Public Disagreement*

Rajiv Sethi[†] Muhamet Yildiz[‡]

November 9, 2009

Abstract

Members of different social groups often hold widely divergent public beliefs regarding the nature of the world in which they live. We develop a model that can accommodate such public disagreement, and use it to explore questions concerning the aggregation of distributed information and the consequences of social integration. The model involves heterogeneous priors, private information, and repeated communication until beliefs become public information. We show that when priors are correlated, all private information is eventually aggregated and public beliefs are identical to those arising under observable priors. When priors are independently distributed, however, some private information is never revealed and the expected value of public disagreement is greater when priors are unobservable than when they are observable. If the number of individuals is large, communication breaks down entirely in the sense that disagreement in public beliefs is approximately equal to disagreement in prior beliefs. Interpreting integration in terms of the observability of priors, we show how increases in social integration can give rise to less divergent public beliefs on average.

*We thank the Institute for Advanced Study for financial support and hospitality, and Danielle Allen, Roland Benabou, Sylvain Chassang, Glenn Loury, Eric Maskin, Debraj Ray, and seminar participants at IAS, ISI, NYU, Princeton, Rutgers and the University of Illinois at Urbana-Champaign for helpful comments.

[†]Department of Economics, Barnard College, Columbia University and the Institute for Advanced Study (rs328@columbia.edu)

[‡]Department of Economics, MIT and the Institute for Advanced Study (myildiz@mit.edu)

1 Introduction

Members of different social groups often hold widely divergent beliefs regarding the nature of the world in which they live. In many instances such beliefs are not openly expressed, and hence knowledge of the disparity remains confined to a relatively small set of observers. There are times, however, when a high profile event triggers public reactions that make knowledge of the divergence inescapable. A dramatic example of this occurred on October 3, 1995, when a nation transfixed by the criminal trial of O.J. Simpson tuned in to hear the announcement of the verdict. The following report describes the scene in New York's Times Square (Allen et al., 1995):

> "In the moments before the O.J. Simpson verdict was announced, the crowd moved as one, heads all tilted upwards, eyes trained on the giant video screen. But when the verdict was delivered, the crowd split into two distinct camps one predominantly black, the other white and each with a vastly different response. Many blacks... reacted with jubilation. Many whites wore faces of shock and anger directed not only at the verdict, but at the reaction from blacks... Throughout the country, the scene was similar. In Wall Street offices, college campuses, stores, train stations and outside the Los Angeles County Courthouse, the Simpson verdict drew reactions that split along racial lines."

Differences in reaction to the verdict reflected substantial racial differences in beliefs regarding the likelihood that Simpson was guilty. Brigham and Wasserman (1999) tracked such beliefs over the course of a year, starting with the period of jury selection in 1994 and ending three weeks after the announcement of the verdict. During jury selection 54% of whites and 10% of blacks in their sample thought that Simpson was "guilty" or "probably guilty". By the time closing arguments were concluded these numbers had risen to 70% for whites and 12% for blacks, reflecting an even larger racial gap. The final round of the survey, taken several days after the verdict and initial reaction had been made public, showed modest convergence but a significant remaining disparity, with 63% of whites and 15% of blacks declaring a belief in probable or certain guilt.

While reactions to the Simpson verdict may be the most visible manifestation of racial differences in beliefs, there are a number of other dimensions on which stark differences are a matter of public record. A 1990 survey by the New York Times and WCBS found that 29% of black respondents (as compared with 5% of whites) considered it to be true or possibly true that the AIDS virus was "deliberately created in a laboratory in order to infect black people." Almost 60% of blacks believed that it was true or possibly true that the government "deliberately makes sure that drugs are easily available in poor black neighborhoods," and 77% gave credence to the claim that "the government deliberately singles out and investigates black elected officials in order to discredit them in a way it doesn't do with white officials." The corresponding numbers for white respondents were 16% and 34% respectively. These differences cannot be attributed to differences in socioeconomic status or demographic characteristics (Crocker et al., 1999). A June 2008 survey found that while 5% of black respondents believed that Barack Obama was a Muslim, the corresponding figure was 12% for white respondents, and 19% for white evangelical protestants (Pew Research Center, 2008). And in a poll conducted just a few days after the last presidential election, 38% of black respondents but only 8% of whites stated that racial discrimination against blacks in the United States continues to be "a very serious problem" (CNN/Opinion Research 2008).[1]

The persistence of such public disagreement appears to conflict with the standard hypothesis in economic theory that differences across individuals in beliefs are due solely to differences in information. If this view were correct, then disagreement itself would be informative and lead to revised beliefs and eventual convergence (Geanakoplos and Polemarchakis, 1982). This is the insight underlying Aumann's (1976) theorem, which states that two individuals who have common priors and are commonly known to be rational must have identical posterior beliefs if these beliefs are themselves common knowledge, no matter how different their information may be. As suggested by Aumann (1976), the widespread public disagreement that one observes in practice can be attributed either to differences of priors on the underlying parameter that is being estimated or to systematic biases in computing probabilities, i.e., to differences of priors on the broader state space in which individuals update their beliefs.

In this paper, we develop a tractable framework which allows for public disagreement and can be used to explore questions concerning the aggregation of distributed information and the consequences of social integration. We consider a finite population of individuals who differ with respect to both their priors and their information about the state of the world. All priors and signals are assumed to be normally distributed; priors may or may not be correlated, and signals are independent. The profile of priors in the population may or may not itself be commonly known; we consider both cases. Given their priors and their information, individuals form beliefs and these beliefs are publicly and truthfully announced. The announcements are informative, and individuals update their beliefs based on them. This results in a further round of announcements, which may also be informative. The sequence of announcements continues until no further belief revision occurs. At the end of the process, all beliefs become public information; we call these *public beliefs*. We are interested in whether or not all distributed information is incorporated into public beliefs through the process of communication, and the manner in which the extent of disagreement in public beliefs is affected by patterns of social integration.

Given the heterogeneity of priors, public beliefs would involve some level of disagreement even if priors were observable, so that each person's signal could be deduced from his announcement. When priors are unobservable, the possibility arises that the process of communication may fail to aggregate all distributed information, resulting in different levels of disagreement relative to the case of observable priors. This happens because unobservability of priors gives rise to a natural

[1] There also exist striking belief differences across groups defined by political affiliation. For instance, a July 2009 poll by Research 2000 found that 93% of Democrats but only 47% of Republicans agreed with the statement that "Barack Obama was born in the United States of America". Such differences are harder to interpret since party affiliation is clearly endogenous with respect to such beliefs. Based on unpublished poll internals, however, Weigel (2009) estimates that 97% of black respondents but less than 30% of Southern whites agreed with the statement that Obama was born in the US.

signal-jamming problem. An individual's first announcement is a convex combination of his prior and his signal. Since other individuals observe neither the prior nor the signal, they can only extract partial information about each of these from the announcement. At the end of the first round of communication, therefore, beliefs do not reflect all distributed information. We show that when priors are uncorrelated, none of the subsequent announcements has any informational value. As a result, some distributed information remains uncommunicated, despite potentially unlimited rounds of communication. Public disagreement now stems not only from heterogeneous priors but also from informational differences that are induced by the fact that priors are privately observed.

The problem becomes especially acute in a large society. We show that when a fixed amount of information is distributed among a large number of individuals, unobservability of priors leads to a breakdown in communication: the difference between the public beliefs of any two individuals is approximately equal to the difference in their prior beliefs, as though no information had been received and communicated. Hence, in a large society, public disagreement is greater under unobservable priors than under observable priors at almost all realizations of priors and signals. On the other hand, in small groups, unobservability of priors can lead to smaller levels of public disagreement at some realizations, simply because a more optimistic person may receive a more pessimistic signal and cannot communicate his information fully. Even in this case, however, we show that the *expected value* of public disagreement must be larger when priors are unobservable and uncorrelated.

With correlated priors the situation is more complex. As long as each individual's prior is correlated with that of at least one other individual, we show that (subject to a regularity condition that is generically satisfied) all distributed information is fully incorporated into public beliefs *even if priors are unobservable*. While individuals may agree to disagree, their eventual beliefs are precisely what they would be if they had been able to observe each other's signals. This happens because the manner in which an individual responds to the announced beliefs of others reveals his beliefs about the priors of others, which in turn reveals his own prior. As a consequence, public beliefs in the case of unobserved (but correlated) priors are identical to those resulting from observable priors. However, convergence to public beliefs takes longer when priors are unobserved, and involves levels of statistical sophistication that far exceed those required for convergence under observable priors. Hence, we view this result as a benchmark, suggesting that the beliefs that emerge in the long run are invariant to the manner in which information is distributed in society and the pattern of observability of priors. The beliefs held before convergence has been attained, which we interpret as medium-run beliefs, exhibit all of the properties of public beliefs under independently distributed priors.

One interpretation of the assumption that priors are observable is that individuals understand the thought processes and perspectives of others, even if they do not share them. Such understanding could arise through social integration and mutual understanding that goes beyond the announcement of posterior beliefs. Under this interpretation, our results enable us to investigate the relationship between social integration and public disagreement. We do this by exploring a variant of the model with uncorrelated priors, two social groups and three possible information structures. We say that society is fragmented if no priors are observable, segregated if each individual observes only the priors of those within his own social group, and integrated if all priors are observed. Our earlier results imply that expected disagreement is greater under fragmentation than under integration. A segregated society with uncorrelated priors behaves in a manner similar to a fragmented society with correlated priors: all distributed information is aggregated in the long run. In the medium run, however, some novel effects arise. We show that the ex-ante expectation of public disagreement held by individuals in a minority group is smaller than the same expectation held by members of a majority group. Furthermore, the expected magnitude of public disagreement in the medium run is greater under segregation than under integration, and greater under fragmentation than under segregation.

When the population size is large, medium-run beliefs under segregation exhibit a number of intriguing characteristics. First, the bias is state dependent, and can be much larger than the ex-ante difference across groups in prior beliefs. Hence differences in priors can become *amplified* through communication under segregation. In fact, even if there is no ex-ante difference in prior beliefs, there will be disagreement after the first round of communication. Second, individuals belonging to a minority group face a disadvantage under segregation even though all individuals receive equally precise signals and have access to the same belief announcements. The disadvantage arises in the *interpretation* of public announcements. Since minorities (by definition) observe the priors of a smaller segment of the total population, their inability to extract full information from the announcements of others can be very costly. As a result, medium-run beliefs of majority group members are more closely aligned with reality (interpreted as the true state) than are the beliefs of minority group members. Finally, we show that when both groups are composed of ex-ante identical individuals, realized belief differences under segregation are greater than such differences under either integration or fragmentation. Segregation tends to homogenize within group beliefs at the expense of amplifying between group beliefs.

Our work contributes to a growing literature that allows for heterogeneity in prior beliefs.[2] In particular, Banerjee and Somanathan (2001), Van den Steen (2004), and Che and Kartik (2008) explore strategic communication under observable heterogeneous priors. Since heterogeneous priors lead to heterogeneous preferences, some information cannot be communicated (as in Crawford and Sobel, 1982). Our work differs in allowing priors not only to be heterogeneous, but also to be *unobserved*. Furthermore, communication in our model is truthful, non-strategic and two-sided. We consider non-strategic communication in order to focus on the role of unobservability of priors in communication. (Moreover, in the applications we have in mind, individuals do not face strong incentives to misrepresent their opinions.) In this we follow Geanakoplos and Polemarchakis (1982), who show how the agreement predicted by Aumann (1976) could arise through a sequence of truthful belief announcements. We adopt the same model of sequential announcement introduced there, but

[2] Heterogeneous priors play a role in many applications, including work on asset pricing (Harrison and Kreps 1978; Morris 1996; Scheinkman and Xiong 2003), political economy (Harrington 1993), bargaining (Yildiz 2003, 2004), organizational performance (Van den Steen 2005), political polarization (Dixit and Weibull 2007) and mechanism design (Morris 1994; Eliaz and Spiegler 2007; Adrian and Westerfield 2007).

apply it to the case of heterogeneous and possibly unobserved priors. Our work is also related to the literature on learning with heterogeneous priors (e.g. Freedman 1965; Acemoglu et al. 2008), which focuses on learning from external sources rather than from communication.

Our work may also be seen as part of a literature dating back to Loury (1977) on the economic effects of social integration; see Chaudhuri and Sethi (2008) and Bowles et al. (2009) for recent contributions. While this literature examines the effects of integration on income differences, our focus here is on disparities in beliefs. Such disparities can have significant welfare consequences. As Crocker et al. (1999) note, blacks and whites "exist in very different subjective worlds" and "a chasm remains... in the ways they understand and think about racial issues and events." Such differences in beliefs can make "communication and interaction across racial lines painful and difficult", as blacks find "their construal of reality flatly denied" and whites feel hurt or outraged that blacks give credence to conspiracy theories that they find bizarre or outlandish. In addition, beliefs affect responses to government policies such as public health initiatives aimed at reducing the spread of communicable diseases or the promotion of birth control. Most fundamentally, differences in beliefs about the fairness of the justice system or the extent of racial discrimination in daily life can have corrosive effects on the functioning of a democracy and erode confidence and participation in the political process. While a serious analysis of such welfare effects is beyond the scope of this paper, our analysis is motivated in part by the sense that persistent public disagreement can be welfare reducing in subtle but substantial respects.

The remainder of the paper is structured as follows. We introduce the model in Section 2, and explore a special two-person case in Section 3. When there are just two individuals, correlated priors result in the same limiting beliefs (and hence the same levels of expected disagreement) as commonly known priors. The general case is examined in Section 5, where it is shown that the irrelevance of observability result continues to hold as long as the primitives of the model satisfy a genericity condition. The case of uncorrelated priors (which fails this condition) is explored in Section 4, where we identify conditions under which observability of priors lowers expected disagreement relative to unobservability. Section 6 uses our results to explore the relationship between social integration and public disagreement, and Section 7 concludes.

2 The Model

There are n individuals $i \in N = \{1, 2, \ldots, n\}$ and an unknown real-valued parameter θ, which we call the state of the world. Individuals differ with respect to both their prior beliefs and their private information about the state of the world. Before the receipt of any information, individual i believes that θ is normally distributed with mean μ_i and unit variance:[3]

$$\theta \sim_i N(\mu_i, 1).$$

[3] We use the subscript i to denote the belief of i. For example, E_i and $E_i[\cdot|\cdot]$ denote the ex-ante- and the conditional- expectation operators under i's beliefs. We omit the subscript when all individuals agree. For example, $X \sim N(0,1)$ means that all individuals agree that X has the standard normal distribution. Likewise, E denotes the expectation operator when all individuals agree; e.g., $E[X]$ means that $E_i[X] = E_j[X] = E[X]$ for all $i, j \in N$.

Given these (possibly heterogeneous and privately observed) prior beliefs, each individual i observes a private signal x_i that is informative about θ with additive idiosyncratic noise ε_i:

$$x_i = \theta + \varepsilon_i.$$

All individuals agree that $\theta, \varepsilon_1, \ldots, \varepsilon_n$ are independently distributed, and that

$$\varepsilon_i \sim N(0, \tau^2).$$

Observing x_i, individual i updates[4] his belief about θ to a normal distribution with mean

$$A_{i,1} = \alpha\mu_i + (1 - \alpha)\, x_i \tag{2}$$

and variance

$$\alpha = \frac{\tau^2}{1 + \tau^2}. \tag{3}$$

Hence, one can think of μ_i as the manner in which individual i processes his information x_i, about which other individuals are uncertain. One can also think of x_i as the component of the belief of i that is perceived to be informative about θ by other individuals, and μ_i as the residual component, which is perceived by others to contain no information about θ. We refer to the pair (μ_i, x_i) as i's *type*, assuming that (μ_i, x_i) is privately known by i unless we explicitly specify that μ_i is observable, in which case μ_i will be common knowledge.

The priors $(\mu_1, \ldots, \mu_n)$ are distributed normally with mean $(\bar{\mu}_1, \ldots, \bar{\mu}_n)$ and variance-covariance matrix Σ with entries σ_{ij} for $i, j \in N$. A crucial assumption is that conditional on μ_i, individual i believes that the state θ, the others' priors $\mu_{-i} = (\mu_j)_{j \neq i}$, and the noise terms ε_j, $j \in N$, are all stochastically independent. That is, player i thinks that there is some uncertainty about how each individual j processes his information x_j, but does not think that the manner in which j updates his beliefs reflects any information about θ.

Within this framework, we consider a model of deliberation involving truthful communication of beliefs in a sequence of stages, as in Geanakoplos and Polemarchakis (1982). Once signals are received, beliefs are made public in period 1 by simultaneous (and truthful) announcements $A_{i,1}$, $i \in N$, where $A_{i,1}$ denotes player i's expectation of θ conditional on the prior μ_i and the signal x_i. After observing all announcements, individuals update their beliefs and simultaneously announce these updated beliefs $A_{i,2}$, $i \in N$, in period 2. Here $A_{i,2}$ denotes i's expectation of θ conditional on his own prior μ_i, his own signal x_i, and the others' initial announcements $A_{-i,1} = (A_{j,1})_{j \neq i}$. Individuals continue to update and announce their beliefs indefinitely. The limiting values of the sequence of announcements is denoted $A_{i,\infty}$ for $i \in N$. We call $A_{i,\infty}$ the *public belief of* i, emphasizing the fact that this belief becomes public information (i.e. common knowledge) at the

[4] Throughout the paper, we use the following well-known formula. If $\theta \sim N(\mu, \sigma^2)$ and $\varepsilon \sim N(0, v^2)$, then conditional on signal $s = \theta + \varepsilon$, θ is normally distributed with mean

$$E\left[\theta | s\right] = \frac{v^2}{\sigma^2 + v^2}\mu + \frac{\sigma^2}{\sigma^2 + v^2}s \tag{1}$$

and variance $\sigma^2 v^2 / \left(\sigma^2 + v^2\right)$.

end of the communication process. We assume that everything we have described to this point is common knowledge.

Remark 1. Since $(\mu_1, \ldots, \mu_n)$ may be correlated, i may think that μ_i is correlated with both μ_{-i} and θ, but μ_{-i} and θ are independent conditional on μ_i. Such seemingly inconsistent beliefs arise naturally as follows. Suppose that all potentially relevant historical facts are represented by a family $\{X_m\}_{m \in M}$ of random variables. Each individual i considers a set $\{X_m \mid m \in R_i\}$ of random variables to be relevant for understanding θ for some $R_i \subset M$; he considers the remaining random variables X_m with $m \notin R_i$ irrelevant. His conditional expectation of θ given $\{X_m \mid m \in R_i\}$ is μ_i, which is all the relevant information about θ in $\{X_m\}_{m \in M}$ according to i. Consequently, conditional on μ_i, μ_{-i} does not affect his beliefs about θ; i.e., he considers μ_{-i} and θ to be independent. On the other hand, at the ex-ante stage, if i assigns positive probability to $R_i \cap R_j \neq \varnothing$ for some $j \neq i$, then i considers μ_i and μ_j to be stochastically dependent.

Remark 2. Our model of deliberation presumes that individuals cannot directly communicate their information, or the manner in which they have incorporated their information into their beliefs, although they can fully communicate their resulting beliefs. This is because we think of information as a complex object consisting of many small bits and pieces, and the manner in which these are incorporated into one's final opinion is itself a complex process that involves interpretation in light of one's upbringing and experience. For simplicity, we represent this process using two real numbers: x_i, which incorporates everything that others find relevant, and μ_i, which represents everything that others find irrelevant.

We conclude this section by describing the two environments that we will investigate. We say that *priors are observable* if $(\mu_1, \ldots, \mu_n)$ is common knowledge (although drawn from an ex-ante distribution). We say that *priors are unobservable* if μ_i is privately known by i for each i. We use superscripts ck and u to denote variables in the observable and unobservable priors cases, respectively. For example, we write $A^{ck}_{i,k}$ or $A^{u}_{i,k}$ for the announcement of i at round k, depending on whether priors are observable or unobservable, respectively.

Under observable priors, public beliefs can be easily computed. Each individual i can deduce the signal x_j of any other individual j from the first round announcements. (Specifically, from (2), we have $x_j = \left(1+\tau^2\right) A_{j,1} - \tau^2 \mu_j$.) Hence, individuals extract the entire relevant signal

$$(x_1 + \cdots + x_n)/n = \theta + (\varepsilon_1 + \cdots + \varepsilon_n)/n,$$

where the noise has variance

$$\bar{\tau}^2 = \tau^2/n. \tag{4}$$

Using this signal, they form their public beliefs as follows:

$$A^{ck}_{i,\infty} = A^{ck}_{i,2} = \frac{\tau^2}{n+\tau^2}\mu_i + \frac{n}{n+\tau^2}\sum_{j=1}^{n}\frac{x_j}{n}. \tag{5}$$

Here, the expression for $A^{ck}_{i,2}$ follows from (1). Since all the available information is revealed by the first announcements, the updating stops at round 2. The difference between the public beliefs of any two individuals $i, j \in N$ is therefore simply

$$A_{i,\infty}^{ck} - A_{j,\infty}^{ck} = \frac{\tau^2}{n + \tau^2} \left(\mu_i - \mu_j \right) = \frac{\bar{\tau}^2}{1 + \bar{\tau}^2} \left(\mu_i - \mu_j \right). \tag{6}$$

Holding constant $\bar{\tau}^2$, this difference in beliefs is independent of n. That is, under observable priors, differences in public beliefs between any pair of individuals are due only to differences in priors, which are scaled down according to the precision $1/\bar{\tau}^2$ of the distributed information. This difference in beliefs serves as the benchmark against which we measure belief differences under unobservable priors.

3 Examples

Before proceeding to more general results, we consider the case of two individuals. We assume without loss of generality that $\mu_i \geq \mu_j$.

3.1 Observable Priors

Consider the case in which the priors μ_i and μ_j are common knowledge. Since $n = 2$, (6) implies that the difference in public beliefs is

$$A_{i,\infty}^{ck} - A_{j,\infty}^{ck} = \frac{\tau^2}{2 + \tau^2} \left(\mu_i - \mu_j \right). \tag{7}$$

Note that although each individual's public belief depends on the other's initial announcement, the difference in beliefs is independent of both initial announcements, and the individuals agree on the distribution of this difference.

3.2 Unobservable Independent Priors

Next consider the case in which the priors μ_i and μ_j are not observed, and are independently distributed, each with variance σ^2. First round beliefs and announcements are exactly as in the case of observable priors:

$$A_{i,1}^{u} = \alpha \mu_i + (1 - \alpha) x_i.$$

Observing $A_{j,1}^{u}$, all i can infer is that $\alpha \mu_j + (1 - \alpha) x_j$ is equal to $A_{j,1}^{u}$, and cannot know the specific values of each variable. Hence, he attributes some of the variation in $A_{j,1}^{u}$ to variation in μ_j and some to variation in x_j. More precisely, he observes an additional signal

$$\left(1 + \tau^2\right) A_{j,1}^{u} - \tau^2 \bar{\mu}_j = \theta + \tau^2 \left(\mu_j - \bar{\mu}_j \right) + \varepsilon_j \tag{8}$$

with additive noise $\tau^2 \left(\mu_j - \bar{\mu}_j \right) + \varepsilon_j$. The noise term has mean 0 and variance $\sigma^2 \tau^4 + \tau^2$. He then updates his beliefs to a normal distribution with mean

$$\begin{aligned} A_{i,2}^{u} &= \frac{\sigma^2 \tau^4 + \tau^2}{\alpha + \sigma^2 \tau^4 + \tau^2} A_{i,1}^{u} + \frac{\alpha}{\alpha + \sigma^2 \tau^4 + \tau^2} \left(\left(1 + \tau^2\right) A_{j,1}^{u} - \tau^2 \bar{\mu}_j \right) \\ &= \frac{1}{1 + \left(1 + \sigma^2 \tau^2\right) \left(1 + \tau^2\right)} \left(\left(1 + \sigma^2 \tau^2\right) \left(1 + \tau^2\right) A_{i,1}^{u} + \left(1 + \tau^2\right) A_{j,1}^{u} - \tau^2 \bar{\mu}_j \right). \end{aligned}$$

Here, the first equality is obtained by updating according to (1) starting from $\theta \sim N(A^u_{i,1}, \alpha)$ and using the signal in (8), and the second equality is by (3). Note that (unlike the case of commonly known priors) i puts greater weight on his own announcement than on that of j. This is because i does not know j's prior. When j announces a higher expectation $A^u_{j,1}$, i believes that with some probability j has obtained a higher value of the signal x_j, motivating i to increase his own expectation of θ too. He also thinks that, with some probability, the high announcement may be due to a bias towards higher values (i.e. larger μ_j), in which case i would not want to increase his expectation of θ. Consequently, each player's beliefs become less sensitive to the other's announcement than in the case of commonly known priors.

Even after the second round announcements, i does not know x_j, so there remains some relevant asymmetric information. In other words, some of the distributed information is not aggregated at the first round. One might hope that further announcements communicate more private information, resulting in the aggregation of the remaining distributed information. This is not the case, however. Since $A^u_{i,1}$ and $A^u_{j,1}$ are sufficient statistics for $A^u_{i,2}$ and $A^u_{j,2}$, the second round announcements provide no additional information, and

$$A^u_{i,2} = A^u_{i,3} = ... = A^u_{i,\infty}.$$

The difference in public beliefs is

$$\begin{aligned} A^u_{i,\infty} - A^u_{j,\infty} &= \frac{1}{1 + (1+\sigma^2\tau^2)(1+\tau^2)} \left(\sigma^2\tau^2\left(1+\tau^2\right)(A_{i,1} - A_{j,1}) + \tau^2\left(\bar{\mu}_i - \bar{\mu}_j\right)\right) \\ &= \frac{\tau^2}{1 + (1+\sigma^2\tau^2)(1+\tau^2)} \left(\left(\bar{\mu}_i - \bar{\mu}_j\right) + \sigma^2\tau^2\left(\mu_i - \mu_j\right) + \sigma^2\left(\varepsilon_i - \varepsilon_j\right)\right). \end{aligned} \tag{9}$$

The difference of opinion has three sources: the difference in the means of the distributions from which priors are drawn $\left(\bar{\mu}_i - \bar{\mu}_j\right)$, the difference in the realized values of the priors $\left(\mu_i - \mu_j\right)$, and the difference in information $\left(\varepsilon_i - \varepsilon_j\right)$. Since communication never completely eliminates informational differences, these differences affect public beliefs. Communication does, however, decrease the role of differential information as the coefficient of $(A_{i,1} - A_{j,1})$ is strictly less than 1. That is, differences in information play a larger role in affecting initial announcements than in affecting public beliefs. As in the common knowledge case, all individuals agree on the distribution of the difference in public beliefs.

Note from (9) that the two individuals will generally agree to disagree even if they have *identical* priors $\left(\mu_i = \mu_j\right)$, since they cannot deduce from the announcements that their priors are in fact identical. This makes transparent the obvious but sometimes overlooked fact that the standard common prior assumption requires not only that the players have the same prior, but also that this fact is itself commonly known. Furthermore, even if both individuals have identical priors and receive identical signals $\left(\varepsilon_i = \varepsilon_j\right)$, they may disagree once their beliefs have been communicated, provided that the priors themselves are not drawn from identical distributions. Communication can therefore lead to increased polarization when priors are unobserved even when individuals receive exactly the same information.[5]

[5] This phenomenon arises also when priors are observed (but heterogeneous), provided that they fail to satisfy the

In conclusion, uncertainty about the manner in which other individuals process information hinders the communication of relevant private information through the announcement of beliefs. Consequently, individuals hold different beliefs both because they have (possibly) different priors and because of different information.

3.3 A Comparison of Belief Differences

Note that $A_{i,\infty} - A_{j,\infty}$ measures the amount that i overestimates θ relative to j at the end of deliberation. Hence we call $A_{i,\infty} - A_{j,\infty}$ *public bias* of i relative to j. Since uncertainty regarding priors leads to less communication of information, one may think that it also leads to greater public bias. This is not the case. It may so happen that the individuals have very different priors, and knowledge of this may lead to a very large difference of opinion. Indeed, when the priors are not observed, by (9), any amount of public bias is possible, including no bias at all. In contrast, when the priors are common knowledge, by (7), the amount of public bias is constant, depending only on the difference in realized priors.

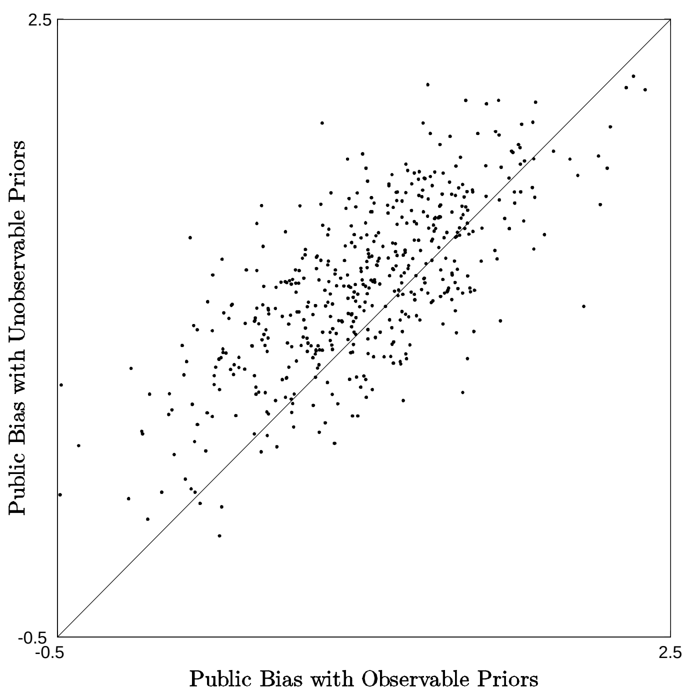

Figure 1. Public Bias with Observable and Unobsevable Priors

Figure 1 plots the values of public bias under observed and unobserved priors, respectively, for a set of randomly drawn type realizations.[6] Here, for each realization, the horizontal coordinate is $A^{ck}_{i,\infty} - A^{ck}_{j,\infty}$ and the vertical coordinate is $A^{u}_{i,\infty} - A^{u}_{j,\infty}$. In the realizations that lie below the diagonal, public beliefs differ more when the priors are observable. Hence the figure demonstrates that making priors observable may lead to greater disagreement in many cases.

monotone likelihood ratio property (Dixit and Weibull, 2007). Similarly, optimal behavior with bounded memory can give rise to increased belief polarization even when individuals receive the same signals (Wilson, 2003).

[6] The figure is based on 500 realizations of type profiles for parameter values $\sigma^2 = \tau^2 = 1$, $\bar{\mu}_i = 3$, and $\bar{\mu}_j = 0$.

While observability of priors can result in greater public bias for particular type realizations, observability always lowers the ex-ante expected value of public bias, $E\left[A_{i,\infty}-A_{j,\infty}\right]$. To see this, note that when priors are observable, by (7), the expected bias in public beliefs is

$$E[A^{ck}_{i,\infty}-A^{ck}_{j,\infty}]=\frac{\tau^2}{2+\tau^2}\left(\bar{\mu}_i-\bar{\mu}_j\right).$$

On the other hand, when the priors are not observable, by (9), the expected public bias is

$$E[A^{u}_{i,\infty}-A^{u}_{j,\infty}]=\frac{\tau^2\left(1+\sigma^2\tau^2\right)}{1+\left(1+\sigma^2\tau^2\right)\left(1+\tau^2\right)}\left(\bar{\mu}_i-\bar{\mu}_j\right).$$

If $\bar{\mu}_i=\bar{\mu}_j$ then the expected public bias is zero in both cases. If $\bar{\mu}_i>\bar{\mu}_j$, however, then $\sigma^2>0$ implies

$$E[A^{u}_{i,\infty}-A^{u}_{j,\infty}]>E[A^{ck}_{i,\infty}-A^{ck}_{j,\infty}]>0.$$

That is, the expected public bias is higher when priors are not observable than when they are observable. This is intuitive because unobservability of priors impedes the full aggregation of the distributed information through deliberation. This result is useful in comparing the difference between the average opinions of various groups. For example, it implies that differences across groups in beliefs about the incidence of police brutality or racial profiling would narrow on average if members of each group were to observe each other's priors and therefore understand how their information is incorporated into beliefs. We return to this point in Section 6.

3.4 Unobservable Correlated Priors

Under the assumption that the priors are uncorrelated, we have so far illustrated that unobservability of priors may impede the aggregation of distributed information through deliberation and affect the amount of public disagreement. We now show that when priors are correlated, all distributed information is aggregated and hence the observability of priors has no effect on public beliefs.

Assume that μ_i and μ_j are correlated:

$$\begin{pmatrix}\mu_i\\ \mu_j\end{pmatrix}\sim N\left(\begin{pmatrix}\bar{\mu}_i\\ \bar{\mu}_j\end{pmatrix},\sigma^2\begin{bmatrix}1&\rho\\ \rho&1\end{bmatrix}\right),$$

where $\rho\neq 0$. Observing μ_i, i believes that μ_j is distributed normally with mean

$$E_i[\mu_j|\mu_i]=\bar{\mu}_j+\rho\left(\mu_i-\bar{\mu}_i\right)$$

and variance

$$Var_i\left(\mu_j|\mu_i\right)=\sigma^2\left(1-\rho^2\right).$$

That is, $E_i[\mu_j|\mu_i]$ is a one-to-one function of μ_i. As before, we have $A^u_{i,1}=A_{i,1}$ and $A^u_{j,1}=A_{j,1}$. Now, for i, the announcement $A^u_{j,1}$ of j in the first round yields an additional noisy signal

$$\left(1+\tau^2\right)A^u_{j,1}-\tau^2E_i[\mu_j|\mu_i]=\tau^2\left(\mu_j-E_i[\mu_j|\mu_i]\right)+x_j=\theta+\tau^2\left(\mu_j-E_i[\mu_j|\mu_i]\right)+\varepsilon_j. \qquad (10)$$

The additive noise $\tau^2\left(\mu_j - E_i[\mu_j|\mu_i]\right) + \varepsilon_j$ has mean 0 and variance $\sigma^2\left(1-\rho^2\right)\tau^4 + \tau^2$. Updating his belief, in the second round i announces

$$A^u_{i,2} = KA^u_{i,1} + LA^u_{j,1} - \alpha L E_i[\mu_j|\mu_i],$$

where K and L are known strictly positive constants.[7] The crucial observation here is that $A^u_{i,2}$ is strictly decreasing in $E_i[\mu_j|\mu_i]$, which is i's expectation of j's prior once i has observed his own prior. Player j, having observed $A^u_{i,1}$ and $A^u_{j,1}$ from the previous round, can therefore use $A^u_{i,2}$ to deduce that

$$E_i[\mu_j|\mu_i] = (\alpha L)^{-1}\left(KA^u_{i,1} + LA^u_{j,1} - A^u_{i,2}\right).$$

Moreover, since $E_i[\mu_j|\mu_i] = \bar{\mu}_j + \rho\left(\mu_i - \bar{\mu}_i\right)$ and $\rho \neq 0$, there is a one-to-one mapping between μ_i and $E_i[\mu_j|\mu_i]$. Hence j correctly infers that

$$\mu_i = \bar{\mu}_i + \rho^{-1}\left((\alpha L)^{-1}\left(KA^u_{i,1} + LA^u_{j,1} - A^u_{i,2}\right) - \bar{\mu}_j\right).$$

That is, at the end of second round, all prior beliefs are revealed, and all signals can be inferred. The announcements in all subsequent rounds are therefore precisely the same as in the common knowledge case:

$$A^u_{i,3} = \cdots = A^u_{i,\infty} = A^{ck}_{i,\infty} = \frac{1+\tau^2}{2+\tau^2}\left(A_{i,1} + A_{j,1}\right) - \frac{\tau^2}{2+\tau^2}\mu_j.$$

Accordingly, when priors are correlated, both individuals can infer each other's prior beliefs from the manner in which they *react* to the initial announcements. All distributed information is therefore aggregated through communication, and the resulting public bias is fully attributable to differences in prior beliefs:

$$A^u_{i,\infty} - A^u_{j,\infty} = \frac{\tau^2}{2+\tau^2}\left(\mu_i - \mu_j\right).$$

We show in Section 5 that this is true under broad conditions. First, however, we consider the case of uncorrelated priors.

4 Public Biases

In this section we explore the impact of observability of priors on the degree of bias in public beliefs under the assumption that priors are independently and identically distributed:

Assumption 1. *The variance-covariance matrix for priors is $\Sigma = \sigma^2 I$.*

[7] One applies (1), starting from $\theta \sim N\left(A^u_{i,1}, \alpha\right)$ and using the signal in (10), to obtain

$$A^u_{i,2} = \frac{\sigma^2\left(1-\rho^2\right)\tau^4 + \tau^2}{\alpha + \sigma^2\left(1-\rho^2\right)\tau^4 + \tau^2}A^u_{i,1} + \frac{\alpha}{\alpha + \sigma^2\left(1-\rho^2\right)\tau^4 + \tau^2}\left(\left(1+\tau^2\right)A^u_{j,1} - \tau^2 E_i\left[\mu_j|\mu_i\right]\right).$$

The desired equation is obtained by letting K and L respectively denote the coefficients of $A^u_{i,1}$ and $A^u_{j,1}$.

That is, for all distinct pairs i and j, the priors μ_i and μ_j are independent (i.e. $\sigma_{ij} = 0$) and the variances of priors are equal (i.e. $\sigma_{ii} = \sigma^2$ for all i).

Consider two individuals, i and j. At the end of deliberation, j thinks that the expected value of θ is $A_{j,\infty}$. He also knows that i thinks that the expected value of θ is $A_{i,\infty}$. Therefore, j thinks that i overestimates θ by an amount $A_{i,\infty} - A_{j,\infty}$. This leads to our notion of public bias.

Definition 1. For any $i, j \in N$, the public bias of i relative to j is $A_{i,\infty} - A_{j,\infty}$.

Similarly, the *ex-ante bias* of i relative to j is $\bar{\mu}_i - \bar{\mu}_j$. The bias after i and j have observed their own priors but before they observe any information is $\mu_i - \mu_j$, which we call the *prior bias* of i relative to j. Note that the ex-ante bias is known to all players, and the public bias comes to be known through communication, but the prior bias may never be revealed.

We know from (6) that when priors are common knowledge, the only source of public bias is the difference in realized priors, $\mu_i - \mu_j$, which is scaled down through communication. The following lemma identifies the amount of public bias when priors are unobservable, generalizing the analysis of Section 3 to n individuals.

Lemma 1. *Under Assumption 1, for any i and j, the public bias of i relative to j under unobservable priors is*

$$A^u_{i,\infty} - A^u_{j,\infty} = \frac{\tau^2}{\gamma}\left(\bar{\mu}_i - \bar{\mu}_j\right) + \frac{\tau^4\sigma^2}{\gamma}\left(\mu_i - \mu_j\right) + \frac{\tau^2\sigma^2}{\gamma}\left(\varepsilon_i - \varepsilon_j\right), \tag{11}$$

where $\gamma = \left(1+\tau^2\right)\left(1+\tau^2\sigma^2\right) + n - 1$.

Under unobservable priors, public bias has three sources: ex-ante bias $(\bar{\mu}_i - \bar{\mu}_j)$, prior bias $(\mu_i - \mu_j)$, and informational difference $(\varepsilon_i - \varepsilon_j)$. The informational difference contributes to public bias because unobservability of priors impedes the full aggregation of information. Ex-ante bias affects public bias because, without full aggregation, individuals use ex-ante information on priors to estimate the information of others.

By Lemma 1, the magnitude of public biases does not depend on θ. Hence all individuals agree on the distribution of these biases (although they disagree on the distribution of public beliefs). Our next result establishes that, if the priors are drawn from distinct distributions, the expected bias is necessarily larger under unobservable priors. (The expected bias is always zero when priors are drawn from the same distribution.)

Proposition 1. *Under Assumption 1, if $\bar{\mu}_i > \bar{\mu}_j$, then*

$$E[A^u_{i,\infty} - A^u_{j,\infty}] > E[A^{ck}_{i,\infty} - A^{ck}_{j,\infty}] > 0.$$

Consider two individuals i and j. Suppose that $\bar{\mu}_i > \bar{\mu}_j$ so that at the ex-ante stage i overestimates θ relative to j, although the actual prior μ_i of i may or may not turn out to be larger than μ_j. After each player k forms his prior and receives his information, all individuals deliberate, communicating their beliefs as in our model. At the end of the deliberation, their beliefs become public. Proposition 1 establishes that all individuals expect that, at the end of the deliberation, i overestimates θ less vis-a-vis j when priors are observable. That is, $E_k[A^{ck}_{i,\infty} - A^{ck}_{j,\infty}] < E_k[A^u_{i,\infty} - A^u_{j,\infty}]$

according to each $k \in N$. Therefore, making priors observable decreases public biases on average. This suggests that social integration, interpreted as an increased understanding of the manner in which other people think, should result in lower levels of public disagreement on average. We explore these issues further in Section 6.

We conclude this section with a discussion of the manner in which increases in population size affect the aggregation of distributed information. When information is distributed among a large number of individuals, unobservability of priors becomes detrimental for communication, so much so that the bias at the end of the deliberation process is approximately the same as the bias before deliberation begins. Towards establishing this, recall from (4) that the distributed information in society has variance $\bar{\tau}^2 = \tau^2/n$. If one fixes τ and varies n, as n gets large, the distributed information becomes very precise. Consequently, the individuals approximately learn θ from each other. In order to disentangle the effect of group size n from the effect of the information available to the group, we now fix the precision of the distributed information and let n vary.

In particular, consider a family of models $\left(\tau_n^2, \sigma^2, \bar{\boldsymbol{\mu}}_n, n\right)$, indexed by the number of individuals n, where $\bar{\boldsymbol{\mu}}_n = (\bar{\mu}_1, \ldots, \bar{\mu}_n)$ is the vector of means for the priors; $\mu_i \sim N\left(\bar{\mu}_i, \sigma^2\right)$ for each $i \leq n$. We assume that the variance τ_n^2/n approaches some positive value $\bar{\tau}^2$ as $n \to \infty$. (A special case of this arises if the variance of distributed information is independent of n, i.e., the total amount of information is fixed. In that case, $\tau_n^2 = n\bar{\tau}^2$ for some fixed $\bar{\tau} > 0$.) For any distinct individuals i and j, and any pair of realized priors μ_i and μ_j, this family of models defines a sequence of random variables $(A_{i,\infty}^u - A_{j,\infty}^u)_n$. Our next result shows that under unobservable priors, as the number of individuals n becomes large, this sequence of random variables converges in distribution to $\mu_i - \mu_j$.

Proposition 2. *Under Assumption 1, for any family $\left(\tau_n^2, \sigma^2, \bar{\boldsymbol{\mu}}_n, n\right)$ of models, any distinct individuals i and j, and any realized priors $\left(\mu_i, \mu_j\right)$,*

$$\left(A_{i,\infty}^u - A_{j,\infty}^u\right)_n \xrightarrow{\mathcal{D}} \mu_i - \mu_j.$$

Proof. By Lemma 1,

$$E\left[\left(A_{i,\infty}^u - A_{j,\infty}^u\right)_n \middle| \mu_i, \mu_j\right] = \tau_n^2 \sigma^2 \eta \left(\mu_i - \mu_j\right) + \eta \left(\bar{\mu}_i - \bar{\mu}_j\right).$$

where

$$\eta = \frac{\tau_n^2/n}{n\left(\tau_n^2/n\right)^2 \sigma^2 + \left(\tau_n^2/n\right)\left(1+\sigma^2\right) + 1}.$$

As $n \to \infty$ and $\tau_n^2/n \to \bar{\tau}^2 > 0$, η goes to 0, while $\tau_n^2\sigma^2\eta$ goes to 1. Hence

$$\lim_{n\to\infty} E\left[\left(A_{i,\infty}^u - A_{j,\infty}^u\right)_n \middle| \mu_i, \mu_j\right] = \mu_i - \mu_j.$$

To complete the proof we need to show that the variance of $(A_{i,\infty}^u - A_{j,\infty}^u)_n$ goes to 0 as $n \to \infty$. Since μ_i and μ_j are given, from Lemma 1 we have

$$Var\left[\left(A_{i,\infty}^u - A_{j,\infty}^u\right)_n \middle| \mu_i, \mu_j\right] = 2\tau_n^2\sigma^4\eta^2 = 2\tau_n^2\sigma^2\eta\left(\sigma^2\eta\right).$$

Since η goes to 0 and $\tau_n^2\sigma^2\eta$ goes to 1 as $n \to \infty$, $\lim_{n\to\infty} Var[(A_{i,\infty}^u - A_{j,\infty}^u)_n \,|\mu_i, \mu_j] = 0$. □

Hence, when the number n of individuals is large, the public bias of i relative to j is approximately equal to the prior bias of i relative to j. In the limit, all distributed information, no matter how precise, is entirely dissipated.[8]

Under observable priors, individuals use all distributed information efficiently. Hence, their public beliefs and the bias in those beliefs do not depend on how information is distributed. When priors are unobservable, however, even if individuals have very precise information as a group and announce their beliefs sincerely, they cannot communicate any significant information: at the end of the deliberation process, their beliefs are as they were at the outset. The intuition is that each individual has such a small amount of information that their announcements reveal little more than their priors. Recall from (6) that

$$A^{ck}_{i,\infty} - A^{ck}_{j,\infty} = \frac{\bar{\tau}^2}{1+\bar{\tau}^2}\left(\mu_i - \mu_j\right)$$

so the difference in beliefs under observable priors is independent of n holding $\bar{\tau}$ fixed. An immediate implication of Proposition 2 is therefore the following: as the population size becomes large, so that a given amount of information is distributed among an increasingly large number of individuals, public bias under unobservability is greater *not only in expectation but also for almost all type realizations*. This is illustrated in Figure 2, which repeats the exercise depicted in Figure 1 but for three different values of group size. As n gets large, type realizations for which observability results in greater public bias (which lie below the diagonal) become increasingly rare.[9]

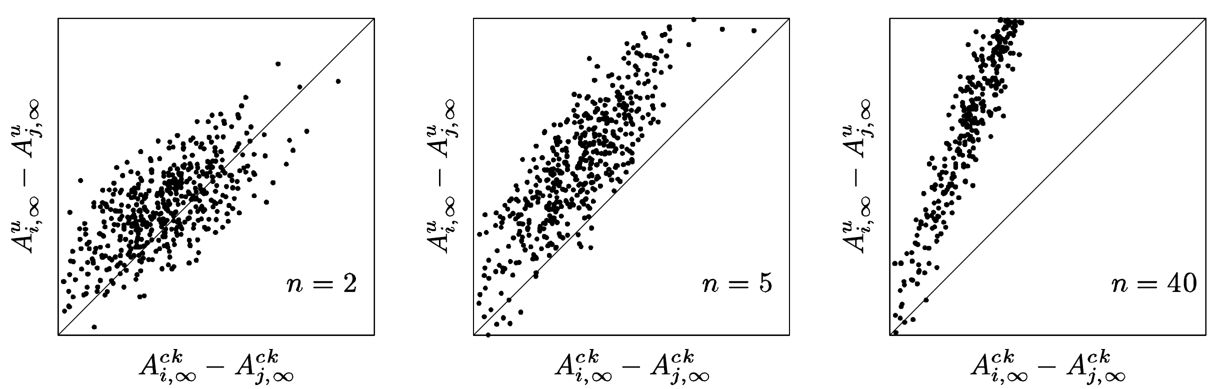

Figure 2. Public Bias with Observable and Unobsevable Priors for various n.

Note that for large n, Proposition 2 and (6) imply that

$$A^{u}_{i,\infty} - A^{u}_{j,\infty} = \frac{1+\bar{\tau}^2}{\bar{\tau}^2}\left(A^{ck}_{i,\infty} - A^{ck}_{j,\infty}\right).$$

[8] Note that since $\mu_i - \mu_j$ is a constant (conditional on the realized priors μ_i and μ_j), convergence in distribution implies convergence in probability. Hence we also have plim $A^{u}_{i,\infty} - A^{u}_{j,\infty} = \mu_i - \mu_j$.

[9] Each plot is based on 500 realizations of type profiles for parameter values $\sigma^2 = 1$, $\bar{\mu}_i = 3$, $\bar{\mu}_j = 0$, and $\tau_n^2/n = \bar{\tau}^2 = 1/2$ for all n. Only type realizations at which $\mu_i \geq \mu_j$ (and public biases lie in the positive quadrant) are shown. For realizations at which $\mu_i < \mu_j$ our results imply that as n gets large, public biases will lie below the diagonal in the negative quadrant.

Hence, the public bias under unobservability is a linear function of the public bias under observability, with slope greater than 1. This is illustrated for the case of $n = 40$ in the right panel of Figure 2.

5 Aggregation of Distributed Information

We now turn to the general model with unobservable priors, and provide a near-characterization of the cases in which the private information of an individual is revealed through deliberation. We show that, roughly speaking, if an individual's prior is correlated with the prior of any other individual, his private information is revealed by the end of the second round; otherwise his information is never revealed. Hence, except for certain knife-edge cases (as in the example of independent priors), the process of sequential belief announcements leads to the aggregation of all distributed information.

The idea that an individual's private information is revealed through communication is formalized as follows.

Definition 2. We say that the private information of individual i is *revealed by (the end of) round* k if (μ_i, x_i) is measurable with respect to $\{A_{j,m}\}_{j \in N, m \leq k}$. If the private information of individual i is not revealed by round k for any k, we say that his private information is *never revealed*.

That is, the private information of i is revealed by the end of round k if, by observing all announcements up to and including those in round k, one can compute his prior belief μ_i and signal x_i. In that case, his private information will be common knowledge at any round $m > k$:

$$A_{j,m} = E_j \left[\theta \mid \mu_i, x_i, \mu_j, x_j, \{A_{i',l}\}_{i' \in N \setminus \{i,j\}, l \leq m} \right] \qquad (\forall j \in N).$$

To present our characterization, we introduce the following notation. For any $i \in N$, we define column vectors $\mu_{-i} = (\mu_j)_{j \neq i}$ and $\sigma_{-i,i} = (\sigma_{j,i})_{j \neq i}$ and write $\Sigma_{-i,-i} = (\sigma_{j,k})_{j \neq i, k \neq i}$ for the variance covariance matrix of μ_{-i}. We write $1_{k \times l}$ for the $k \times l$-dimensional matrix with entries 1 and I for the identity matrix. Finally, we define the row vector M_i as follows:

$$M_i = 1_{1 \times n-1} \left(\alpha 1_{n-1 \times n-1} + \tau^2 I + \tau^4 \left(\Sigma_{-i,-i} - \sigma_{ii}^{-1} \sigma_{-i,i} \sigma_{-i,i}^{\top} \right) \right)^{-1}.$$

Note that M_i depends only on the primitives of the model and is therefore independent of all type realizations. The next definition provides the terminology of the characterization.

Definition 3. We say that i is *isolated* if $\sigma_{-i,i} = 0$. We say that i is *regular under* (τ^2, Σ) if $M_i \sigma_{-i,i} \neq 0$. We say that (τ^2, Σ) is *regular* if every i is regular under (τ^2, Σ).

Note that i is isolated if and only if μ_i is independent of all other priors μ_j. In this case, i cannot infer any information about the priors of others from his own prior. Consequently, others cannot learn about i's prior from the way he reacts to their announcements, and it is not possible to uncover all of his private information. The regularity condition rules out this case and some other knife-edge cases in which $M_i \sigma_{-i,i} = 0$ (without requiring that $\sigma_{-i,i} = 0$). Note that $M_i \sigma_{-i,i} = 0$

is a non-trivial linear equality restriction on the variances (τ^2, Σ), and hence is satisfied only on a lower-dimensional subspace of the space of all variances (τ^2, Σ). In particular, the set of regular parameters (τ^2, Σ) has full Lebesgue measure and is open and dense.

Our characterization establishes that whether the private information of an individual is revealed depends on whether he is regular or isolated.

Proposition 3. *Assume that the priors are not observable. If i is regular, then his private information is revealed by the end of round 2. Conversely, if i is isolated, then his private information is never revealed.*

An immediate implication of this is:

Corollary 1. *If (τ^2, Σ) is regular, then all private information is revealed by the end of round 2.*

This result establishes the *irrelevance of observability*: public beliefs under unobservable priors are identical to public beliefs under common knowledge of priors as long as (τ^2, Σ) is regular. All information is aggregated no matter how little individuals know about each other's way of thinking. Moreover, as in the two person case, this process requires just two rounds of communication.

In order to prove Proposition 3, in the Appendix, we compute the announcements (see Lemma 2). After the first round, the announcement of an individual i is an affine function of the first round announcements of all individuals, the priors of the individuals whose information has been revealed, and the prior μ_i of i himself. In the second round announcement, the coefficient of μ_i is proportional to $M_i\sigma_{-i,i}$. Hence, when $M_i\sigma_{-i,i} \neq 0$, other individuals can compute μ_i using the publicly available information and $A_{i,2}$. In that case, the private information of i is revealed by the end of the second round. Moreover, in any round after the first, the coefficient of μ_i is proportional to $\sigma_{-i,i}$. When $\sigma_{-i,i} = 0$, the announcement of i does not contain any new information because it is a function of publicly available information, namely the first round announcements and the priors that have already been revealed. In that case, i's private information is never revealed.

Since all private information is revealed by the end of the second round when (τ^2, Σ) is regular, the difference in public beliefs with unobservable priors is identical to the difference with observable priors: $A^u_{i,k} = A^{ck}_{i,k}$ for all i and $k \geq 3$. Using (6), we therefore have

$$A^u_{i,\infty} - A^u_{j,\infty} = A^{ck}_{i,\infty} - A^{ck}_{j,\infty} = \frac{\tau^2}{n + \tau^2}\left(\mu_i - \mu_j\right) = \frac{\bar{\tau}^2}{1 + \bar{\tau}^2}\left(\mu_i - \mu_j\right).$$

That is, under the regularity assumption, regardless of whether priors are observable or unobservable, differences in public beliefs are due only to differences in priors, scaled down by a factor that depends on the precision $1/\bar{\tau}^2$ of the distributed information.

The regularity assumption is weaker than genericity and contains many interesting "non-generic" cases, as the following example illustrates.

Example 1. Take $N = B \cup W$ consisting of two groups $B = \{1, 2\}$ and $W = \{3, 4\}$. For each i, $\sigma_{ii} = \sigma^2$ and for all distinct individuals i and j, $\sigma_{ij} > 0$ if i and j are in the same group and $\sigma_{ij} = 0$ otherwise. That is, from his own prior, an individual can learn about the other individual's prior in his own group, but he cannot learn anything about the other group. Nevertheless, (τ^2, Σ) is regular. One can check that, for any $i \in N$,

$$M_i \sigma_{-i,i} \propto 1 + \tau^2 + \sigma^2\tau^2 + \sigma^2\tau^4 + \sigma^2\tau^2\rho + \sigma^2\tau^4\rho \neq 0.$$

This example illustrates that even in a segregated society with no correlation in priors across groups, all distributed information is incorporated into public beliefs.

Proposition 3 implies that public beliefs are discontinuous with respect to the correlation of priors. When the priors are correlated, no matter how small the correlation may be, public beliefs incorporate all private information. When priors are independent, however, a substantial amount of private information remains private. This is true even for the third round announcements. The discontinuity stems from our assumption that individuals can communicate their beliefs precisely. In reality, individuals have only noisy information about the beliefs of others. For example, public polls reveal only partial information about the beliefs of a few randomly selected individuals. With imperfect observability, beliefs at any round will be continuous with respect to the correlation parameter. Accordingly, for sufficiently low levels of correlation between priors, a substantial amount of private information will remain uncommunicated at each round.

Full aggregation after only two rounds is an artifact of the two-dimensional model we use for tractability. Geanakoplos and Polemarchakis (1982) show that even under the common prior assumption beliefs may take arbitrarily long to stabilize. Hence the complete aggregation of distributed information could take an arbitrarily large number of communication rounds in a more general setting. Accordingly, we view Proposition 3 to be demonstrating that in the long run all information is aggregated under correlated priors. That is, we interpret third round announcements as corresponding to the long run.

Complete aggregation of distributed information in the long run relies on the assumption that all individuals have high levels of statistical sophistication. Not only are they able to make rational inferences based on the initial beliefs of others, they are also able to make rational inferences based on the manner in which others adjust their beliefs after hearing each successive round of announcements. This requires that individuals assume that beliefs are as described in the model, and assume that all individuals assume that beliefs are as described in the model, and update their beliefs accordingly... up to high orders. When such strong assumptions fail, individuals may fail to aggregate distributed information fully, and long-run behavior may resemble the case of independent priors, where individuals do not make inferences based on the manner in which others react to information.

6 Social Structure

As an illustration of the theory developed in the previous sections, we now analyze the amount of bias between two groups under three alternative social structures, which we call fragmentation, integration, and segregation. Fragmentation corresponds to a structure in which no individual observes the prior of any other. Under integration, each individual observes the prior of every other individual. And under segregation, each individual observes the priors of all those belonging to the same group, but none of the priors of those in the other group.

More formally, let $N = B \cup W$, where B and W are disjoint sets with $n_b \geq 2$ and $n_w \geq 2$ members, respectively. We maintain the assumption that $\Sigma = \sigma^2 I$, so priors are independently distributed, and we assume that for some $\bar{\mu}_b > \bar{\mu}_w$,

$$\bar{\mu}_i = \bar{\mu}_b \text{ and } \bar{\mu}_j = \bar{\mu}_w \qquad (\forall i \in B, j \in W).$$

That is, ex-ante, members of B overestimate θ relative to members of W.[10] An individual member of B, of course, may turn out to have a higher expectation than an individual member of W once each observes his own prior. We assume that opinions are communicated by successive belief announcement as before and that all announcements are observable.[11] Define the average opinion within each group in period k as follows:

$$\hat{A}_{b,k} = \frac{1}{n_b} \sum_{i \in B} A_{i,k} \text{ and } \hat{A}_{w,k} = \frac{1}{n_w} \sum_{j \in W} A_{j,k}.$$

We use the same superscript to denote other within-group averages as well, so $\hat{\mu}_b = \frac{1}{n_b} \sum_{i \in B} \mu_i$, $\hat{\varepsilon}_w = \frac{1}{n_w} \sum_{j \in W} \varepsilon_j$, etc. We are interested in the extent to which average opinion in B exceeds that in W at any given round k, defined as follows:

$$\beta_k \equiv \hat{A}_{b,k} - \hat{A}_{w,k}.$$

We let β_k^F, β_k^I and β_k^S denote the values of this difference under fragmentation, integration and segregation respectively.

For reasons discussed at the end of Section 5, we regard $k = 2$ as the medium run and $k = 3$ as the long run. From the previous section, recall that in both fragmented and integrated societies, $A_{i,k} = A_{i,2}$ for all $k \geq 2$, and the distinction between the medium and long run is not meaningful. However, as we show below, the distinction is important under segregation, since individuals behave as in the correlated priors case (although the priors are in fact independent).

6.1 Fragmentation

In a fragmented society, individuals obtain information, form beliefs, communicate these beliefs to pollsters, and observe the aggregate belief distribution. No individual observes the prior belief of any other individual. Instead, he uses his prior belief about the thinking of the others in order to extract the information revealed in the polls. This is the case of unobservable priors.

[10] This assumption is without loss of generality even if the groups are of unequal size, because if $\bar{\mu}_b < \bar{\mu}_w$, then we can simply reverse the order on θ by considering $-\theta$. Simply put, we are measuring the biases in the direction that, ex ante, members of B overestimate with respect to the members of W.

[11] The results in this section hold without modification even if only the average announcement in each group is publicly observable. To see this, note from (5) that when priors are observable, the public belief of i depends only on his own prior and the aggregate signal. Similarly, when priors are unobservable, the public belief of i depends only on his own initial announcement and the aggregate announcement in the group from (37) in the proof of Lemma 1.

From Lemma 1, for any round $k \geq 2$, the difference in average opinions across groups is

$$\beta_k^F = \frac{\tau^2}{\gamma}\left(\bar{\mu}_b - \bar{\mu}_w\right) + \frac{\tau^4\sigma^2}{\gamma}\left(\hat{\mu}_b - \hat{\mu}_w\right) + \frac{\tau^2\sigma^2}{\gamma}\left(\hat{\varepsilon}_b - \hat{\varepsilon}_w\right). \tag{12}$$

Hence, the bias has three sources: the ex-ante bias between groups $\left(\bar{\mu}_b - \bar{\mu}_w\right)$, the average prior bias between groups $(\hat{\mu}_b - \hat{\mu}_w)$, and the average informational difference between groups $(\hat{\varepsilon}_b - \hat{\varepsilon}_w)$. Recalling the definition of γ from Lemma 1, the expected value of the between-group bias is therefore

$$E[\beta_k^F] = \frac{\tau^2\left(1+\tau^2\sigma^2\right)}{\left(1+\tau^2\right)\left(1+\tau^2\sigma^2\right)+n-1}\left(\bar{\mu}_b - \bar{\mu}_w\right). \tag{13}$$

6.2 Integration

In an integrated society, each individual observes the priors of every other individual. They communicate directly, understanding the manner in which information is incorporated into beliefs. This is the case of observable priors.

From (6), for any round $k \geq 2$, the difference in average opinions across groups is

$$\beta_k^I = \frac{\tau^2}{\tau^2+n}\left(\hat{\mu}_b - \hat{\mu}_w\right). \tag{14}$$

Hence, the difference across groups in average opinion is the difference between their respective average priors, scaled down by a factor that uses all of the distributed information efficiently. The expected value of this is

$$E\left[\beta_k^I\right] = \frac{\tau^2}{\tau^2+n}\left(\bar{\mu}_b - \bar{\mu}_w\right) \tag{15}$$

and hence, from Proposition 1,

$$E[\beta_k^F] > E[\beta_k^I].$$

6.3 Segregation

Now we consider a segregated society partitioned into two components, one for each group. Each component is like an integrated society that is closed to members of the other component; individuals in different groups receive information about each other only through opinion polls. Formally, we assume that the prior of an individual is observable to the members of his own group and unobservable to the members of other group. That is, for each $i \in B$ and $j \in W$, μ_i is common knowledge among B and μ_j is common knowledge among W.

Now, when any $i \in B$ observes the first round announcements of his own group, he extracts all of the relevant information that other members of B have, concluding correctly that

$$\hat{x}_b \equiv \frac{1}{n_b}\sum_{i\in B} x_i = \left(1+\tau^2\right)\hat{A}_{b,1} - \tau^2\hat{\mu}_b. \tag{16}$$

On the other hand, he can extract only limited information from the announcements of the other group. The only relevant information for him is $\left(1+\tau^2\right)\hat{A}_{w,1} = \hat{x}_w + \tau^2\hat{\mu}_w$, where he knows neither

$\hat{x}_w$ nor $\hat{\mu}_w$. Combining these two pieces of information, he updates his belief, and in the second round, he announces

$$A_{i,2}^S = c_b \left(\alpha_b \mu_i + (1-\alpha_b)\, \hat{x}_b\right) + (1-c_b)\left((1+\tau^2)\hat{A}_{w,1} - \tau^2 \bar{\mu}_w\right) \qquad (i \in B)$$

where

$$\alpha_b = \frac{\tau^2}{\tau^2 + n_b} \tag{17}$$

and

$$c_b = \frac{\left(\tau^2 + n_b\right)\left(1+\tau^2\sigma^2\right)}{\left(1+\tau^2\sigma^2\right)\left(\tau^2+n_b\right)+n_w}. \tag{18}$$

Hence the average opinion in B at this stage is

$$\hat{A}_{b,2}^S = c_b \left(\alpha_b \hat{\mu}_b + (1-\alpha_b)\, \hat{x}_b\right) + (1-c_b)\left((1+\tau^2)\hat{A}_{w,1} - \tau^2 \bar{\mu}_w\right). \tag{19}$$

It turns out that, together with the first round announcements, the second round announcements reveal all relevant information. To see this, consider any $j \in W$. From the average first round announcements of the other group, j deduces that $\left(1+\tau^2\right)\hat{A}_{b,1} = \hat{x}_b + \tau^2 \hat{\mu}_b$, and in the second round deduces (19). Since $n_b > 1$, j can solve these two independent linear equations, thereby computing $\hat{x}_b$ and $\hat{\mu}_b$. That is, j does not need to know how members of B think: knowing that members of B know how each other thinks, j can infer all relevant information from the manner in which members of the B react to each others' announcements. As a result, by the end of the second round, all distributed information is aggregated, and in the long run, segregated and integrated societies are identical:

Proposition 4. *For each $i \in N$ and $k \geq 3$, $A_{i,k}^S = A_{i,k}^I$ and $\beta_k^S = \beta_k^I$.*

This illustrates the power of the argument behind Proposition 3. When some individuals have information about other individuals (through correlation in Proposition 3 and observation here), third parties can extract that information from the manner in which these individuals react to each other's announcements.

We now turn to the medium-run beliefs in a segregated society. From (16) and (19), we obtain

$$\hat{A}_{b,2}^S = (1-\alpha_b c_b)\,\theta + \alpha_b c_b \hat{\mu}_b + (1-c_b)\,\tau^2\left(\hat{\mu}_w - \bar{\mu}_w\right) + (1-\alpha_b)\, c_b \hat{\varepsilon}_b + (1-c_b)\,\hat{\varepsilon}_w. \tag{20}$$

Similarly,

$$\hat{A}_{w,2}^S = (1-\alpha_w c_w)\,\theta + \alpha_w c_w \hat{\mu}_w + (1-c_w)\,\tau^2\left(\hat{\mu}_b - \bar{\mu}_b\right) + (1-\alpha_w)\, c_w \hat{\varepsilon}_w + (1-c_w)\,\hat{\varepsilon}_b \tag{21}$$

where α_w and c_w are defined analogously to (17) and (18).

If $n_b = n_w$ then $\alpha_b c_b = \alpha_w c_w$. In that case, the medium-run bias, $\beta_2^S = \hat{A}_{b,2}^S - \hat{A}_{w,2}^S$, does not depend on θ, and all individuals have the same expectation:

$$E\left[\beta_2^S\right] = \frac{\tau^2\left(1+\tau^2\sigma^2\right)}{\left(1+\tau^2\sigma^2\right)\left(\tau^2+n/2\right)+n/2}\left(\bar{\mu}_b - \bar{\mu}_w\right).$$

It is easily verified that for any $n > 2$,

$$E[\beta_2^I] < E[\beta_2^S] < E[\beta_2^F].$$

That is, when groups are of equal size, they agree about the value of the medium-run bias under all three information structures, and the bias is greatest under fragmentation, least under integration, and intermediate under segregation.

When groups are of unequal size, however, the medium-run bias does depend on θ, and hence the members of different groups will have different expectations about it. Our next result establishes that, in a segregated society, ex-ante, members of a minority group will expect a smaller medium-run bias than the members of a majority group. Despite this, it further establishes that they all agree that the expected medium-run bias under segregation is higher than that under integration, and lower than that under fragmentation:

Proposition 5. *If $n_b < n_w$, then, for all $i \in B$ and $j \in W$,*

$$E[\beta_2^I] < E_i[\beta_2^S] < E_j[\beta_2^S] < E[\beta_2^F]$$

To gain some intuition for the finding that minorities expect lower levels of medium-run bias, consider a highly skewed population with a very large majority group and a very small minority. Then the minority expects the majority to approximately learn θ, and hence to converge to a belief that is close to the minority group prior. The minority therefore expects the initial bias to diminish substantially. In contrast, the majority expects the minority to learn very little, and hence to maintain beliefs that are distant from majority group priors, with little narrowing of the initial bias. Roughly speaking, the minority expects the majority to come around to their way of thinking, while the majority expects no such convergence.

In summary, expected biases are always highest under fragmentation. Expected biases are higher under segregation than under integration in the medium run, but the two social structures are identical in the long run. This is intuitive, since individuals have the least ability to process information under fragmentation and the greatest ability to process information under integration.

6.4 Large Societies

We have so far compared the expected value of biases under three social structures for arbitrary values of the population size n. In large societies idiosyncratic differences cancel each other out and we can compare the magnitudes of *actual* biases under various social structures state by state. Doing so reveals that our analysis of expectations misses an interesting and potentially disturbing fact about medium-run beliefs: segregation puts minorities at a disadvantage in processing public information and consequently results in biases even when groups are formed from ex-ante identical individuals.

In order to compare biases in large societies (as in Section 4), we consider a family of models indexed by n, such that

$$\text{as } n \to \infty,\ \tau^2/n \to \bar{\tau}^2 \text{ and } n_b/n \to r \tag{22}$$

for some $\bar{\tau}^2 > 0$ and $r \in \left(0, \frac{1}{2}\right)$. That is, we adopt the convention that B is the minority group. In a large fragmented society, by (12), the bias is approximately as great as the ex-ante bias:

$$\lim_{n\to\infty} \beta_k^F = \bar{\mu}_b - \bar{\mu}_w \qquad \text{almost surely, for all } k \geq 2. \tag{23}$$

By (14), in a large integrated society, the bias is smaller, to a degree that depends on the precision of the distributed information:

$$\lim_{n\to\infty} \beta_k^I = \frac{\bar{\tau}^2}{\bar{\tau}^2 + 1} \left(\bar{\mu}_b - \bar{\mu}_w\right) \qquad \text{almost surely, for all } k \geq 2. \tag{24}$$

In a large segregated society, the bias is identical to that under integration in the long run, as we have seen above:

$$\lim_{n\to\infty} \beta_k^S = \frac{\bar{\tau}^2}{\bar{\tau}^2 + 1} \left(\bar{\mu}_b - \bar{\mu}_w\right) \qquad \text{almost surely, for all } k \geq 3.$$

In the long run, both segregated and integrated societies use all available information efficiently.

In the medium run, under segregation, information is not fully aggregated. This does not, however, mean that the magnitude of the bias lies strictly between the corresponding magnitudes under fragmentation and integration respectively. To see this, note from (20) and (21) that average group beliefs in the medium run are given by:

$$\lim_{n\to\infty} \hat{A}_{b,2}^S = \frac{\bar{\tau}^2}{\bar{\tau}^2 + r} \bar{\mu}_b + \frac{r}{\bar{\tau}^2 + r} \theta \tag{25}$$

$$\lim_{n\to\infty} \hat{A}_{w,2}^S = \frac{\bar{\tau}^2}{\bar{\tau}^2 + 1 - r} \bar{\mu}_w + \frac{1 - r}{\bar{\tau}^2 + 1 - r} \theta. \tag{26}$$

Notice that neither group processes information as efficiently as in an integrated society. In effect, a representative member of the minority group faces a noisy signal with variance $\bar{\tau}^2/r$, and a representative member of the majority group faces a noisy signal with variance $\bar{\tau}^2/\left(1 - r\right)$. Under integration, each individual obtains a noisy signal with variance $\bar{\tau}$, which is clearly smaller than both $\bar{\tau}^2/r$ and $\bar{\tau}^2/\left(1 - r\right)$. Furthermore, under segregation, minorities are disadvantaged in processing public information, since $\bar{\tau}^2/r > \bar{\tau}^2/\left(1 - r\right)$. As a result, in the medium run, the majority belief puts greater weight on the true state (and less weight on the prior) when compared with the minority group belief. This disadvantage becomes more pronounced as group sizes become more unequal.[12]

Note that the medium-run bias under segregation depends on θ:

$$\lim_{n\to\infty} \beta_2^S = \frac{\bar{\tau}^2}{\bar{\tau}^2 + r} \bar{\mu}_b - \frac{\bar{\tau}^2}{\bar{\tau}^2 + 1 - r} \bar{\mu}_w - \left(\frac{\bar{\tau}^2}{\bar{\tau}^2 + r} - \frac{\bar{\tau}^2}{\bar{\tau}^2 + 1 - r} \right) \theta \qquad \text{almost surely.} \tag{27}$$

Because of this dependence, the bias can take any value. In particular, in the medium run, the difference in beliefs under segregation may *increase* (relative to the ex-ante belief difference) and

[12] Individuals belonging to a minority within any population tend to have a smaller number of affiliates in friendship networks (Currarini et al., 2009), which should reinforce this effect. On the other hand, segregation itself tends to be endogenously increasing in the size of the minority group (Sethi and Somanathan, 2004), which suggests that the extent of public disagreement may not vary monotonically with the size of the minority group.

therefore exceed the difference under fragmentation. This will occur if θ turns out to be very different from ex-ante expectations of it.

An interesting special case arises when the groups have identical ex-ante beliefs: $\bar{\mu}_b = \bar{\mu}_w = \bar{\mu} \neq \theta$ for some $\bar{\mu}$. That is, the two groups start out with identical priors, and the true state happens to be different from the priors. Then, from (23) and (24), the medium and long run bias is negligible under fragmentation and integration: $\lim_{n\to\infty} \beta_k^F = \lim_{n\to\infty} \beta_k^I = 0$ almost surely for $k \geq 2$. However, from (27), the medium-run bias under segregation is strictly positive:

$$\lim_{n\to\infty} \beta_2^S = \left(\frac{\overline{\tau}^2}{\overline{\tau}^2 + r} - \frac{\overline{\tau}^2}{\overline{\tau}^2 + 1 - r} \right) (\bar{\mu} - \theta) \qquad \text{almost surely.}$$

Furthermore, from (25-26), the majority group belief in the medium run is closer to the true state. These two facts may be stated as follows.

Proposition 6. *Suppose that $\bar{\mu}_b = \bar{\mu}_w \neq \theta$. Then,*

$$\lim_{n\to\infty} \left| \hat{A}_{b,2}^S - \theta \right| > \lim_{n\to\infty} \left| \hat{A}_{w,2}^S - \theta \right|$$

and

$$\lim_{n\to\infty} \left| \beta_2^S \right| > \lim_{n\to\infty} \left| \beta_k^F \right| = \lim_{n\to\infty} \left| \beta_k^I \right| = 0.$$

To summarize, when the two groups are composed of ex-ante identical individuals, the beliefs of the majority group are more closely aligned with reality than are the beliefs of the minority group in the medium run. Also, the level of medium-run bias is greater under segregation than under either integration or fragmentation. The former result arises directly from the fact that majority group members have an advantage in the interpretation of public information. The latter result arises because segregation tends to homogenize beliefs *within* groups, which has the effect of creating belief heterogeneity *across* groups. This effect does not arise under either fragmentation or integration.

7 Conclusions

If a group of individuals share a common prior and are commonly known to be Bayesian rational (in the sense that each member of the group forms beliefs using Bayes' rule according to the common prior) then public disagreement cannot arise. Accounting for such disagreement therefore requires a departure from one or both of these hypotheses. We have chosen here to explore the implications of heterogeneous priors, while maintaining stringent assumptions regarding Bayesian rationality. Two main results follow from this. First, we find that for generic values of the model's primitives, the extent of public disagreement is independent of whether or not priors are observable, and public beliefs involve the aggregation of all distributed information in the long run. Second, we find that when priors are uncorrelated, the expected value of public bias is lower in an integrated society than in a fragmented one. For large societies, a stronger result holds: public bias is greater in a fragmented society relative to an integrated one under almost all realizations of priors and information. This suggests that social integration (in the sense of better understanding of the priors of others) should result in diminished public disagreement, especially in large populations.

Our results depend on the ability of individuals to make highly sophisticated statistical inferences, based not only on the initial beliefs of others but also on the manner in which these beliefs are adjusted over time on the basis of earlier announcements. If cognitive limitations prevent individuals from making inferences based on the manner in which one person responds to another's announcement, then our medium-run analysis applies, and the expected value of bias across social groups depends systematically on the extent of social integration. Expected bias is smallest in integrated societies (where priors are observable both within and between social groups) and largest in fragmented societies (where priors are unobservable even within social groups). Intermediate levels of expected bias arise under segregation, when priors are observable within but not across groups. Hence integration both within and across social groups tends to reduce expected levels of public bias.

In large populations, realized biases may be greater under segregation than under either fragmentation or integration, as belief differences are compressed within groups but amplified across groups. Communication in segregated societies can cause initial biases to be amplified, and new biases to emerge where none previously existed. Despite the fact that all announcements are public and all signals equally precise, members of a minority group face a disadvantage in the interpretation of public information that results in beliefs that are less closely aligned with the true state. If majority group members (or outside observers) fail to appreciate this effect, they may regard the views of minorities as "bizarre" or "outlandish", attributing them to failures in reasoning rather than to structural factors such as the demographic composition and constraints on information exchange induced by the heterogeneity and unobservability of prior beliefs.

A Appendix—Proofs

A.1 Aggregation of Distributed Information

In this subsection, we prove Proposition 3. The proof requires the use of the following well-known formula. For any two random vectors X and Y, if

$$\begin{pmatrix} X \\ Y \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_X \\ \mu_Y \end{pmatrix}, \begin{pmatrix} \Sigma_X & \Sigma_{X,Y} \\ \Sigma_{Y,X} & \Sigma_Y \end{pmatrix}\right),$$

then conditional on Y, X is distributed with $N\left(E\left[X|Y\right], Var\left(X|Y\right)\right)$ where

$$\begin{aligned} E\left[X|Y\right] &= \mu_X + \Sigma_{X,Y}\Sigma_Y^{-1}\left(Y - \mu_Y\right) \\ Var\left(X|Y\right) &= \Sigma_X - \Sigma_{X,Y}\Sigma_Y^{-1}\Sigma_{Y,X}. \end{aligned} \tag{28}$$

We also need to introduce some more notation. For any subset $N' \subset N$, we use subscript N' to denote the column vector obtained by stacking up all the values for $j \in N'$. For example, we write $\mu_{N'} = (\mu_j)_{j\in N'}$, $A_{N',k} = \left(A_{j,k}\right)_{j\in N'}$, and $\sigma_{N',i} = \left(\sigma_{j,i}\right)_{j\in N'}$. For any subsets N' and N'' of N and any matrix $X = \left(x_{i,j}\right)_{i,j\in N}$, we write $X_{N',N''}$ for the submatrix with entries from N' and N'', i.e., $X_{N',N''} = \left(x_{i,j}\right)_{i\in N', j\in N''}$. We use subscript $-i$ instead of $N\backslash\{i\}$, e.g., $\mu_{-i} = (\mu_j)_{j\neq i}$ and $\Sigma_{-i,-i} = \left(\sigma_{j,k}\right)_{j\neq i, k\neq i}$. We write $1_{k\times l}$ for the $k \times l$-dimensional matrix with entries 1 and I for the identity matrix. We write

$$\begin{aligned} \tilde{\mu}_{-i} &\equiv E\left[\mu_{-i}|\mu_i\right] = \bar{\mu}_{-i} + \sigma_{ii}^{-1}\sigma_{-i,i}\left(\mu_i - \bar{\mu}_i\right) \\ \tilde{\Sigma}_{-i,-i} &\equiv Var\left(\mu_{-i}|\mu_i\right) = \Sigma_{-i,-i} - \sigma_{ii}^{-1}\sigma_{-i,i}\sigma_{-i,i}^{\top}. \end{aligned} \tag{29}$$

Using the definitions of R and H in Lemma 2 below, we also write

$$\begin{aligned} \hat{v} &= \tau^2/\left(\tau^2 + 1 + |R|\right), \\ \alpha &= \tau^2/\left(\tau^2 + 1\right), \\ M_R &= 1_{1\times|H|}\left(\hat{v}1_{|H|\times|H|} + \tau^2 I + \tau^4\tilde{\Sigma}_{H,H} - \tau^4\tilde{\Sigma}_{H,R}\tilde{\Sigma}_{R,R}^{-1}\tilde{\Sigma}_{R,H}\right)^{-1} \\ M_i &= 1_{1\times n-1}\left(\alpha 1_{n-1\times n-1} + \tau^2 I + \tau^4\Sigma_{-i,-i} - \tau^4\sigma_{ii}^{-1}\sigma_{-i,i}\sigma_{-i,i}^{\top}\right)^{-1}. \end{aligned} \tag{30}$$

We compute the announcements in the following lemma.

Lemma 2. *Assume that the priors are not observable. For any $i \in N$ and any round k, let $R \subseteq N\backslash\{i\}$ be the set of other individuals whose private information is revealed by the end of round $k-1$, and let $H = N\backslash\left(R \cup \{i\}\right)$. Then,*

$$\begin{aligned} A_{i,k}^u = {} & \frac{\tau^2+1}{\tau^2+1+|R|}\left(1 - \hat{v}M_R 1_{|H|\times 1}\right)\sum_{j\in R\cup\{i\}} A_{j,1} + \left(1+\tau^2\right)\hat{v}M_R A_{H,1} \\ & - \frac{\tau^2}{\tau^2+1+|R|}1_{1\times|R|}\mu_R - \tau^2\hat{v}M_R\left(\bar{\mu}_H - \tilde{\Sigma}_{H,R}\tilde{\Sigma}_{R,R}^{-1}\left(\mu_R - \bar{\mu}_R\right)\right) \\ & - \tau^2\sigma_{ii}^{-1}\hat{v}M_R\left(\sigma_{H,i} - \tilde{\Sigma}_{H,R}\tilde{\Sigma}_{R,R}^{-1}\sigma_{R,i}\right)\left(\mu_i - \bar{\mu}_i\right) \end{aligned} \tag{31}$$

when $R \neq \varnothing$ *and*

$$A^u_{i,k} = (1 - \alpha M_i 1_{n-1\times 1}) A_{i,1} + \tau^2 M_i A_{-i,1} - \tau^2 \alpha M_i \bar{\mu}_{-i} - \tau^2 \sigma_{ii}^{-1} \alpha M_i \sigma_{-i,i} (\mu_i - \bar{\mu}_i) \tag{32}$$

when $R = \varnothing$.

Proof. We will use mathematical induction on k. We first compute $A^u_{i,2}$, showing that the statement is true for $k = 2$. For each j, since $A_{j,1} = \alpha\mu_j + (1-\alpha) x_j$,

$$\left(1 + \tau^2\right) A_{j,1} = \theta + \varepsilon_j + \tau^2 \mu_j.$$

Hence,

$$E_i \left[\left(1 + \tau^2\right) A_{j,1} \mid \mu_i, x_i\right] = A_{i,1} + \tau^2 E_i[\mu_j|\mu_i].$$

Substituting (29) in this equality, we obtain

$$E_i \left[\left(1 + \tau^2\right) A_{-i,1} \mid \mu_i, x_i\right] = 1_{n-1\times 1} A_{i,1} + \tau^2 \bar{\mu}_{-i} + \tau^2 \sigma_{ii}^{-1} \sigma_{-i,i} (\mu_i - \bar{\mu}_i). \tag{33}$$

Now, the first round of announcements provides i a new vector $\left(1 + \tau^2\right) A_{-i,1} = \theta 1_{n-1\times 1} + \varepsilon_{-i} + \tau^2 \mu_{-i}$ of signals with additive normal noise. Notice that, conditional on (x_i, μ_i), the variance of $\theta 1_{n-1\times 1} + \varepsilon_{-i} + \tau^2 \mu_{-i}$ is

$$\alpha 1_{n-1\times n-1} + \tau^2 I + \tau^4 \left(\Sigma_{-i} - \sigma_{ii}^{-1} \sigma_{-i,i} \sigma_{-i,i}^\top\right).$$

Hence, updating his belief according to (28), in the second round i announces

$$\begin{aligned}
A^u_{i,2} &= E_i \left[\theta \mid \mu_i, x_i, \left(1 + \tau^2\right) A_{-i,1}\right] \\
&= A_{i,1} + \alpha M_i \left(\left(1 + \tau^2\right) A_{-i,1} - E_i \left[\left(1 + \tau^2\right) A_{-i,1} | \mu_i, x_i\right]\right) \\
&= (1 - \alpha M_i 1_{n-1\times 1}) A_{i,1} + \tau^2 M_i A_{-i,1} - \alpha \tau^2 M_i \bar{\mu}_{-i} - \tau^2 \sigma_{ii}^{-1} \alpha M_i \sigma_{-i,i} (\mu_i - \bar{\mu}_i)
\end{aligned} \tag{34}$$

where the second equality is by (28) and the definition of M_i, and the last equality is by (33). Now suppose that the proposition is true for rounds $k' \leq k - 1$ and for all j. Then, if

$$M_R \left(\sigma_{H,j} - \tilde{\Sigma}_{H,R} \tilde{\Sigma}_{R,R}^{-1} \sigma_{R,j}\right) = 0$$

for R defined for k' and j, no new information is revealed by the announcement $A^u_{j,k'}$ because it is measurable with respect to the public information at the end of round $k' - 1$. On the other hand, if

$$M_R \left(\sigma_{H,j} - \tilde{\Sigma}_{H,R} \tilde{\Sigma}_{R,R}^{-1} \sigma_{R,j}\right) \neq 0,$$

then we can solve for μ_j from (31) for k' and j. That is, either the private information of j is revealed by the end of round $k - 1$, i.e., $j \in R$, or i knows only that $A_{j,1} = \alpha\mu_j + (1-\alpha) x_j$. Now, if $R = \varnothing$, i has not learned any new information after the first round. In that case, $A^u_{i,k} = A^u_{i,2}$, and (32) is equivalent to (34). Now suppose that $R \neq \varnothing$. Individual i knows (μ_i, x_i), $\left(\mu_j, x_j\right)$ for $j \in R$ and that $A_{j,1} = \alpha\mu_j + (1-\alpha) x_j$ for $j \notin R$. We compute conditional distributions sequentially, first conditioning on (μ_i, x_i), then on (μ_R, x_R), and finally on $A_{H,1} = \alpha\mu_H + (1-\alpha) x_H$, i.e.,

$$\left(1 + \tau^2\right) A_{H,1} = 1_{|H|\times 1} \theta + \varepsilon_H + \tau^2 \mu_H. \tag{35}$$

Conditional on (μ_i, x_i), $\left(\theta, \mu_{-i}, \varepsilon_{-i}\right)$ are independently and normally distributed with $\theta \sim N\left(A_{i,1}, \alpha\right)$, $\mu_{-i} \sim N(\tilde{\mu}_{-i}, \tilde{\Sigma}_{-i,-i})$, and $\varepsilon_{-i} \sim N\left(0, \tau^2 I\right)$. Then, from (μ_R, x_R), he obtains a new signal $x_R = 1_{|R|\times 1}\theta + \varepsilon_R$ about θ and also potentially new information about μ_H from μ_R. Conditioning on $x_R = 1_{|R|\times 1}\theta + \varepsilon_R$, he updates his belief about θ to $N(\hat{\mu}_i, \hat{v})$ where

$$\begin{aligned}
\hat{\mu}_i &= \frac{\tau^2+1}{\tau^2+1+|R|} A_{i,1} + \frac{1}{\tau^2+1+|R|} 1_{1\times|R|} x_R \\
&= \frac{\tau^2+1}{\tau^2+1+|R|} \sum_{j \in R \cup \{i\}} A_{j,1} - \frac{\tau^2}{\tau^2+1+|R|} 1_{1\times|R|}\mu_R \\
\hat{v} &= \frac{\tau^2}{\tau^2+1+|R|}.
\end{aligned}$$

Conditioning on μ_R, he updates his belief about μ_H to $N(\hat{\mu}_H, \hat{\Sigma}_H)$ where

$$\begin{aligned}
\hat{\mu}_H &= \tilde{\mu}_H + \tilde{\Sigma}_{H,R}\tilde{\Sigma}_{R,R}^{-1}\left(\mu_R - \tilde{\mu}_R\right) \\
\hat{\Sigma}_H &= \tilde{\Sigma}_{H,H} - \tilde{\Sigma}_{H,R}\tilde{\Sigma}_{R,R}^{-1}\tilde{\Sigma}_{R,H}.
\end{aligned}$$

Now, i conditions on (35) starting from $\theta \sim N(\hat{\mu}_i, \hat{v})$. Given the conditionings so far, by (35),

$$\left(1+\tau^2\right) A_{H,1} \sim N\left(\hat{\mu}_i 1_{|H|\times 1} + \tau^2\hat{\mu}_H, \hat{v}1_{|H|\times|H|} + \tau^4\hat{\Sigma}_H + \tau^2 I\right).$$

Using (28), he therefore obtains

$$\begin{aligned}
A_{i,k} &= E\left[\theta | \mu_i, x_i, \mu_R, x_R, \left(1+\tau^2\right) A_{H,1} = 1_{|H|\times 1}\theta + \varepsilon_H + \tau^2\mu_H\right] \\
&= \hat{\mu}_i + \hat{v}1_{1\times|H|}\left(\hat{v}1_{|H|\times|H|} + \tau^4\hat{\Sigma}_H + \tau^2 I\right)^{-1}\left(\left(1+\tau^2\right)A_{H,1} - \hat{\mu}_i 1_{|H|\times 1} - \tau^2\hat{\mu}_H\right) \\
&= \left(1 - \hat{v}M_R 1_{|H|\times 1}\right)\hat{\mu}_i + \left(1+\tau^2\right)\hat{v}M_R A_{H,1} - \tau^2\hat{v}M_R\hat{\mu}_H \\
&= \frac{\tau^2+1}{\tau^2+1+|R|}\left(1 - \hat{v}M_R 1_{|H|\times 1}\right)\sum_{j\in R\cup\{i\}} A_{j,1} - \frac{\tau^2}{\tau^2+1+|R|}\left(1 - \hat{v}M_R 1_{|H|\times 1}\right)1_{1\times|R|}\mu_R \\
&\quad + \left(1+\tau^2\right)\hat{v}M_R A_{H,1} \\
&\quad - \tau^2\hat{v}M_R\left(\bar{\mu}_H + \sigma_{ii}^{-1}\sigma_{H,i}\left(\mu_i - \bar{\mu}_i\right) + \tilde{\Sigma}_{H,R}\tilde{\Sigma}_{R,R}^{-1}\left(\mu_R - \bar{\mu}_R - \sigma_{ii}^{-1}\sigma_{R,i}\left(\mu_i - \bar{\mu}_i\right)\right)\right),
\end{aligned}$$

where the second equality is by (28); the third is by arrangement of terms using the definition of M_R, and the last by substituting the values of $\hat{\mu}_i$ and $\hat{\mu}_H$. By rearranging terms, we obtain the equality in the proposition. □

Using Lemma 2, we can now prove Proposition 3.

Proof of Proposition 3. Assume first that i is regular, i.e., $M_i\sigma_{-i,i} \neq 0$. Then, since no individual's private information is revealed by the end of round 2, by (32),

$$\mu_i = \bar{\mu}_i + \frac{\left(1 - \alpha M_i 1_{n-1\times 1}\right)A_{i,1} + \tau^2 M_i A_{-i,1} - \tau^2\alpha M_i\bar{\mu}_{-i} - A_{i,2}}{\alpha\tau^2\sigma_{ii}^{-1}M_i\sigma_{-i,i}},$$

i.e., μ_i is measurable with respect to $A_{i,1}$, $A_{-i,1}$, and $A_{i,2}$. Moreover, since $A_{i,1} = \alpha\mu_i + (1-\alpha_i)\, x_i$, we can further compute that

$$x_i = \left(1+\tau^2\right) A_{i,1} - \tau^2 \left(\bar{\mu}_i + \frac{\left(1-\alpha M_i 1_{n-1\times 1}\right) A_{i,1} + \tau^2 M_i A_{-i,1} - \tau^2 \alpha M_i \bar{\mu}_{-i} - A_{i,2}}{\alpha \tau^2 \sigma_{ii}^{-1} M_i \sigma_{-i,i}}\right),$$

showing that x_i is measurable with respect to $A_{i,1}$, $A_{-i,1}$, and $A_{i,2}$. Therefore, the private information of i is revealed by round 2. Conversely, suppose that i is isolated, i.e., $\sigma_{-i,i} = 0$. (Note that, in that case, $\tilde{\mu}_{-i} = \mu_{-i}$ and $\tilde{\Sigma}_{-i,-i} = \Sigma_{-i,-i}$.) Hence, by Lemma 2, for any $k > 1$, if $R \neq \varnothing$, then the coefficient of μ_i is

$$-\tau^2 \sigma_{ii}^{-1} \hat{\upsilon} M_R \left(\sigma_{H,i} - \Sigma_{H,R} \Sigma_{R,R}^{-1} \sigma_{R,i}\right) = 0$$

because $\sigma_{H,i} = 0$ and $\sigma_{R,i} = 0$. If $R = \varnothing$, the coefficient is again $\tau^2 \sigma_{ii}^{-1} \alpha M_i \sigma_{-i,i} = 0$. Thus, $A_{i,k}$ is measurable with respect to the information at the end of round $k-1$, revealing no new information. On the other hand, since $\sigma_{-i,i} = 0$, (x_R, μ_R) does not provide any information about μ_i, either. It only reduces the variance of x_i without revealing it. Hence, the private information of i is not revealed at any round. □

A.2 Public Bias

Proof of Lemma 1. By Proposition 3, since the priors are independent, no information is revealed. Hence, by Lemma 2, $A^u_{i,\infty} = A^u_{i,2}$, and $A^u_{i,2}$ satisfies (32). To compute $A^u_{i,2}$ from (32), first define

$$\varphi = \left(1+\tau^2\right)\left(1+\tau^2\sigma^2\right)$$

and note that

$$\begin{aligned}
M_i &= 1_{1\times n-1}\left(\alpha 1_{n-1\times n-1} + \left(\tau^2 + \tau^4\sigma^2\right) I\right)^{-1} \\
&= \alpha^{-1} 1_{1\times n-1}\left(1_{n-1\times n-1} + \varphi I\right)^{-1} \\
&= \frac{1}{\alpha\varphi\left(\varphi+n-1\right)} 1_{1\times n-1}\left(\left(\varphi+n-1\right) I - 1_{n-1\times n-1}\right) \\
&= \frac{1}{\alpha\left(\varphi+n-1\right)} 1_{1\times n-1}. \qquad (36)
\end{aligned}$$

Here, the first equality is obtained by substituting $\Sigma = \sigma^2 I$ in (30), and the second equality is by simple algebra. In the third equality, we invert the matrix $1_{n-1\times n-1} + \varphi I$. It can be easily verified that

$$\left(1_{n-1\times n-1} + \varphi I\right)^{-1} = \frac{1}{\varphi\left(\varphi+n-1\right)}\left(\left(\varphi+n-1\right) - 1_{n-1\times n-1} I\right),$$

yielding the third line. Finally, by adding up the rows of the matrix $\left(\left(\varphi+n-1\right) I - 1_{n-1\times n-1}\right)$, we obtain (36). Substituting (36) in (32), we then obtain

$$\begin{aligned}
A^u_{i,2} &= \left(1-\alpha M_i 1_{n-1\times 1}\right) A_{i,1} + \tau^2 M_i A_{-i,1} - \alpha\tau^2 M_i \bar{\mu}_{-i} \\
&= \left(1 - \frac{1_{1\times n-1} 1_{n-1\times 1}}{\varphi+n-1}\right) A_{i,1} + \frac{\tau^2}{\alpha\left(\varphi+n-1\right)} 1_{1\times n-1} A_{-i,1} - \frac{\tau^2}{\varphi+n-1} 1_{1\times n-1} \bar{\mu}_{-i} \\
&= \frac{1}{\varphi+n-1} A_{i,1} + \frac{1+\tau^2}{\varphi+n-1} \sum_{j\neq i} A_{j,1} - \frac{\tau^2}{\varphi+n-1} \sum_{j\neq i} A_{j,1} \bar{\mu}_j.
\end{aligned}$$

Here, the first equality is simply (32) for $\sigma_{-i,i} = 0$, and the second equality is just by the substitution of the value of M_i from (36). The last equality is by straightforward algebra. By adding and subtracting new terms with $A_{i,1}$ and μ_i, we obtain

$$A^u_{i,\infty} = A^u_{i,2} = \frac{\varphi - \left(1+\tau^2\right)}{\varphi + n - 1} A_{i,1} + \frac{1+\tau^2}{\varphi+n-1} \sum_{j=1}^{n} A_{j,1} + \frac{\tau^2}{\varphi+n-1}\bar{\mu}_i - \frac{\tau^2}{\varphi+n-1}\sum_{j=1}^{n}\bar{\mu}_j. \tag{37}$$

Terms with summations do not depend on i, and hence are cancelled out in the difference, yielding

$$\begin{aligned}
A^u_{i,\infty} - A^u_{j,\infty} &= \frac{\varphi - \left(1+\tau^2\right)}{\varphi+n-1}\left(A_{i,1} - A_{j,1}\right) + \frac{\tau^2}{\varphi+n-1}\left(\bar{\mu}_i - \bar{\mu}_j\right) \\
&= \frac{\tau^2\sigma^2\left(1+\tau^2\right)}{\gamma}\left(A_{i,1} - A_{j,1}\right) + \frac{\tau^2}{\gamma}\left(\bar{\mu}_i - \bar{\mu}_j\right) \\
&= \frac{\tau^2\sigma^2}{\gamma}\left(\tau^2\left(\mu_i - \mu_j\right) + \left(x_i - x_j\right)\right) + \frac{\tau^2}{\gamma}\left(\bar{\mu}_i - \bar{\mu}_j\right) \\
&= \frac{\tau^4\sigma^2}{\gamma}\left(\mu_i - \mu_j\right) + \frac{\tau^2\sigma^2}{\gamma}\left(\varepsilon_i - \varepsilon_j\right) + \frac{\tau^2}{\gamma}\left(\bar{\mu}_i - \bar{\mu}_j\right).
\end{aligned}$$

Here the second equality is by substitution of the definitions $\varphi = \left(1+\tau^2\right)\left(1+\tau^2\sigma^2\right)$ and $\gamma = \varphi + n - 1$; the third equality is by $\left(1+\tau^2\right)A_{i,1} = \tau^2\mu_i + x_i$ and the last is by $x_i - x_j = \varepsilon_i - \varepsilon_j$. □

Proof of Proposition 1. Note from (6) and (11) that

$$\begin{aligned}
E[A^{ck}_{i,\infty} - A^{ck}_{j,\infty}] &= \frac{\tau^2}{\tau^2+n}\left(\bar{\mu}_i - \bar{\mu}_j\right) \\
E[A^u_{i,\infty} - A^u_{j,\infty}] &= \frac{\tau^2\left(1+\tau^2\sigma^2\right)}{(1+\tau^2)(1+\tau^2\sigma^2)+n-1}\left(\bar{\mu}_i - \bar{\mu}_j\right)
\end{aligned}$$

$E[A^{ck}_{i,\infty} - A^{ck}_{j,\infty}]$ is independent of σ^2 while $E[A^u_{i,\infty} - A^u_{j,\infty}]$ is increasing in σ^2. Since $E[A^u_{i,\infty} - A^u_{j,\infty}] = E[A^{ck}_{i,\infty} - A^{ck}_{j,\infty}]$ for $\sigma^2 = 0$, we have $E[A^u_{i,\infty} - A^u_{j,\infty}] > E[A^{ck}_{i,\infty} - A^{ck}_{j,\infty}]$ for $\sigma^2 > 0$. □

A.3 Social Groups

Proof of Proposition 5. For any $i \in B$, $E_i\left[\hat{\mu}_b\right] = E_i\left[\theta\right] = \bar{\mu}_b$ and $E_i\left[\hat{\mu}_w\right] = \bar{\mu}_w$. Hence, by (20) and (21),

$$E_i\left[\beta_2^S\right] = \alpha_w c_w \left(\bar{\mu}_b - \bar{\mu}_w\right). \tag{38}$$

Similarly, for any $j \in W$,

$$E_j\left[\beta_2^S\right] = \alpha_b c_b \left(\bar{\mu}_b - \bar{\mu}_w\right). \tag{39}$$

Now,

$$\begin{aligned}
\alpha_b c_b &= \frac{\tau^2\left(1+\tau^2\sigma^2\right)}{\tau^2\left(1+\tau^2\sigma^2\right) + n_b\tau^2\sigma^2 + n} \\
\alpha_w c_w &= \frac{\tau^2\left(1+\tau^2\sigma^2\right)}{\tau^2\left(1+\tau^2\sigma^2\right) + n_w\tau^2\sigma^2 + n}.
\end{aligned}$$

Since $n_b < n_w$, we have $c_b\alpha_b > c_w\alpha_w$, showing that $E_i\left[\beta_2^S\right] < E_j\left[\beta_2^S\right]$. To see that $E[\beta_2^I] < E_i[\beta_2^S]$, observe that (15) can be obtained from (38) by setting $\sigma^2 = 0$, and $\alpha_w c_w$ is increasing in σ^2. To see that $E_j[\beta_2^S] < E[\beta_2^I]$, observe that (13) can be obtained from (39) by setting $n_b = 1$, and $\alpha_b c_b$ is decreasing in n_b. $\square$

References

[1] Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz (2008). "Fragility of Asymptotic Agreement under Bayesian Learning", Unpublished Manuscript, MIT.

[2] Adrian, Tobias and Mark Westerfield (2007). "Disagreement and Learning in a Dynamic Contracting Model", *Review of Financial Studies,* forthcoming.

[3] Allen, Michael O., Patrice O'Shaughnessy, and Dean Chang (1995). "Tremor Hits Racial Divide." *New York Daily News*, October 8, 1995.

[4] Aumann, Robert J. (1976). "Agreeing to Disagree." *Annals of Statistics* 4: 1236-1239.

[5] Banerjee, Abhijit V. and Rohini Somanathan (2001). "A Simple Model of Voice." *Quarterly Journal of Economics* 116: 189-227.

[6] Bowles, Samuel, Glenn C. Loury and Rajiv Sethi. (2009). "Group Inequality." Economics Working Papers 88, Institute for Advanced Study.

[7] Brigham, John C. and Adina W. Wasserman (1999). "The Impact of Race, Racial Attitude, and Gender on Reactions to the Criminal Trial of O.J. Simpson." *Journal of Applied Social Psychology* 29: 1333-1370.

[8] Chaudhuri, Shubham and Rajiv Sethi (2008). "Statistical Discrimination with Peer Effects: Can Integration Eliminate Negative Stereotypes?" *Review of Economic Studies* 75: 579-596.

[9] Che, Yeon-Koo and Navin Kartik (2008). "Opinions as Incentives." Unpublished Manuscript, Columbia University.

[10] Crawford, Vincent and Joel Sobel (1982). "Strategic Information Transmission", *Econometrica* 50: 1431-1451.

[11] Currarini, Sergio, Matthew O. Jackson and Paolo Pin (2009). "An Economic Model of Friendship: Homophily, Minorities and Segregation," *Econometrica* 77: 1003-1046.

[12] CNN/Opinion Research Inc. (2008). "Post-Election Race Relations." Accessed at www.cnn.com/2008/POLITICS/11/11/obama.poll/ on November 12, 2008.

[13] Crocker, Jennifer, Riia Luhtanen, Stephanie Broadnax, and Bruce E. Blaine (1999). "Belief in U.S. Government Conspiracies Against Blacks Among Black and White College Students: Powerlessness or System Blame?" *Personality and Social Psychology Bulletin* 25: 941-953.

[14] DeParle, Jason (1990). "Talk of Government Being Out to Get Blacks Falls on More Attentive Ears" *New York Times*, October 29, 1990.

[15] Dixit, Avinash K. and Jorgen W. Weibull (2007). "Political Polarization." *PNAS* 104: 7351-7356.

[16] Eliaz, Kfir and Rani Spiegler (2007). "A Mechanism Design Approach to Speculative Trade" *Econometrica* 75 (3), 875-884.

[17] Freedman, David (1965): "On the Asymptotic Behavior of Bayes Estimates in the Discrete Case II" *Annals of Mathematical Statistics*, 36, 454-456.

[18] Geanakoplos, John D. and Heraklis M. Polemarchakis (1982). "We Can't Disagree Forever." *Journal of Economic Theory* 28: 192-200.

[19] Harrington, Joseph E. (1993). "Economic Policy, Economic Performance, and Elections." *American Economic Review* 83: 27-42.

[20] Harrison, J. Michael and David M. Kreps (1978). "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *Quarterly Journal of Economics* 111: 1111-1133.

[21] Loury, Glenn C. (1977). "A Dynamic Theory of Racial Income Differences." In *Women, Minorities and Employment Discrimination*, edited by Phyllis Wallace and Annette LaMond. Lexington: Lexington Books.

[22] Morris, Stephen (1994). "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica* 62, 1327-1347.

[23] Morris, Stephen (1996). "Speculative Investor Behavior and Learning," *Quarterly Journal of Economics* 111: 1111-1133.

[24] Pew Research Center for the People & the Press (2008). "12% Still Think Obama is Muslim." Accessed at pewresearch.org/databank/dailynumber/?NumberID=509 on October 27, 2008.

[25] Scheinkman, Jose and Wei Xiong (2003). "Overconfidence and Speculative Bubbles." *Journal of Political Economy* 111: 1183-1219.

[26] Sethi, Rajiv and Rohini Somanathan (2004). "Inequality and Segregation." *Journal of Political Economy* 112: 1296-1321.

[27] Van den Steen, Eric (2004). "The Costs and Benefits of Homogeneity, with an Application to Culture Clash." Unpublished Manuscript, MIT.

[28] Van den Steen, Eric (2005). "Organizational Beliefs and Managerial Vision," *Journal of Law, Economics, and Organization* 21: 256-283

[29] Weigel, David (2009). "How Many Southern Whites Believe Obama Was Born in America?" *Washington Independent*, July 31, 2009.

[30] Wilson, Andrea (2003). "Bounded Memory and Biases in Information Processing." Unpublished Manuscript, Princeton University.

[31] Yildiz, Muhamet (2003). "Bargaining without a Common Prior: An Immediate Agreement Theorem." *Econometrica* 71: 793-811.

[32] Yildiz, Muhamet (2004). "Waiting to Persuade." *Quarterly Journal of Economics* 119: 223-248.