations $\left(s_t, x_t\right)$ infinitely many times, and this vector is i.i.d. over time with a distribution parameterized by (ρ, Σ). Similarly, the single-homing agent observes realizations $\left(s_{jt}, x_t\right)$ for each j infinitely many times, and this vector is i.i.d. over time with a distribution parameterized by ρ_j. For the multi-homing agent, we take the limit as $t \to \infty$ once; for the single-homing agent, we take the limit as $t \to \infty$ separately for each j.

Note that (1) the maps $\theta^{obs} \mapsto P_{\theta^{obs}}$ are continuous in total variation norm (due to unit-normality); (2) $P_{\theta^{obs}}$ is differentiable in quadratic mean (due to unit-normality); (3) Θ^{obs} is σ-

compact (since Θ^{obs} is compact); (4) $P_{\theta^{obs}}$ has a nonsingular information matrix at θ_F^{obs} (since $\theta_F^{obs} = \theta_0^{obs} \in \text{int}\left(\Theta^{obs}\right)$); (5) the agent's prior μ_0 is absolutely continuous in a neighborhood of θ_F^{obs} with continuous positive density (by continuity assumption); and (6) $P_{\theta^{obs}} \neq P_{\theta^{obs\prime}}$ if $\theta^{obs} \neq \theta^{obs\prime}$, i.e. the model is identified. Therefore, by the Bernstein-von Mises theorem (van der Vaart 1998, Theorem 10.1 and Lemma 10.6), there exists a sequence of estimators that is uniformly consistent on Θ^{obs} for estimating θ^{obs}. Since uniform consistency is defined to be uniform convergence in distribution to the point mass at the true value, this proves the lemma. □

Together, the above two lemmas imply that $\mu(\theta \mid D)$ converges uniformly to a limit

$$\mu_\infty(\theta) = \begin{cases} \tilde{\mu}_0\left(\theta^{unobs} \mid \theta^{obs}\right) & \text{if } \theta^{obs} = \theta_F^{obs} \\ 0 & \text{otherwise} \end{cases}$$

for any $\theta \in \Theta^{prior}$, which completes the proof.

A.4 Proof of Proposition 2

Define θ^{obs}, θ^{unobs}, θ_F^{obs}, and Θ_F^{prior} as in the second paragraph of the proof of Proposition 1. Take any $\theta \in \Theta_F^{prior}$. From the proof of proposition 1, we have that $\mu_{\infty,n}(\theta) = \tilde{\mu}_{0,n}\left(\theta^{unobs} \mid \theta^{obs}\right) = \mu_{0,n}(\theta) / \int_{\Theta_F^{prior}} \mu_{0,n}(\theta)\, d\theta$. First, observe that it is not possible for $a < a_F^*$ by lemma 2 and 3. Second, observe that if $a > a_F^*$, we have $|a - a_0| > |a_F^* - a_0|$. By proposition 1, there exists $\theta' \in \Theta_F^{prior}$ with $a' = a_F^*$. Therefore, $\lim_{n\to\infty} \mu_{0,n}(\theta) = 0$ by the definition of $\mu_{0,n}$. It follows that $\mu_\infty^*(\theta) := \lim_{n\to\infty} \mu_{\infty,n}(\theta)$ is equal to $\lim_{n\to\infty} \tilde{\mu}_{0,n}\left(\theta^{unobs} \mid \theta^{obs}\right)$ if both $\theta^{obs} = \theta_F^{obs}$ and $a = a_F^*$ and is equal to zero otherwise, where the limit exists due to the continuity of $\mu_{0,n}$ and the fact that $\mu_{0,n}$ has full support over Θ^{prior}. The uniqueness of α in the support of μ_∞^* follows from the uniqueness of a in the support of μ_∞^*, since all $\theta \in \Theta_F^{prior}$ satisfies $\alpha = \rho / a$. The uniqueness of ρ (and Σ in the multi-homing case) follows from the condition that $\theta^{obs} = \theta_F^{obs}$.

A.5 Proof of Proposition 3

Suppose a single-homing agent observes source j and no feedback in period t. Under the limiting posterior μ_∞^* associated with $F \in \mathscr{F}$, the state ω_t and signal s_{jt} are joint distributed

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

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

Now suppose a multi-homing agent observes the vector of realizations s_t and no feedback in period t. Under the limiting posterior μ_∞^* associated with $F \in \mathscr{F}$, the state ω_t and signals s_t are joint distributed

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

By the properties of the multivariate normal distribution, the conditional distribution of ω_t given s_t is therefore $N\left(\alpha_F^{*\prime}\left(\Sigma_F^*\right)^{-1} s_t, 1-\alpha_F^{*\prime}\left(\Sigma_F^*\right)^{-1}\alpha_F^*\right)$, confirming the second part of the proposition.

A.6 Proof of Lemma 4

Suppose the distribution of (s_t, x_t), denoted by $F \in \mathscr{F}$, is parameterized by $y = \begin{bmatrix} a_0 & b \end{bmatrix}'$ and $z_j = \begin{bmatrix} \alpha_j & \beta_j \end{bmatrix}'$ for each j. Let Z be the $2 \times J$ matrix where the jth column is z_j. Let $P_Z := Z\left(Z^T Z\right)^{-1} Z'$ be the projection onto the column space of Z.

Lemma A3. *Under single-homing, if there exists j that is the agent's trust-maximizing source, then the agent's confidence is $a^* = \sqrt{a_0^2+b^2}$; otherwise, $a^* \in \left[a_0, \sqrt{a_0^2+b^2}\right)$. Under multi-homing, if $y \in span(Z)$, then the agent's confidence is $a^* = \sqrt{a_0^2+b^2}$.*

Proof. First consider a single-homing agent. If there exists a source j which maximizes the agent's trust among all feasible sources, it follows that for any $j' \neq j$, $\left|\rho_{Fj'}\right| \leq \left|\rho_{Fj}\right| = \left|z_j' y\right| = a_0\left(\frac{a_0}{\sqrt{a_0^2+b^2}}\right) + b\left(\frac{b}{\sqrt{a_0^2+b^2}}\right) = \sqrt{a_0^2+b^2}$. It follows that $\underline{a}_F = \max_j\left\{\left|\rho_{Fj}\right|\right\} = \sqrt{a_0^2+b^2}$. If there does not exist a trust-maximizing source, then $\left|\rho_{Fj'}\right| < \left|\rho_{Fj}\right| = \sqrt{a_0^2+b^2}$ for all j'. Therefore,

$a_0 \leq \underline{a}_F < \sqrt{a_0^2 + b^2}$. Now consider a multi-homing agent. Suppose $y \in \text{span}(Z)$. Then $y = P_Z y$. Note that $\rho_F = Z'y$ and $\Sigma_F = Z'Z$. Therefore, $\sqrt{(\rho_F)'(\Sigma_F)^{-1}\rho_F} = \sqrt{y'P_Z y} = \sqrt{y'y} = \sqrt{a_0^2 + b^2}$. Furthermore, for any j, $|\rho_j| \leq \sqrt{a_i^2 + b_i^2}$ by the same logic as in the previous paragraph. Therefore, $\underline{\underline{a}}_F = \max\left\{\underline{a}_F, \sqrt{(\rho_F)'(\Sigma_F)^{-1}\rho_F}\right\} = \sqrt{a_0^2 + b^2}$. Lemma A3 follows by proposition 2. □

By remark 2, $y \in \text{span}(Z)$ if at least one source is the agent's trust-maximizing source. Lemma 4 then follows from applying lemma A3 and proposition 2.

A.7 Proof of Lemma 5

First, consider single-homing. If a single-homing agent observes her trust-maximizing source j and no feedback in period t, her confidence is $a^* = \sqrt{a_0^2 + b^2}$ by lemma A3. Furthermore, her trust for that source is $\alpha_j^* = 1$ by proposition 2. By proposition 3, her posterior belief about ω_t is therefore a point mass at $\overline{\omega}_t = \alpha^{max}\omega_t + \beta^{max} r_t$. The multi-homing case proceeds in the same manner and yields the equivalent posterior belief about ω_t, since confidence a^* and hence trust for the trust-maximizing source are unaffected by the addition of sources.

A.8 Proof of Lemma 6

Lemma A4. *Suppose a multi-homing agent observes the vector of realizations s_t and no feedback in period t. Under the limiting posterior μ_∞^* associated with $F \in \mathscr{F}$, the conditional distribution of the state ω_t given signals s_t is $N(\overline{\omega}_t, v)$, where*

$$\begin{aligned} \overline{\omega}_t &:= \frac{1}{a^*} y'Z\left(Z'Z + K\right)^{-1}\left(Z'w_t + \varepsilon_t\right) \\ v &:= 1 - \frac{1}{a^{*2}} y'Z\left(Z'Z + K\right)^{-1} Z^T y \end{aligned}$$

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

Proof. By proposition 3, $\overline{\omega}_t = \frac{1}{a^*}\rho_F'\Sigma_F^{-1}s_t$ and $v = 1 - \frac{1}{a^{*2}}\rho_F'\Sigma_F^{-1}\rho_F$. In our special parameterization of ideological bias, we have that $\rho_F = Z'y$ and $\Sigma_F = Z'Z + K$. The lemma follows from noting that $s_t = Z'w_t + \varepsilon_t$. □

Now suppose the market contains exactly two frontier sources with distinct biases. This means that $y \in \text{span}(Z)$. Hence, $y = P_Z y$. By lemma A3, the agent's confidence is $a^* = \sqrt{a_0^2 + b^2} = \sqrt{y'y}$. Her trust is ρ_0/a^* by proposition 2. Note that $K = 0$ since both sources are frontier. Thus by lemma A4, $\overline{\omega}_t = \frac{1}{a^*} y' P_Z w_t = \frac{1}{a^*} y' w_t = \alpha^{max}\omega_t + \beta^{max} r_t$. Similarly, $v = 1 - \frac{1}{a^{*2}} y' P_Z y = 1 - \frac{1}{a^{*2}} y' y = 0$. These confidence, trust, and distribution of beliefs are the same as those in the case where the market contains the agent's trust-maximizing source. The generalization to at least two such sources is immediate, because the addition of sources does not affect confidence a^* nor trust α^* for those two sources and hence does not affect inference about ω_t.

A.9 Proof of Lemma 7

Lemma A5. *Suppose conditions (ii) and (iii) in definition 3 holds. Let Z be the $2 \times J$ matrix where the jth column is $z_j = \begin{bmatrix} \alpha_j & \beta_j \end{bmatrix}'$. Then $\text{span}(Z) = \mathbb{R}^2$ for a sufficiently large J.*

Proof. Suppose $\text{span}(Z) \neq \mathbb{R}^2$. Then the set of vectors $\{z_j\}$ are linearly dependent, so $(\alpha_j, \beta_j) = (c_j\alpha, c_j\beta)$ for $\alpha \in (-1,1)$, $\beta \in (-1,1)$, and some $c_j \in (0,1)$ for each j. If $\alpha \neq 0$ and $\beta \neq 0$, then $d_{\alpha\beta} = \left(\frac{\beta}{\alpha}\right) d_{\alpha\alpha} = \left(\frac{\alpha}{\beta}\right) d_{\beta\beta}$ and $d_{\alpha\beta}^2 / d_{\alpha\alpha} d_{\beta\beta} = 1$ for all J. This contradicts the assumption for a sufficiently large J. If $\alpha = 0$ or $\beta = 0$, then either $d_{\alpha\alpha} = 0$ or $d_{\beta\beta} = 0$, which also contradicts the assumption. □

We now show that, in the limit where there are many and diverse sources, the multi-homing agent's confidence converges to $\sqrt{a_0^2 + b^2}$, the agent's trust to $\rho_0 / \sqrt{a_0^2 + b^2}$, and the agent's posterior belief about ω_t converges in distribution to a point mass at $\overline{\omega}_t = \alpha^{max}\omega_t + \beta^{max} r_t$.

By Woodbury's matrix identity, we can write $Z(Z'Z + K)^{-1} = (I - R) Z K^{-1}$, where $R := Q(I + Q)^{-1}$ and $Q := Z K^{-1} Z' = \begin{bmatrix} d_{\alpha\alpha} & d_{\alpha\beta} \\ d_{\alpha\beta} & d_{\beta\beta} \end{bmatrix}$. It is easy to check that

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

It is also easy to check that $Z(Z'Z + K)^{-1} Z' = (I - R) Q = R$. By assumption $d_{\alpha\alpha} \to \infty$, $d_{\beta\beta} \to \infty$, and $d_{\alpha\beta}^2 / d_{\alpha\alpha} d_{\beta\beta} \to 0$ as $J \to \infty$. Therefore, $\lim R = I$. This proves that $\lim Z(Z'Z + K)^{-1} Z' = I$.

Furthermore, note that $\lim Z\left(Z'Z+K\right)^{-1}=\lim\left(ZK^{-1}-RZK^{-1}\right)=0$. We can then conclude, by the weak law of large numbers (since $\vec{\varepsilon}_t$ is a vector of independent random variables with zero mean and bounded variance), that $Z\left(Z^TZ+K\right)^{-1}\vec{\varepsilon}_t\to_p 0$.

By lemma A5, Z spans $\mathbb{R}^2$ for sufficiently large J. Therefore $\lim a^*=\sqrt{a_0^2+b^2}$ by lemma A3. Her trust is $\rho_0/\lim a^*$ by proposition 2. It follows from lemma A4 and derivations in the previous paragraph that

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

and

$$v=1-\frac{1}{a^{*2}}y'Z\left(Z'Z+K\right)^{-1}Z^Ty\to_p 1-\frac{y'y}{\lim a^{*2}}=0,$$

which completes the proof.

A.10 Proof of Proposition 4

Everything except the two last sentences of the proposition is immediate from lemmas 4, 5, 6, and 7. The first part of the second to last sentence follows from lemma A3. The second part is a straightforward application of proposition 2. The final sentence can be proved by noting that as $\left|\alpha_j\right|\to 1$ and as $\left|\beta_j\right|\to 0$ for all j, the single-homing agent has confidence $a^*\to a_0$, trust $\alpha_j\to\alpha_{0j}$, and posterior mean $\overline{\omega}_t\to\omega_t$, whereas the multihoming agent has $\overline{\omega}_t\to_p\alpha^{max}\omega_t+\beta^{max}r_t$ by lemma 7.