A Bayesian Considers Echo Chambers, Group Viewpoints, and Information Accuracy*

Thomas A. Rietz
Tippie College of Business
University of Iowa

January 13, 2026

*Preliminary and incomplete. Please do not quote. Thank you to Amrita Nain, Fotis Grigoris, Richard Peter and Ashish Tiwari for comments and suggestions.

A Bayesian Considers Echo Chambers, Group Viewpoints, and Information Accuracy

Abstract

I develop a Bayesian model of information acquisition that incorporates documented tendencies of groups to have more similar views within group than across groups. Suppose n observers of two types (e.g., Democrats and Republicans, or stock analysts from two different firms) see a signal about a variable of interest (e.g., a future election outcome or future stock returns). First, an unbiased latent signal is drawn for each type that underlies the type's initial viewpoint. Then, observers receive individual specific signals that are unbiased relative to their type's latent signal. Bayesian posteriors based on these signals are unbiased and achieve the lowest posterior variance when the fractions of observer types are equal. However, the posterior variance cannot be reduced to zero unless there are more types of observers with independent initial viewpoints, i.e., viewpoint diversification.

JEL Classifications: C11, D72, D82, G41, M14 D83

Keywords: Bayesian inference, information aggregation, echo chambers, diversity, corporate governance

This paper is motivated by observations from trader behavior and forecasts in prediction markets. In the first paper in modern prediction markets, Forsythe, Nelson, Neumann, and Wright (1992) show that partisan traders appear to have different expectations and interpret information differently across parties. They have different viewpoints. By inferring implied distributions from prices, Berg, Geweke, and Rietz (2010) show that elections seem to be characterized by residual uncertainty. Why, after all the polls, forecasts, and markets, does there remain residual uncertainty? Put another way, why might the "wisdom of the crowd" (Surowiecki (2005) as embodied in prediction markets) *NOT* work?

This paper also arises from observing calls that governance boards be ideologically diversified, meaning roughly even representation across the two major political parties (e.g., Chait, Holland, and Taylor, 1993; Yockey, 2024, among others). Hemel and Feinstein (2018) show that balance requirements are common in theory, if not in practice. Feinstein and Hemel (2024, p. 1,206) show "qualified support for the view that Delaware's constitutional commitment to a politically balanced judiciary adds value to Delaware-incorporated firms." Why should balanced boards across two parties add value?

Similarly, it is motivated by calls to diversify corporate boards across a range of demographic attributes, areas of expertise, etc. This is an area of particularly intense research as shown by an excellent review by Knyazeva, Knyazeva, and Naveen (2021). However, they show that most of the research (briefly reviewed below) is empirical and results are mixed. Further, with some exceptions, there is a notion that more diversity is good, but seldom a straightforward theoretical reason for the notion. Is there a simple information based reason that populating corporate boards with different board member types should be good without resorting to conflicts of interest, adverse selection, or potential managerial malfeasance?

I develop a Bayesian model of information acquisition and aggregation that is consistent with these observations and provides answers for these questions, but can apply in other areas as well.

Consider a situation where a panel of people ("observers") with different initial viewpoints ("types") gather information about an unknown variable with a commonly known prior distribution. For each observer type, there is a common, latent, unobserved, but unbiased, piece of information about the variable that determines a baseline belief or initial viewpoint for all observers of that type.

There are many examples of such initial viewpoint types. Consider well documented behavioral differences between men and women in economic decision-making (Croson and Gneezy, 2009). While one may ascribe differences in behavior to differences in preferences, Harris and Jenkins (2006) argue that different initial beliefs drive different behaviors, not differences in preferences. They show that women have a "greater perceived likelihood of negative outcomes and lesser expectation of enjoyment" then men when assessing risky choices (p. 48). Thus, women start from a different place than men when assessing the likely outcomes of risky decisions. In this case, types correspond to gender.

Political expectations are another example. In an election, Democrats may share some common beliefs, assumptions, and information that form a common viewpoint for their beliefs. Republicans may share different common beliefs, assumptions and, information that form a different common viewpoint for their beliefs. For example, Peterson, Goel, and Iyengar (2021) show that Democrats and Republicans selectively expose themselves to different news sources. So, they have different information. Forsythe et al. (1992) argue that a false consensus effect (Ross, Greene, and House, 1977) cause different expectations across Republican and Democratic traders in political prediction markets. A false consensus effect can arise because partisan individuals gather with similar individuals and have some common expectations that they project onto others (Bauman and Geher, 2002), instead of combining their information with others. Forsythe, Rietz, and Ross (1999) argue that a wishful thinking effect influences how Republican and Democratic traders interpret information and trade in prediction markets. Thus, Democrats and Republicans start from a different place when assessing the likely results of an upcoming election. If one were to poll voters on what they think the outcome a future election will be (i.e., "citizen forecasting," Lewis-Beck and Skalaban, 1989), the analysis here shows how the views of different types of voters can be most efficiently weighted. In these cases, a type corresponds to a party affiliation.

Another example is home country bias in investments. Coval and Moskowitz (1999) argue that the reason for the home country bias (a preference for investing in geographically close companies) is due to information effects. Investors have local knowledge about nearby businesses. The local knowledge is likely to be commonly held among investors in one region and differ from information of investors in other regions. In this case, a type corresponds to a geographic region. There is a similar home team bias in sports. Love, Kopeć, and Guest (2015) suggest that fans of sports teams and sports writers who follow specific teams interpret the same information differently. Given the same past records, fans of a team and writers who follow the team project higher future wins than fans and writers who follow other teams. In this case, type corresponds to a team's fan base.

With respect to corporate boards, Knyazeva et al. (2021) provide a thorough review for the interested reader, so I will be very brief here. They document that researchers have investigated diversity in the form of demographics (e.g., race, gender, and cultural backgrounds), professional backgrounds, and expertise. They cite two papers on political diversity on boards (Kim, Pantzalis, and Park, 2013; Lee, Lee, and Nagarajan, 2014) both suggesting diversity in political beliefs is valuable. They review research on the determinants of board diversity based on firm type, business conditions, firm structure, and other firm specific factors. There is also considerable research on firm outcomes and Knyazeva et al. (2021) find the results are quite mixed.[1] Most of the research Knyazeva et al. (2021) cite is empirical, but based on the ideas that diversity brings in more viewpoints and this should be good. Here, assuming that more accurate information leads to better decisions, I provide a theoretical underpinning for this idea. The theory does not rely on any biases, board politics, any kind of adverse selection or principal agent problems. All that is required is a diversity of initial viewpoints.[2]

The viewpoint effect I model may arise from different experience, backgrounds, and expertise. It can also arise whenever groups are operating in "echo chambers" (when groups share and reinforce their own beliefs to the exclusion of others) or with confirmation bias (when people seek out information that confirms their beliefs, but ignore information that can dispute them). Cinelli, De Francisci Morales, Galeazzi, Quattrociocchi, and Starnini (2021) show that such effects are heavily reinforced by social media and that "users online tend to prefer information adhering to their worldviews, ignore dissenting information, and form polarized groups around shared narratives." Polarized groups can have different initial viewpoints, but even strong beliefs can be unbiased.

My model assumes only that initial viewpoints are different, not that they are systematically biased in any way. Thus, this paper is not about bias (either implicit or explicit). It is also not about politics or dishonesty. It is not about traditional demographic-based notions of diversity or ethical

[1]For example, Kang, Kim, and Oh (2020) find that higher diversity leads to more dissent which leads to higher firm value and better governance. In contrast, Donaldson, Malenko, and Piacentino (2020) find that higher diversity leads to more deadlock and worse outcomes.

[2]This implies that diversity in the form of demographics and expertise *may* constitute a valuable "type" if it is associated with different initial viewpoints, but it will not be valuable without a difference in viewpoints.

arguments for diversity. All observers honestly reveal their information and their types are known, so it is not about situations where observers have incentives to hide or misrepresent information. It is also not about groups that may have special or inside information. The model below only requires differences in (unbiased) initial viewpoints. Given this, how does one incorporate the information of observers optimally in forecasting a variable relevant for decision making?

I model a case where each observer's type is known and taken into account when assessing the value of their information (e.g., we can see their gender, voter registrations, what region they are from, what team jerseys they wear, their specific corporate expertise, etc.). The analysis follows a structure similar to latent variables Bayesian inference models found in the literature (e.g., Gelman, Carlin, Stern, and Rubin, 1995, Chapter 5). While the theory is not particularly novel, the implications are intuitive and important for understanding the properties of posterior distribution mean (i.e., forecast based on the available information) and variance (i.e., the level of residual forecasting error) in such situations. Specifically:

1. If all information is gathered and revealed by the same type of observers, the posterior variance over the variable of interest (i.e., the forecasting error) is higher than when information comes from different types with different initial viewpoints.
2. This problem is worse when initial viewpoints are more polarized.
3. A balance of observers of different types on a panel gives the lowest posterior variance.
4. An imbalance of types requires different weights on information from different types to minimize posterior variance.
5. With a small number of observer types on a panel, there are limits to the ability of the law of large numbers to reduce the forecasting error. The mean of the posterior distribution does not converge to the true underlying value. That is, with a limited number of types, the "wisdom of the crowd" doesn't work.
6. Only through many types of observers with independent initial viewpoints across types can the posterior variance be driven further down.

Thus, panels of observers (e.g, boards, governance bodies, analyst groups, etc.) made up of a single observer type will have the highest posterior variances (i.e., the make the biggest errors). Balanced panels across a few types perform better in reducing posterior variance. The lowest posterior variance is achieved only by panels of many types of observers, i.e., high levels of initial viewpoint diversification. Only then will the mean of the posterior distribution converge to the true underlying value.

To show the potential size of the problem and how it can evolve through time, I use data from the University of Chicago NORC's General Social Survey. For a specific example, I use survey responses from 1974 through 2004 on political orientation and questions about how the government should spend money. Then, I show how to use simple dummy variable regressions to assess the degree of polarization across two types: conservative and liberal respondents. Then, for various numbers of types and observers, I calculate the impact of the polarization on the posterior variance (or equivalently, the mean absolute forecasting error), while holding overall signal variance constant. In 1974, political viewpoints explained little of the variance in opinions on government spending for defense and welfare. Polarization was low. However, by 2021, the impact of political viewpoints rose by 7.4 to 9.4 fold. Thus, the level of political polarization has risen dramatically. The calculations back theoretical results in the paper that (1) the posterior variance is lower when there is less disagreement in initial viewpoints and (2) adding more observers had limited power to solve the problem and (3) adding more observer types more effectively reduces the problem. Finally, it shows the problem is magnified when polarization is high.

If better information leads to better decisions, then this paper makes another important contribution: It does not require any bias on the part of any observer or group to justify balancing panels or expanding them to take advantage of viewpoint diversification.

1 Common Model Elements

Consider a situation where a panel of observers with different initial viewpoints (i.e., types) gather information about an unknown variable with a common prior distribution. Let V be the variable of interest and let

$$V \sim N(0, \sigma_V^2) \text{ with precision } \lambda_V = 1/\sigma_V^2. \tag{1}$$

V might be the vote shares in an upcoming election, the future return on a stock, the future performance of a sports team, or the best strategy for a corporation, etc.[3]

[3]Consider a concrete example. Berg et al. (2010) focus on forecasting the major party vote shares in U.S. Presi-

Observers of different types either access different baseline information or interpret the same baseline information differently, resulting in different initial viewpoints. For each type, a latent, common, unbiased, but unobserved, signal about V is drawn:

$$y_\tau = V + \eta_\tau \text{ where } \eta_\tau \sim N(0, \sigma_\eta^2) \text{ with a common precision } \lambda_\eta = 1/\sigma_\eta^2, \tag{2}$$

where y_τ is the latent signal for all individuals of type τ and forms an initial viewpoint.[4] There is no bias in the latent signals, just noise.

Then, each observer of type τ receives an individual signal:

$$y_{i,\tau} = y_\tau + \epsilon_i = V + \eta_\tau + \epsilon_i \text{ where } \epsilon_i \sim N(0, \sigma_\epsilon^2) \text{ with a common precision } \lambda_\epsilon = 1/\sigma_\epsilon^2, \tag{3}$$

where $y_{i,\tau}$ is the signal received by observer i (who is of type τ) and ϵ_i is independent of η_τ. Thus, an observer's overall signal has expected value V plus of a combination of the common, but unobserved, initial viewpoint (η_τ) plus individual variation from that initial viewpoint (ϵ_i).[5]

Note that, conditional on V, each type's latent signal is distributed $N(V, \sigma_\eta^2)$ and each individual signal is distributed $N(V, \sigma_\eta^2 + \sigma_\epsilon^2)$. So, both are effectively unbiased forecasts of V. Thus, each observer has an overall unbiased signal.

2 Two Types, Two Individuals Cases

Suppose there are two Types denoted by $\tau \in \{D, R\}$. While D and R can be thought of as Democrat and Republican, they can be any two types of observers who have different latent initial information/viewpoints across types, but common within types. Suppose there are two individual observers on the panel. The Bayesian posterior distributions vary based on whether the observers are of the same type or different types.

dential elections. They model a normal prior distribution for $V = ln(f/(1-f))$ where f is the fraction of the two party vote taken by the Democratic candidate. They parameterize their prior distribution by fitting it to historical popular vote shares.

[4]The common precision simplifies the exposition, but the intuition generalizes to different precisions.

[5]The false consensus idea leads to another way to think about this. People form groups based on the closeness of their private signal. As Bauman and Geher (2002, p. 314) put it "people tend to be friends with people who hold similar attitudes." This self-selection effectively makes groups share a correlation that is modeled by the common latent signal here.

2.1 Case 1: Two types, two observers, same type

Without loss of generality, suppose both observers are type $\tau = D$. The individual signals are:

$$y_i = V + \eta_D + \epsilon_i, \quad i = 1, 2. \tag{4}$$

The two observed signals share the same latent noise η_D constituting the initial viewpoint with precision λ_η. They each have an independent noise of ϵ_i with precision λ_ϵ. If the two individuals were both Type R, they would have the same precisions for the latent and independent noise. So, one can drop the type designations and condition simply on the fact that they are the same type. In that case, the precision contributed by the two signals from the same type is:

$$\Lambda_{\text{same}} = \frac{2}{2\sigma_\eta^2 + \sigma_\epsilon^2} = \frac{2}{\frac{2}{\lambda_\eta} + \frac{1}{\lambda_\epsilon}} = \frac{2\lambda_\eta\lambda_\epsilon}{\lambda_\eta + 2\lambda_\epsilon}, \tag{5}$$

where "same" denotes conditioning on the two observers being the same type.

Then, the posterior distribution is:

$$V|y_{\text{same}} \sim N\left(\frac{\Lambda_{\text{same}}}{\lambda_V + \Lambda_{\text{same}}} \times \bar{y}, \frac{1}{\lambda_V + \Lambda_{\text{same}}}\right), \tag{6}$$

where $\bar{y}$ is the average signal.

The fact that the two observers have the same initial latent viewpoint affects the precision contributed by the two signals. In turn, this affects the posterior expected value because it is a weighted average of the prior expected value (0) and the sample mean where the weights depend on the precision contributed by the signals. The posterior precision is the sum of the initial precision and the precision contributed by the signals. Therefore, as the inverse of the posterior precision, the posterior variance is affected by the same factors.

2.2 Case 2: Two types, two observers, different types

Now suppose the panel is made up of one type D observer and one type R. In this case, the observer signals are:

$$y_1 = V + \eta_D + \epsilon_1 \text{ and } y_2 = V + \eta_R + \epsilon_2. \tag{7}$$

The two signals have independent latent noise creating independent initial viewpoints, but they have the same precision λ_η. They each have an independent noise of ϵ_i with precision λ_ϵ. So, again, one can drop the type designations and condition simply on the fact that they are different. In that case, the precision contributed by the two signals is:

$$\Lambda_{\text{diff}} = \frac{2}{\sigma_\eta^2 + \sigma_\epsilon^2} = \frac{2}{\frac{1}{\lambda_\eta} + \frac{1}{\lambda_\epsilon}} = \frac{2\lambda_\eta\lambda_\epsilon}{\lambda_\eta + \lambda_\epsilon}, \tag{8}$$

where "diff" denotes conditioning on the two observers being of different types.

Then, the posterior distribution is:

$$V|y_{\text{diff}} \sim N\left(\frac{\Lambda_{\text{diff}}}{\lambda_V + \Lambda_{\text{diff}}} \times \bar{y}, \frac{1}{\lambda_V + \Lambda_{\text{diff}}}\right), \tag{9}$$

where, again, $\bar{y}$ is the average signal.

The precision contributed by the two signals is no longer affected by the same latent initial viewpoint. The signals are now completely independent. Relative to the same type case, this changes the mean update and the posterior variance. .

2.3 Comparing Case 1 and Case 2

First, note that the sensitivity of the posterior mean to the data and the level of the posterior variance are monotone in the precision of the data. The posterior *update weight* on the data, $\Lambda_{\text{data}}/(\lambda_V + \Lambda_{\text{data}})$, is strictly increasing in Λ_{data}. To see this, ake the partial derivative of the update weight with respect to Λ_{data}, giving:

$$\frac{\partial}{\partial \Lambda_{\text{data}}}\left(\frac{\Lambda_{\text{data}}}{\lambda_V + \Lambda_{\text{data}}}\right) = \frac{\lambda_V}{(\lambda_V + \Lambda_{\text{data}})^2} > 0. \tag{10}$$

Further, the posterior variance, $1/(\lambda_V + \Lambda_{\text{data}})$, is strictly decreasing in Λ_{data}. Taking the partial derivative with respect to Λ_{data} gives:

$$\frac{\partial}{\partial \Lambda_{\text{data}}}\left(\frac{1}{\lambda_V + \Lambda_{\text{data}}}\right) = -\frac{1}{(\lambda_V + \Lambda_{\text{data}})^2} < 0. \tag{11}$$

Given this, we can easily compare the posterior mean and variance based on same or different type observers, which leads to the following proposition:

Proposition 1. *When the two observers are of the same type, (1) the posterior expected value is less responsive to the data than when the observers are of different types and (2) the posterior variance is higher than the observers are of different types.*

Proof. Because, $\lambda_\eta + 2\lambda_\epsilon > \lambda_\eta + \lambda_\epsilon$:

$$\Lambda_{\text{same}} = \frac{2\lambda_\eta\lambda_\epsilon}{\lambda_\eta + 2\lambda_\epsilon} < \frac{2\lambda_\eta\lambda_\epsilon}{\lambda_\eta + \lambda_\epsilon} = \Lambda_{\text{diff}}. \tag{12}$$

Because two observers of different types yields a higher precision than two observers of the same type, the posterior update weight is higher (by equation 10) and the posterior variance is lower (by equation 11). □

The posterior update weights multiplying the sample means in equations 9 and 6 show how much the mean of the posterior distribution changes when the data are observed. Signals from two different types have more information content than two signals from the same type, so heavier weight is placed on the sample mean (i.e., the data). As a result, the posterior mean moves more.

The differential information contents also affects the posterior variance. The higher information resulting from signals form two types increases the precision of the data and decreases the posterior variance. Note that the mean absolute prediction error for the posterior normal distribution equals $\sqrt{\frac{2}{\pi}}$ times the posterior standard deviation. So, the lower the posterior variance, the lower the prediction errors will be based on the posterior distribution. Alternatively, the higher posterior variance resulting from a single type of observer implies higher residual forecasting uncertainty.

Figure 1 illustrates the wedges that occur in the posterior variance and mean resulting from having two individuals of the same type versus different types. Panel A shows the posterior variance. The larger the variance in the latent type signals (i.e., the greater the variance in initial viewpoints), the larger the wedge. Panel B shows the posterior mean for a mean signal of 1. The higher the latent variance, the smaller the update. When, there are two observers of one type, the dampening effect of the latent variance is magnified. Thus, if we are becoming more polarized (as asserted by Cinelli et al., 2021), then it is even more valuable to have different viewpoints among the observers.

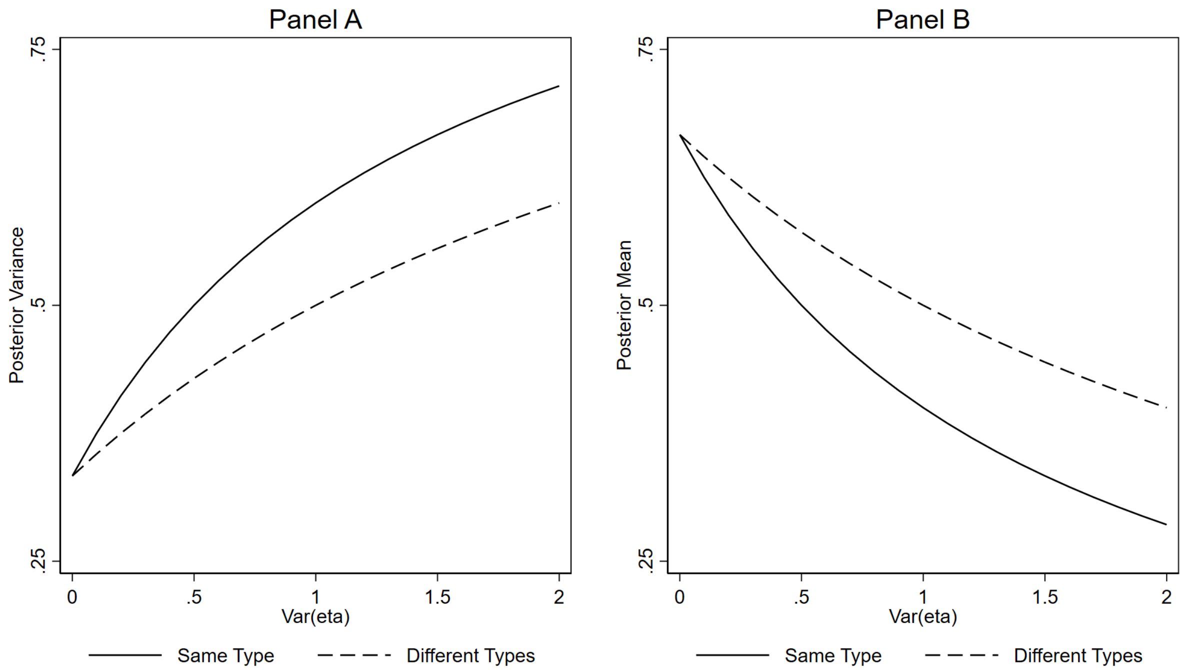

Figure 1. Posterior variance and mean when $\sigma_V^2 = \sigma_\epsilon^2 = 1$ and σ_η^2 ranges from 0 to 2. Panel A shows the posterior variance. Panel B shows the posterior mean for a mean data signal of 1.

3 Generalization to n observers of two types

Now suppose that the panel consists of n observers and there are two types of observers. Let n_D observers see $y_{i,D} = V + \eta_D + \epsilon_i$ and n_R see $y_{i,R} = V + \eta_R + \epsilon_i$ and let $n = n_D + n_R$. Thus, the panel consist of n members, n_D Democrats, and n_R Republicans (or whatever two types you want to think of). Denote the sample means of the observers by $\bar{y}_D$ and $\bar{y}_R$ for type D and type R observers, respectively.

Each type shares a common latent viewpoint, so information within a type is correlated. The precision contributed by type D observers is:

$$\Lambda_D = \frac{n_D}{\sigma_\epsilon^2 + n_D \sigma_\eta^2} = \frac{n_D \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_D \lambda_\epsilon}. \tag{13}$$

Similarly for type R observers:

$$\Lambda_R = \frac{n_R \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_R \lambda_\epsilon}. \tag{14}$$

Thus the total data precision is:

$$\Lambda_{\text{data}} = \Lambda_D + \Lambda_R = \frac{n_D \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_D \lambda_\epsilon} + \frac{n_R \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_R \lambda_\epsilon}. \tag{15}$$

The posterior precision is $\lambda_V + \Lambda_{\text{data}}$ and the posterior distribution becomes:

$$V \mid y \sim N\left(\frac{\Lambda_{\text{data}}}{\lambda_V + \Lambda_{\text{data}}} \cdot \bar{y}, \frac{1}{\lambda_V + \Lambda_{\text{data}}}\right), \tag{16}$$

where $\bar{y} = \frac{\Lambda_D \, \bar{y}_D + \Lambda_R \, \bar{y}_R}{\Lambda_{\text{data}}}$. Thus, the posterior mean optimally weights the signals from the type D and type R observers in forecasting V. The weights come from the relative precision contributed by each signal within each type balanced against the precision contributed by signals across types. Analyzing the results leads to the following proposition.

Proposition 2. *For fixed total sample size $n = n_D + n_R$, the data precision is maximized when n_D and n_R are as equal as possible (i.e., $n_D = n_R$ when n is even, and $|n_D - n_R| = 1$ when n is odd). This (1) minimizes the posterior variance while (2) maximizing the sensitivity of the mean of the posterior distribution to changes in the data.*

Proof. Define $g(m) = \dfrac{m \, \lambda_\eta \, \lambda_\epsilon}{\lambda_\eta + m \, \lambda_\epsilon}$ for $m \geq 0$. Recall, $\lambda_\eta, \lambda_\epsilon > 0$. The first and second derivatives of $g(m)$ are:

$$g'(m) = \frac{\lambda_\eta^2 \lambda_\epsilon}{(\lambda_\eta + m\lambda_\epsilon)^2} > 0 \qquad \text{and} \qquad g''(m) = -\frac{2\lambda_\eta^2 \lambda_\epsilon^2}{(\lambda_\eta + m\lambda_\epsilon)^3} < 0. \tag{17}$$

Thus, $g(m)$ is increasing and strictly concave. With $n = n_D + n_R$ fixed, Jensen's inequality implies:

$$g(n_D) + g(n_R) \leq 2g\left(\frac{n_D + n_R}{2}\right) = 2g\left(\frac{n}{2}\right), \tag{18}$$

with equality if and only if $n_D = n_R$. Thus, the most balanced split maximizes Λ_{data}. Maximizing Λ_{data} both maximizes the posterior update weight (by equation 10) and minimmizes the posterior variance (by equation 11). □

What does this mean? Splitting observers across the two latent viewpoints produces more independent information (less shared latent noise) than concentrating observers on a single latent viewpoint. This maximizes the update in the forecast that results from a given weighted average of signals while minimizing the posterior variance and expected forecasting error.[6]

Figure 2 shows the impact of imbalanced types on the posterior variance (Panel A) and mean for a weighted average signal of 1 (Panel B). With all observers from one type and none from the other, the panel becomes an echo chamber with the largest posterior variance (Panel A) and the lowest sensitivity to the data (Panel B). As the number of type D observers range from 0 to 10 out of 10, Figure 2 shows that the largest impact is going from 0 to 1 or 9 to 10. This arises because, if all observers are one type, all information has the same underlying latent viewpoint and, therefore, is correlated. Just one viewpoint from another type provides the only independent information (from the other underlying latent signal) and, therefore, is heavily weighted.[7]

Breaking up the echo chamber matters a lot. Figure 2 shows that the biggest impact arises from the first observer with an alternative viewpoint. Similarly, the mean of the posterior distribution will shift more when adding one observer of a different type than when adding more. In fact, additional weight on the minority type observations continues to emphasize the value of the uncorrelated initial viewpoints when the numbers are larger.

Observation 1. *On a per capita basis, the smaller group is always overweighted in determining the posterior expected value.*

Proof. Let $w(n_\tau)$ represent the weight per capita for an observer in a type τ with n_τ observers. The weight per capita is:

$$w(n_\tau) = \frac{\Lambda(n_\tau)}{n_\tau} = \frac{\lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_\tau \lambda_\epsilon}. \tag{19}$$

This is decreasing in n_τ since

$$w'(n_\tau) = -\frac{\lambda_\eta \lambda_\epsilon^2}{(\lambda_\eta + n_\tau \lambda_\epsilon)^2} < 0. \tag{20}$$

[6]The perfectly even split arises because both types have the same initial latent viewpoint precision (λ_τ). If one type has more precision than the other, then there will be a heavier weighting on that type. However, observers of the other type will generally still add precision because of the independence of their latent signal. More specifically, if latent precisions differ across types, the contributed precision for type τ with n_τ observers is $\Lambda_\tau = \frac{n_\tau \lambda_{\eta,\tau} \lambda_\epsilon}{\lambda_{\eta,\tau} + n_\tau \lambda_\epsilon}$. For types D and R, total precision is $\Lambda_{\text{data}} = \frac{n_D \lambda_{\eta,D} \lambda_\epsilon}{\lambda_{\eta,D} + n_D \lambda_\epsilon} + \frac{n_R \lambda_{\eta,R} \lambda_\epsilon}{\lambda_{\eta,R} + n_R \lambda_\epsilon}$ and the optimal allocation of observers equates marginal gains from the observers' information: $\frac{\lambda_{\eta,D}^2}{(\lambda_{\eta,D} + n_D \lambda_\epsilon)^2} = \frac{\lambda_{\eta,R}^2}{(\lambda_{\eta,R} + n_R \lambda_\epsilon)^2}$. This implies the weighting should emphasize signals from the type with higher latent precision, but not ignore signals from the other type.

[7]Interestingly, this contrasts with the "critical mass" idea for corporate boards outlined in Schwartz-Ziv (2017). If all board members are rational, they should place the most weight on the first board member with a different viewpoint.

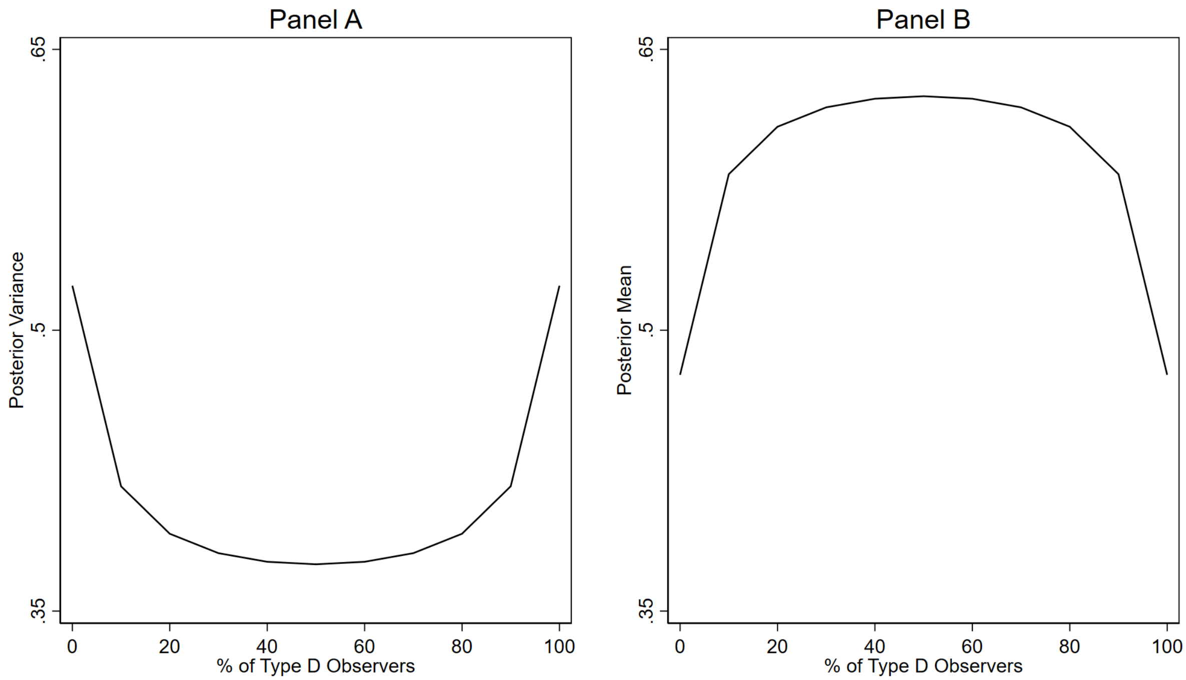

Figure 2. Posterior variance and mean for 10 observers with 0 to 10 of one type when $\sigma_V^2 = \sigma_\eta^2 = \sigma_\epsilon^2 = 1$. Panel A shows the posterior variance. Panel B shows the posterior mean for a mean data signal of 1.

□

Interestingly, if urban and rural voters constitute different types, then this is what the electoral college does by giving more electoral college votes per capita to smaller population rural states.

The next proposition says that there is a limit to how much incorporating two initial latent viewpoints can limit posterior variance.

Proposition 3. *As a board increases in size, with $n_D \to \infty$ and $n_R \to \infty$, (1) the posterior variance around the mean within each group converges to zero but, (2) overall, the variance of the posterior distribution of V converges to $Var(V \mid y) \to 1/(\lambda_V + 2\lambda_\eta)$, which is strictly positive.*

Proof. For two types, the data precision is

$$\Lambda_{\text{data}} = \frac{n_D \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_D \lambda_\epsilon} + \frac{n_R \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_R \lambda_\epsilon}. \tag{21}$$

As $n_D \to \infty$, the first term converges to λ_η because

$$\lim_{n_D \to \infty} \frac{n_D \lambda_\eta \lambda_\epsilon}{\lambda_\eta + n_D \lambda_\epsilon} = \lim_{n_D \to \infty} \frac{\lambda_\eta \lambda_\epsilon}{\frac{\lambda_\eta}{n_D} + \lambda_\epsilon} = \frac{\lambda_\eta \lambda_\epsilon}{\lambda_\epsilon} = \lambda_\eta. \tag{22}$$

Similarly, the second term converges to λ_η. This implies that $\Lambda_{\text{data}} \to 2\lambda_\eta$ and $\text{Var}(V \mid y) = 1/(\lambda_V + \Lambda_{\text{data}})$ converges to $1/(\lambda_V + 2\lambda_\eta)$, which is strictly positive. □

Further, the posterior mean does not coverage to V as shown in the following proposition:

Proposition 4. *As $n_D, n_R \to \infty$, the posterior mean does not converge to the true value V but to a random limit depending on the latent signals η_D, η_R.*

Proof. The posterior mean is:

$$E(V \mid y) = \frac{\Lambda_D \bar{y}_D + \Lambda_R \bar{y}_R}{\lambda_V + \Lambda_D + \Lambda_R}, \tag{23}$$

where $\bar{y}_D$ and $\bar{y}_R$ are the sample means for each type and Λ_D, Λ_R are their contributions to the overall precision. As $n_D \to \infty$, $\bar{y}_D \to V + \eta_D$ and $\Lambda_D \to \lambda_\eta$ because individual noise averages out but latent noise remains. Similarly for type R. Thus,

$$E(V \mid y) \to \frac{\lambda_\eta(V + \eta_D) + \lambda_\eta(V + \eta_R)}{\lambda_V + 2\lambda_\eta} = V + \frac{\lambda_\eta(\eta_D + \eta_R)}{\lambda_V + 2\lambda_\eta}. \tag{24}$$

Since η_D, η_R are independent, zero-mean, and non-degenerate, the limiting posterior mean is random and generally not equal to V. Therefore, even with arbitrarily many observers per type, residual uncertainty persists. □

Intuitively, within each type, averaging across observations wipes out the individual noise, but does not wipe out the type-specific noise in the form of η_τ, which is common to the whole group. Thus even with arbitrarily many observers in each type, the posterior retains randomness coming from the ηs. This prevents converging to the true V and leaves residual uncertainty.

Figure 3 shows how the posterior variance (Panel A) and mean (Panel B, with a weighted average signal of 1) respond to more observers with equal numbers of each observer type ranging from 1 to 10 per type. The posterior variance rapidly falls to a limit equal to $\frac{1}{\lambda_V + 2\lambda_\eta}$. The level of residual uncertainty depends on the precision of the prior and the precision of the two individual latent signals/viewpoints. Similarly, the posterior mean update rises to a limit equal to $\frac{2\lambda_\eta}{\lambda_V + 2\lambda_\eta}$, which is well below the weighted average signal of 1. Increasing the number of observers does not change these limits. Thus, regardless of the number of observers, there will always be residual uncertainty.

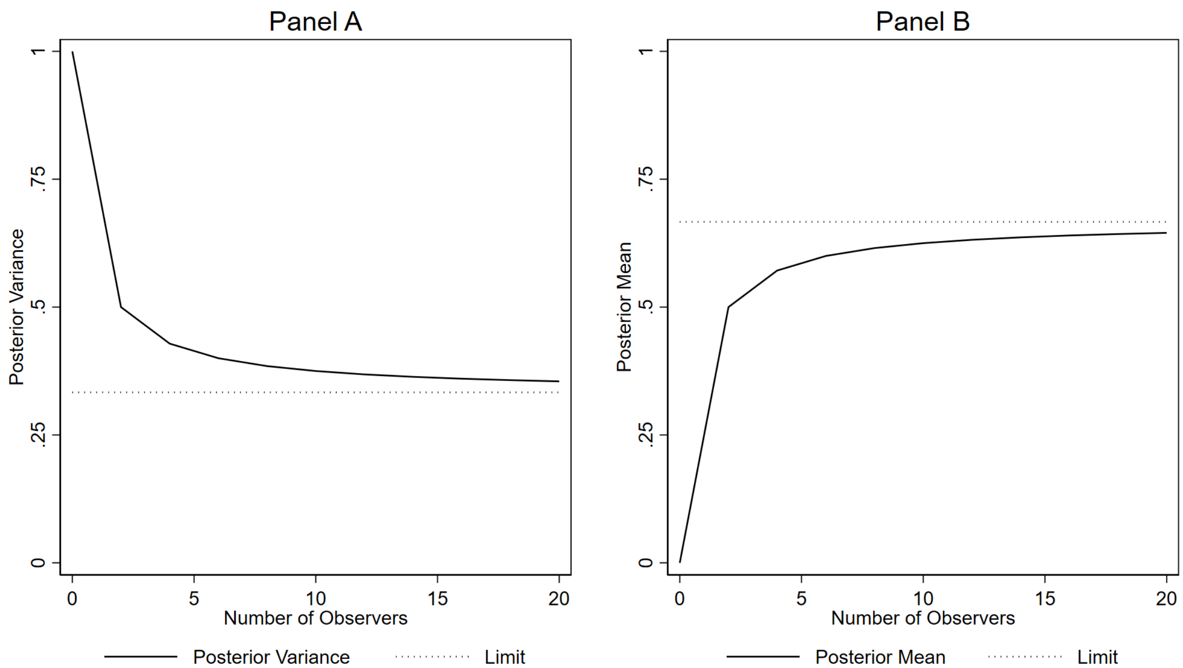

Figure 3. Posterior variance and mean when $\sigma_V^2 = \sigma_\eta^2 = \sigma_\epsilon^2 = 1$ and $n_D = n_R$ ranges from 0 to 20. Panel A shows the posterior variance. Panel B shows the posterior mean for a mean data signal of 1.

For elections, there is data consistent with this lack of convergence: Berg et al. (2010) document residual uncertainty in their posterior distributions of vote shares in Iowa Electronic Markets on U.S. Presidential elections all the way through election day. So do Gruca and Rietz (2021).[8]

4 Generalization to n Types: The value of viewpoint diversification

The analysis above shows, that a balance of types minimizes the residual variance in the posterior distribution of V when there are two (or few) types. But, with a limited number of types, there is a limit to the amount of uncertainty that can be resolved. What if the panel was based on a large

[8]While their data does not go all the way through to the election, Gruca and Rietz (2025) show a similar pattern.

number of types with only one person from each type? That is, what if the panel were designed to capture a larger range of viewpoints instead of being balanced across two viewpoints? The answer shows the value of initial viewpoint diversification.

First, get the latent signals for each type and, becasue there is only one observer of each type, we can index both the observer and type by the same variable, i (i.e., let $i = \tau$). For $i = 1, \dots, n$. Each observer sees the combination of their own latent and individual signals:

$$y_i = V + \eta_i + \epsilon_i, \qquad \eta_i \sim N(0, \sigma_\eta^2), \qquad \text{and} \qquad \epsilon_i \sim N(0, \sigma_\epsilon^2), \tag{25}$$

with all η_i and ϵ_i independent of one another and of V.[9] This means each observer's precision is:

$$\lambda_{\text{obs}} = \frac{1}{\sigma_\eta^2 + \sigma_\epsilon^2} = \frac{1}{\frac{1}{\lambda_\eta} + \frac{1}{\lambda_\epsilon}} = \frac{\lambda_\eta \lambda_\epsilon}{\lambda_\eta + \lambda_\epsilon}, \tag{26}$$

where "obs" denotes a single observer or observation.

Given n observers, the posterior distribution is:

$$V \mid y_1, \dots, y_n \sim N\Big(\frac{n\lambda_{\text{obs}}}{\lambda_V + n\lambda_{\text{obs}}}\bar{y}_i, \frac{1}{\lambda_V + n\lambda_{\text{obs}}}\Big). \tag{27}$$

Now compare the response in the posterior expected value to the same overall set of signals depending on whether the signals are split across two types versus all completely independent. The response is larger when the signals are independent as shown in the following proposition.

Proposition 5. *Suppose there are $n > 2$ observations and either (1) they come from two types with $n_D + n_R = n$ or (2) from n independent types (one observation per type). Then, the posterior precision under n independent types exceeds the posterior precision under any two-type allocation. Consequently, the posterior mean is strictly more responsive to the data in the independent case.*

Proof. For two types, the total precision is maximized by a balanced split $n_D = n_R = n/2$, which gives:

$$\Lambda_{\text{two-types}} = \frac{n\lambda_\eta \lambda_\epsilon}{\lambda_\eta + \frac{n}{2}\lambda_\epsilon}. \tag{28}$$

[9] If there were more than one observer of each type, would go back to being indexed by τ and observers by i and each signal would be $y_{i,\tau} = V + \eta_\tau + \epsilon_i$. The same intuition and similar math would hold for the rest of this section.

For n independent types (one per observation), each contributes $\lambda_{\text{obs}} = \frac{\lambda_\eta \lambda_\epsilon}{\lambda_\eta + \lambda_\epsilon}$ so the overall precision is:

$$\Lambda_{\text{indep}} = n\lambda_{\text{obs}} = n\frac{\lambda_\eta \lambda_\epsilon}{\lambda_\eta + \lambda_\epsilon}. \quad (29)$$

Since $n > 2$ implies $\lambda_\eta + \frac{n}{2}\lambda_\epsilon > \lambda_\eta + \lambda_\epsilon$, the denominator in $\Lambda_{\text{two-type}}$ is strictly larger than that in Λ_{indep}. As a result, $\Lambda_{\text{two-types}} < \Lambda_{\text{indep}}$.

Because n independent types of observers yields a higher precision than the best split of n observers across two types, the posterior update weight is higher (by equation 10) and the posterior variance is lower (by equation 11). □

Intuitively, the condition $\Lambda_{\text{two-types}} < \Lambda_{\text{indep}}$ says that even the *best* allocation across types (the balanced split) yields less total information than n independent types. This means that the precision of the n independent types is higher, which directly translates into greater influence of the data on the posterior mean. Its fairly easy to see that this will generalize to any number of types where the number is less than n.

The next proposition shows that, under the n independent types case, the posterior variance falls to zero and the mean observation becomes a consistent estimator of V when n is large:

Proposition 6. *Suppose the n observations come from n independent types (one observer per type). Then, as $n \to \infty$, (1) the posterior variance converges to zero and (2) the sample mean observation is a consistent estimator of V.*

Proof. Because the n observations are conditionally independent given V, the data precision is additive, so:

$$\Lambda_{\text{data}} = n\lambda_{\text{obs}}, \quad \text{where} \quad \lambda_{\text{obs}} = \frac{\lambda_\eta \lambda_\epsilon}{\lambda_\eta + \lambda_\epsilon} > 0. \quad (30)$$

As a result

$$Var(V \mid y_1, \dots, y_n) = \frac{1}{\lambda_V + \Lambda_{\text{data}}} = \frac{1}{\lambda_V + n\lambda_{\text{obs}}} \xrightarrow[n \to \infty]{} 0, \quad (31)$$

proving part (1).

The mean of the posterior distribution is:

$$E(V \mid y_1, \ldots, y_n) = \frac{n\lambda_{\text{obs}}}{\lambda_V + n\lambda_{\text{obs}}}\bar{y}, \quad \text{where} \quad \bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i. \tag{32}$$

By the law of large numbers, $\bar{y} \to V$ almost surely because $\{y_i\}$ are i.i.d. with mean V and finite variance. Because the weight $\frac{n\lambda_{\text{obs}}}{\lambda_V + n\lambda_{\text{obs}}} \to 1$:

$$E(V \mid y_1, \ldots, y_n) = \frac{n\lambda_{\text{obs}}}{\lambda_V + n\lambda_{\text{obs}}}\bar{y} \xrightarrow[n\to\infty]{\text{a.s.}} V, \tag{33}$$

proving part (2). □

Thus, in contrast to having many observers of few types, having many types of observers with few observers of each type breaks the lower bound on the posterior residual variance and results in estimates that converge to the actual value of V. Figure 4 adds to Figure 3 lines showing the posterior variance (Panel A) and mean (Panel B with a weighted average signal of 1) as the number of types increases instead of increasing the number of observers of only two types. One observer each of three types matches the limits achievable with two types and n observers each. Eventually, the posterior variance converges to zero and the posterior mean converges to the average of the signals which, in turn, converges to the true value. Thus, viewpoint diversification can overcome the lower bound on the posterior variance and allow the posterior mean to converge to the true value of V.

5 Calibration

There are many potential ways groups can form with different initial view points, i.e., common η's within group and different ηs across groups. First, *because* they are a member of a group, individuals may start with common viewpoints. This may arise if they share a common initial information set. For example, Democrats may attend common functions and listen to common news sources *because* they are Democrats. The same may hold for Republicans. This is one rational for the false consensus effect (Ross et al., 1977). Second, individuals may become part of a group *because* they have a common η. This is another rational for the false consensus effect (Bauman and Geher, 2002). As an example, I consider differences in viewpoints along a liberal/conservative spectrum. This likely fits better with the latter idea.

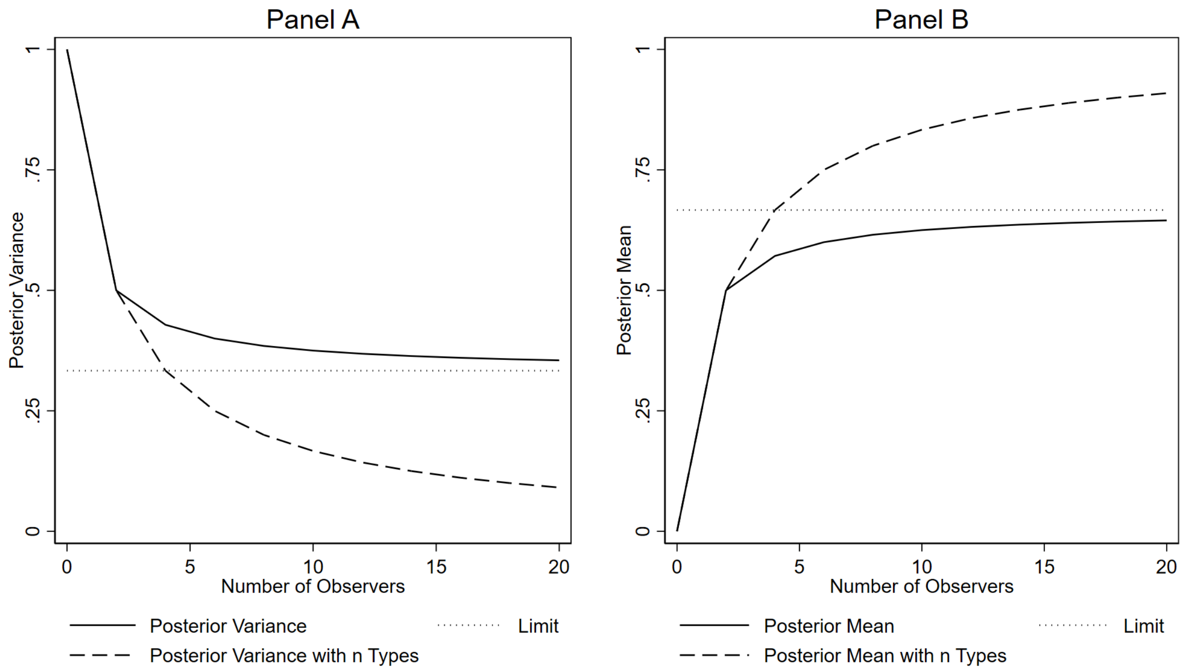

Figure 4. Posterior variance and mean when $\sigma_V^2 = \sigma_\eta^2 = \sigma_\epsilon^2 = 1$. Panel A shows the posterior Variance. The solid line shows posterior variance with $n_D = n_R$ ranging from 0 to 10. The lower dashed line shows the number types ranging from 0 to 20 with one observer of each type. Panel B shows the posterior mean for a mean data signal of 1. The solid line shows posterior mean with $n_D = n_R$ ranging from 0 to 10. The upper dashed line shows the number types ranging from 0 to 20 with one observer of each type.

To show the impact of this using real-world data, I downloaded the University of Chicago NORC's General Social Survey (https://gss.norc.org/). I created two groups based on the respondents' surveyed political viewpoints (the survey's POLVIEWS variable) by assigning respondents who said they were "liberal" or "extremely liberal" a value of 0 and those who said they were "conservative" or "extremely conservative" a value of 1.[10]

Now, consider a regression to explain views of observers with a dummy variable indicating the observer is one of two types. The R^2 of that regression shows the percentage of the total variance in responses due to group membership (i.e., the η's). The remaining variance is due to individual

[10] This measure is for illustrative purposes. Respondents could also respond with "slightly" conservative or liberal or "moderate." Thus, my measure creates two types with a relatively high level of difference, but not the highest. One could blend in "slightly" responses and it dampens the effect to some degree, but leaves the overall point unchanged. Using only "extremely" responses does the opposite. Using the variable PARTYID and grouping over Democrats ("strong Democrat" and "not very strong Democrat") versus Republicans ("strong Republican" and "not very strong republican") gives similar results.

variation (i.e., the ϵ's). This allows one to estimate the relative size of the variances across types versus across individuals with a simple regression. Using data from every survey year from 1974 through 2024, I ran the political views dummy as the independent variable in a regression explaining the respondents' views on whether the government was spending too little, about right, or too much on (1) arms and national defense (NATARMS variable) and (2) welfare (NATFARE variable).[11]

Table 1 shows the results of the dummy variable regressions. The table shows two important things. First, there are generally significant differences in opinion on government spending on arms and defense and on welfare. All of the coefficient estimates are significant, all but two at the 99% level of confidence. Second, the R^2s have been changing through time. Because the R^2s represent the fraction of variance explained by the two ideological group types, the implies that polarization is changing across time. Figure 5 clearly shows that recent opinions are driven more by ideological differences and less by individual variation than they were before the turn of the century. The R^2 for arms in defense in 2021 was 9.4 times what it was in 1974. The R^2 for welfare in 2021 was 7.4 times the 1974 R^2.

Based on each R^2, I calibrate the model keeping the total variance of the signals constant. This can be done by normalizing the variables to mean zero and variances that reflect the R^2. Assume, as before, that $V \sim N(0, 1)$. Then, normalize σ_η^2 and σ_ϵ^2 to sum to 2 and $\frac{\sigma_\eta^2}{\sigma_\eta^2+\sigma_\epsilon^2} = R^2$ for a given year in Table 1. This normalizes the calibrations so that the total signal variance of 3 matches Figures 2 through 4 while splitting the signal variance across the types to match the R^2 of that year's regression. This gives $\sigma_\eta^2 = 2 \times R^2$ and $\sigma_\epsilon^2 = 2 \times (1 - R^2)$. Next, suppose that m is the number of types and N is the total number of observers. A simple extension of Proposition 2 would show that $n = N/m$ observers of each type would minimize the posterior variance. In this case:

$$\Lambda_{\text{data}} = \sum_{\tau=1}^{m} \frac{\frac{N}{m}\lambda_\eta\lambda_\epsilon}{\lambda_\eta + \frac{N}{m}\lambda_\epsilon} = \frac{N\lambda_\eta\lambda_\epsilon}{\lambda_\eta + n\lambda_\epsilon}. \tag{34}$$

Using this, I calculate the posterior variance for 1, 2, 5 and 10 types of observers with 1, 2, 5, 10, 50 and 100 observers split evenly across types using the 1974 data (representative of a low polarization case), 2021 data (the maximum polarization in the data set) and 2024 (the most recent observation).

[11] I note that, because the survey response scale is discrete (1–3), an ordered logit model is technically the correct regression technique to use. This gives similar directional results, but there is not an exact mapping from the psuedo-R^2 to the variance decomposition. So, I use the simpler OLS regression to make the interpretation direct here for illustrative purposes.

Table 1. Dummy variable regressions. In each case, the dummy variable is 0 if the respondent stated they were "liberal" or "extremely liberal" and 1 if they stated they were "conservative" or "extremely conservative." The dependent variable ranged rated spending as 1 ("too little"), 2 ("about right") or 3 ("too much").

Year	Arms and Defense			Welfare		
	Coef.	**Std. Error**	**R2**	**Coef.**	**Std. Error**	**R2**
1974	-0.3149	0.073***	0.0456	0.3858	0.0798***	0.0549
1975	-0.4274	0.0711***	0.0853	0.4456	0.0832***	0.0690
1976	-0.3886	0.0755***	0.0615	0.4776	0.0665***	0.1116
1977	-0.4144	0.0733***	0.0759	0.3653	0.0729***	0.0585
1978	-0.3460	0.0763***	0.0554	0.2969	0.0759***	0.0394
1980	-0.3509	0.0765***	0.0548	0.3217	0.0786***	0.0441
1982	-0.5903	0.069***	0.1269	0.5094	0.0731***	0.0859
1983	-0.6131	0.1087***	0.1360	0.3934	0.1154***	0.0552
1984	-0.5205	0.1277***	0.1199	0.5516	0.1459***	0.1048
1985	-0.1853	0.105*	0.0155	0.3501	0.1069***	0.0501
1986	-0.3323	0.1077***	0.0473	0.4905	0.1112***	0.0929
1987	-0.5172	0.1134***	0.1184	0.4608	0.1303***	0.0742
1988	-0.3769	0.0977***	0.0656	0.3654	0.1126***	0.0471
1989	-0.2843	0.0931***	0.0407	0.4350	0.105***	0.0745
1990	-0.3642	0.0946***	0.0697	0.5020	0.1128***	0.0935
1991	-0.4543	0.0845***	0.1166	0.6620	0.1062***	0.1581
1993	-0.3399	0.0896***	0.0601	0.2658	0.1091**	0.0258
1994	-0.4278	0.0646***	0.0827	0.4708	0.0607***	0.1105
1996	-0.5096	0.0674***	0.1150	0.4875	0.0706***	0.0968
1998	-0.5040	0.0652***	0.1179	0.4147	0.0703***	0.0728
2000	-0.4091	0.0684***	0.0760	0.2952	0.0708***	0.0385
2002	-0.5641	0.0981***	0.1376	0.5007	0.1006***	0.1061
2004	-0.5614	0.1014***	0.1273	0.4690	0.1029***	0.0896
2006	-0.6786	0.0686***	0.1647	0.6255	0.0681***	0.1468
2008	-0.5759	0.082***	0.1265	0.3795	0.0835***	0.0581
2010	-0.7135	0.0776***	0.1935	0.7002	0.0773***	0.1902
2012	-0.6420	0.0769***	0.1692	0.6116	0.0804***	0.1443
2014	-0.7970	0.074***	0.2171	0.7594	0.0728***	0.2072
2016	-0.8278	0.0665***	0.2397	0.7233	0.0666***	0.1963
2018	-0.7850	0.0689***	0.2414	0.8526	0.07***	0.2650
2021	-1.0556	0.0434***	0.4286	1.0696	0.046***	0.4075
2022	-0.9540	0.0532***	0.3243	0.9168	0.0553***	0.2958
2024	-0.8861	0.0558***	0.2872	0.9014	0.0573***	0.2821

*, **, and ** denote significance at the 90%, 95%, and 99% levels of confidence, respectively.

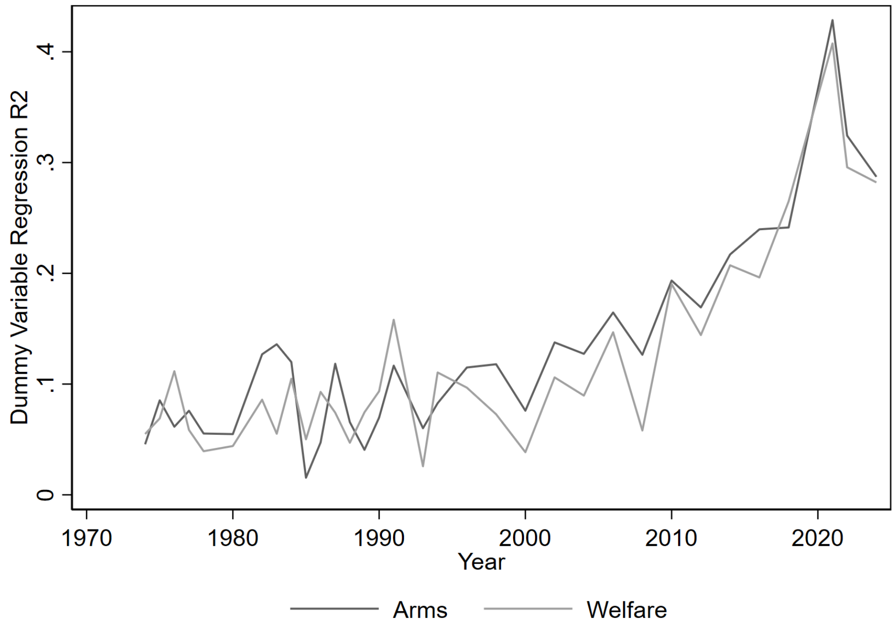

Figure 5. Graph of R^2s from dummy variable regressions. In each case, the dummy variable is 0 if the respondent stated they were "liberal" or "extremely liberal" and 1 if they stated they were "conservative" or "extremely conservative." The dependent variable ranged rated spending as 1 ("too little"), 2 ("about right") or 3 ("too much").

The results of the calibrations are graphed in Figure 6. Graphs of the mean absolute prediction error show the same pattern because the mean absolute prediction error for a normal posterior equals the posterior variance times $(2/\pi)^{0.5}$.

Figure 6 illustrates several important points. First, as we add observers with only one group type represented, the posterior variance falls rapidly to a limit that is strictly above zero and strictly higher than having more than one type represented. Second, for any given number of types, the posterior variance approaches a limit as we add observers, but the limits are always strictly above zero. Third, limits fall as we add types. Fourth, and possibly most important, when polarization is relatively small (i.e., 1974), the posterior variance limits are relatively close to zero regardless of the number of types. When polarization is high (i.e., 2021), the number of viewpoints (types) represented matters a lot. This shows it is much more important to incorporate multiple viewpoints when polarization is high.

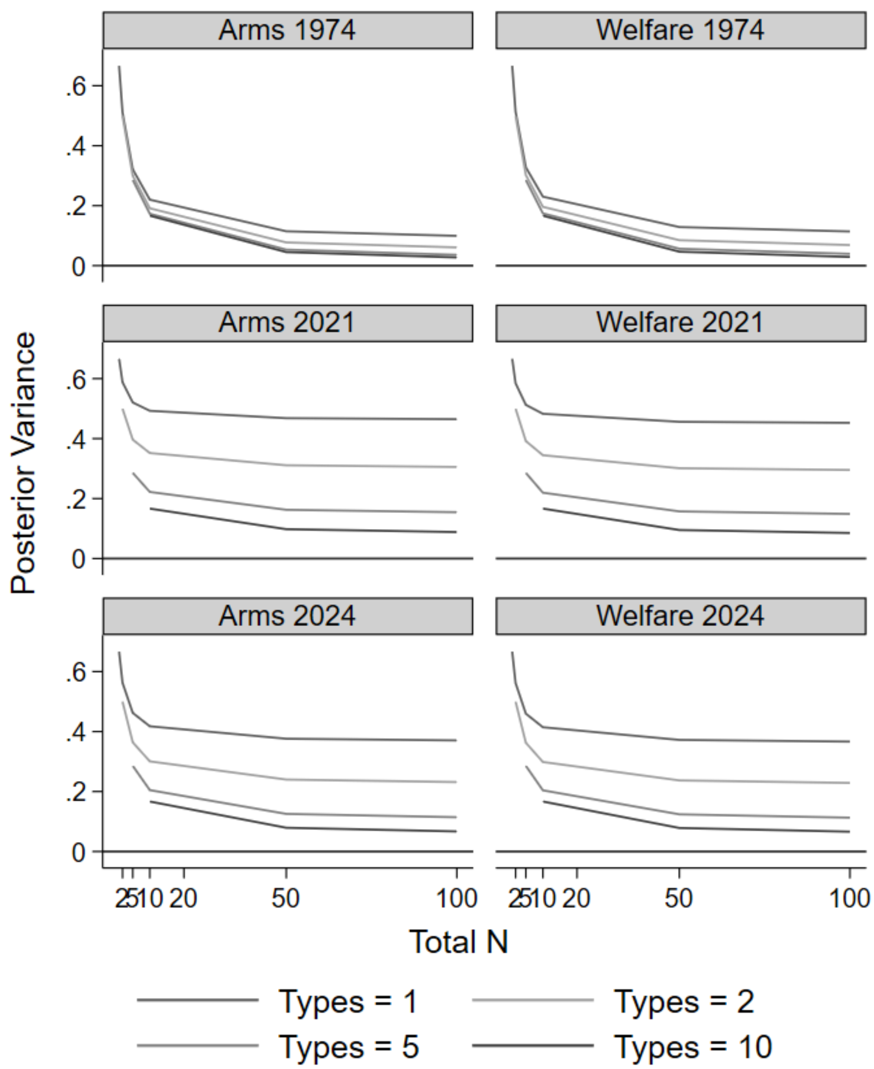

Figure 6. Posterior Variances based on the R^2 from dummy variable regressions in 1974, 2021 and 2024 and various combinations of numbers of types and observers.

6 Additional Observations

Consider another possible way around the limiting factor of two or few types, repeated signals given to observers. There are three possibilities.

First, suppose each observer gets signals in sequence each with a new individual error term. This means their type signal (what defines them as part of a group) stays the same and each new combined signal equals to $y_{i,j,\tau} = V + \eta_\tau + \epsilon_{i,j}$ where j indexes the observers signal number. The information in this set of n such signals for an type τ observer is exactly the same as having n type τ observers each with one signal. As in that case, the posterior variance converges to a strictly positive value and the posterior mean does not converge to the true value.

Second, suppose each observer gets signals in sequence with both new type and individual error terms. Each new combined signal equals $y_{i,j,\tau} = V + \eta_{\tau,j} + \epsilon_{i,j}$ where j indexes the observers signal number. This is effectively the same as having n types. The posterior distribution will converge to the true V. However, it also means that, for no apparent reason, the entire group becomes a new type with every new piece of information. Republicans may become Democrats; men become women; neo-Nazis become ultra-feminists, etc. You could also rig it so that individuals instead of entire groups change. But, neither scenario seems reasonable.

Third, suppose each observer gets signals in sequence with new individual error terms and the variance of η is zero. Each new combined signal equals $y_{i,j,\tau} = V + \epsilon_{i,j}$ where j indexes the observers signal number. This just becomes n observers with random noise. This would also converge, but there is no longer a group viewpoint. This conflicts with evidence that Democrats do consistently differ from Republicans in their initial viewpoints, men do differ from women, Nazis and feminists do differ.

A final possibility is that, using a deterministic or random decay function, people simply forget their types (their initial ηs through time). This would allow observers to become more homogeneous over time, slowly removing the impact of η_τ. Over time, this moves us to the third case above. This means, over time, groups coalesce into a single homogenous group. In today's society, this conflicts with the evidence that we are becoming more polarized, not less.

While none of these panel data solutions seem to solve the problem in a reasonable way, we can still say something about how the posterior distribution reacts to new information in the additional individual error terms case. Suppose that m is the number of types and N is the total number of observers. Rearranging equation 34 slightly gives:

$$\Lambda_{\text{data}} = \sum_{\tau=1}^{m} \frac{\frac{N}{m}\lambda_\eta\lambda_\epsilon}{\lambda_\eta + \frac{N}{m}\lambda_\epsilon} = \frac{N\lambda_\eta\lambda_\epsilon}{\lambda_\eta + \frac{N}{m}\lambda_\epsilon} = \frac{Nm\lambda_\eta\lambda_\epsilon}{m\lambda_\eta + N\lambda_\epsilon}, \tag{35}$$

where, again, N observers are equally divided among m types. It is easy to see that λ_{data} is increasing in m up to N, increasing in N, increasing in λ_η, and increasing in λ_ϵ.[12] All increase the precision of the data and, as a result, decrease the posterior variance and increase the sensitivity of the posterior mean to the data. In standard deviation terms, more polarization across groups (i.e., inter-group noise σ^2_η) and more division withing groups (i.e., intra-group noise, σ^2_ϵ) both increase the posterior variance and dampen the sensitivity of the posterior mean to the data. All of these relationships also hold when the panel is unbalanced.

As Figure 2 shows, adding more observers of a type with fixed N changes the sensitivity of the posterior to the data. Adding observers to the smaller group increases sensitivity, but at a rapidly declining rate. Figure 3 shows how fast the posterior mean and variance approach their limits assuming a balanced panel of observers as the number of observers increases. Equivalently, it shows how fast the posterior mean and variance converge if two observers, one of each type, each receive additional signals.

Unbalanced panels converge more slowly than balanced panels. But, because the updating function weights the smaller group more heavily, it does not matter as much as one might think. Figure 7 shows what happens if the panel is imbalanced with 1 D observer added after every 9 R observers. The unbalanced panel converges somewhat more slowly to the same limit.

A natural and intuitive way to incorporate a range of information in a board is to (1) make the board representative of the population and (2) equally weight the information of each board member. Similarly, a common method of polling expectations for forecasting is to make the sample representative of the population and equally weight all individual forecasts. While this helps, it is suboptimal unless the fractions of types on the board (or poll) just happen to be equal. This observation arises from Figure 2, Panel A and Observation 1. It is fairly easy to show the extent

[12] Simply take partial derivatives with respect to each argument. For example, the partial with respect to N simplifies to $\frac{m^2\lambda^2_\eta\lambda_\epsilon}{(m\lambda_\eta+N\lambda_\epsilon)^2} > 0$. The other partials are similar.

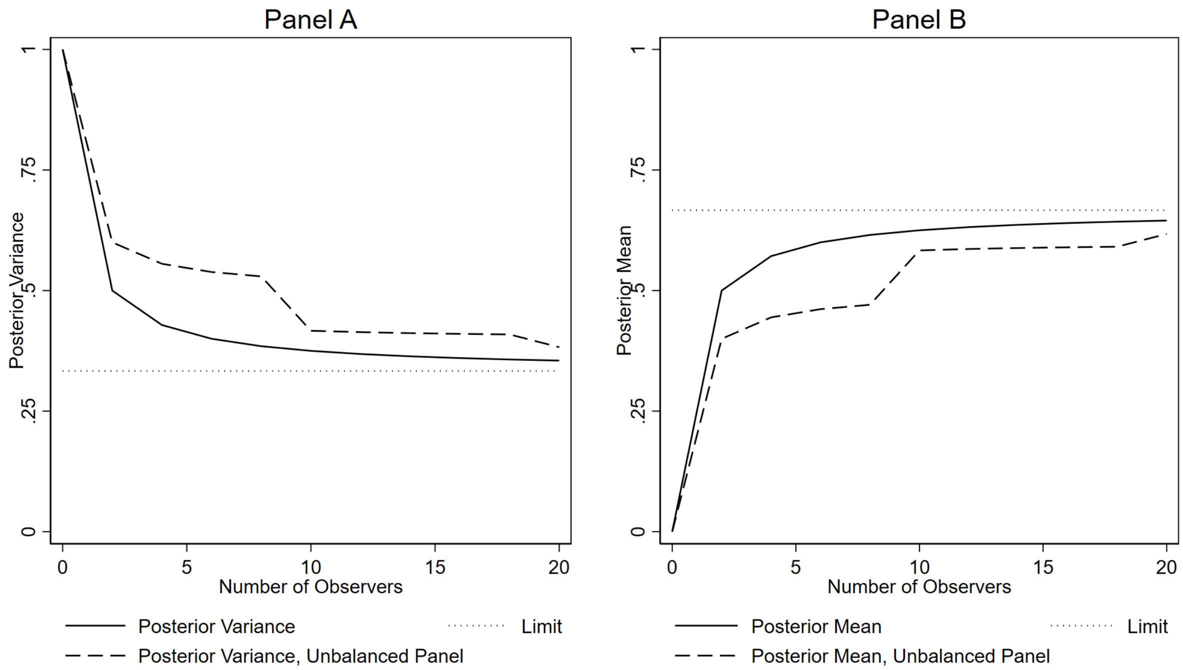

Figure 7. Posterior variance and mean when $\sigma_V^2 = \sigma_\eta^2 = \sigma_\epsilon^2 = 1$ and 1 D is added after very 9 R observers up to 20 observers total versus up to 20 balanced observers. Panel A shows the posterior variance. Panel B shows the posterior mean for a mean data signal of 1.

of the suboptimality by reworking equations 13 through 16 assuming equal weighting instead of optimal weighting.

As before, let n_D denote the number of type-D observers and n_R the number of type-R observers, with $n = n_D + n_R$. The equally weighted mean is:

$$\tilde{y} = \frac{1}{n} \sum_{i=1}^{n} y_i. \tag{36}$$

The variance of $\tilde{y}$ is affected by the n_D observers that have a single common η, the n_R observers that have a different, but commen, η and the $n = n_D + n_R$ independent ϵ's giving:

$$\text{Var}(\tilde{y} \mid V) = \frac{n_D^2 + n_R^2}{n^2} \sigma_\eta^2 + \frac{\sigma_\epsilon^2}{n}, \tag{37}$$

which is minimized only when $n_D = n_R$. Alternatively, in precision notation:

$$\text{Var}(\tilde{y} \mid V) = \frac{n_D^2 + n_R^2}{n^2}\frac{1}{\lambda_\eta} + \frac{1}{n}\frac{1}{\lambda_\epsilon}. \tag{38}$$

Inverting this and a little algebra gives:

$$\Lambda_{\text{EW}} = \frac{n^2 \lambda_\eta \lambda_\epsilon}{n\lambda_\eta + (n_D^2 + n_R^2)\lambda_\epsilon}. \tag{39}$$

Then, as before:

$$\text{Var}_{\text{EW}}(V \mid \tilde{y}) = \frac{1}{\lambda_V + \Lambda_{\text{EW}}}. \tag{40}$$

Comparing (39) to the optimal posterior variance in equation 15 shows that the equal-weight posterior variance is strictly larger for every allocation with $n_D \neq n_R$. Intuitively, equal weighting forces the posterior to behave as if all signals contribute equally to the precision of the data. This causes the posterior to overreact to the over-represented type and underweight the under-represented type. Figure 8 shows the impact of this by comparing the optimally weighted posterior variance from Figure 2 Panel A to the equally weighted posterior variance.

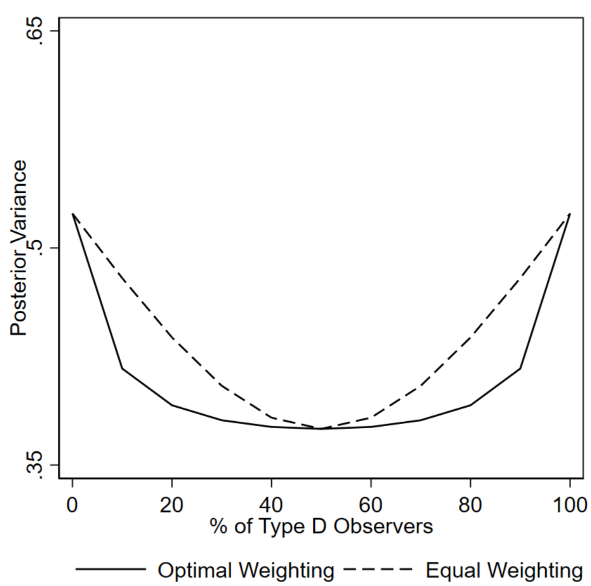

Figure 8. Posterior variance with optimal weights and equal weights for 10 observers with 0 to 10 of one type when $\sigma_V^2 = \sigma_\eta^2 = \sigma_\epsilon^2 = 1$.

Thus, the seemingly "fair" or "natural" rule of adding more panelists and weight them equally is suboptimal. Posterior precision improves fastest not by equal weighting, but by accounting for the contribution of each observer to the overall data precision. This shows clearly the point from Proposition 3: only balancing types *and* weighting them optimally minimizes posterior uncertainty.

7 Discussion

This paper applies straightforward Bayesian analysis to a common situation to gain some interesting insights. Consider the makeup of a panel of individuals tasked with uncovering information to make informed decisions. Suppose the panel members (called "observers" here) include groups (or "types" here) of people that share some common information or beliefs within their type (i.e., common, but unobservable, baseline viewpoints). However, these initial viewpoints differ across types. Starting with the common viewpoint associated with their type, each observer brings some additional private information to the table. In this case, a Bayesian analysis shows the following:

1. If there are two types (groups) of panel members (e.g, Democrat and Republican), then:
    (a) While a highly unbalanced panel may reveal unbiased information, the residual uncertainty about the truth is high.
    (b) A more balanced panel can still reveal unbiased information, while reducing residual uncertainty.
    (c) More panel members can also reduce residual uncertainty to some degree. But, with only two types of panel members, there is a limit to how much residual uncertainty can be reduced regardless of the number of panel members or their balance.
    (d) The (correctly) weighted means of the observations does not give an estimator that converges to the true V.
2. If there are many types (groups) of panel members that each bring their own unique baseline beliefs and viewpoints, then:
    (a) Each member of a panel of a new type brings additional unique information to the panel.
    (b) When there are many types, the posterior mean is more responsive to the observed data than when there are fewer types and it converges to the true V.
    (c) With a sufficient number of types, the residual variance can be reduced to zero.

Thus, the key to the most accurate information, which presumably leads to the best decisions, lies not in more or more balanced representation by a small number of types. Instead, it arises from including a wide range of types with independent initial underlying perspectives, beliefs, and viewpoints. Unfortunately, often, the number of types may be quite limited. Further, none of these results require any bias of any kind from individual observers or types.

References

Bauman, K. P., and G. Geher. 2002. We think you agree: The detrimental impact of the false consensus effect on behavior. *Current Psychology* 21:293–318.

Berg, J. E., J. Geweke, and T. A. Rietz. 2010. Memoirs of an indifferent trader: Estimating forecast distributions from prediction markets. *Quantitative economics* 1:163–186.

Chait, R. P., T. P. Holland, and B. E. Taylor. 1993. *The effective board of trustees*. Bloomsbury Publishing PLC.

Cinelli, M., G. De Francisci Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini. 2021. The echo chamber effect on social media. *Proceedings of the national academy of sciences* 118:e2023301118.

Coval, J. D., and T. J. Moskowitz. 1999. Home Bias at Home: Local Equity Preference in Domestic Portfolios. *The Journal of Finance* 54:2045–2073. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/0022-1082.00181.

Croson, R., and U. Gneezy. 2009. Gender differences in preferences. *Journal of Economic literature* 47:448–474.

Donaldson, J. R., N. Malenko, and G. Piacentino. 2020. Deadlock on the board. *The Review of Financial Studies* 33:4445–4488.

Feinstein, B. D., and D. J. Hemel. 2024. The market value of partisan balance. *Nw. UL Rev.* 119:1201.

Forsythe, R., F. Nelson, G. R. Neumann, and J. Wright. 1992. Anatomy of an experimental political stock market. *The American Economic Review* pp. 1142–1161.

Forsythe, R., T. A. Rietz, and T. W. Ross. 1999. Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior & Organization* 39:83–110.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 1995. *Bayesian data analysis*. Chapman and Hall/CRC.

Gruca, T. S., and T. A. Rietz. 2021. The 2020 (re) election according to the Iowa electronic markets: politics, pandemic, recession, and/or protests? *PS: Political Science & Politics* 54:86–90.

Gruca, T. S., and T. A. Rietz. 2025. Iowa electronic markets: Forecasting the 2024 US presidential election. *PS: political science & politics* 58:258–266.

Harris, C. R., and M. Jenkins. 2006. Gender differences in risk assessment: Why do women take fewer risksthan men? *Judgment and Decision making* 1:48–63.

Hemel, D., and B. D. Feinstein. 2018. Partisan Balance with Bite. *Columbia Law Review* 118:9.

Kang, J.-K., S. Kim, and S. Oh. 2020. Board diversity, director dissent, and monitoring effectiveness. *Work. Pap., Nanyang Bus. Sch./Baruch Coll./Peking Univ* .

Kim, I., C. Pantzalis, and J. C. Park. 2013. Corporate boards' political ideology diversity and firm performance. *Journal of Empirical Finance* 21:223–240.

Knyazeva, A., D. Knyazeva, and L. Naveen. 2021. Diversity on corporate boards. *Annual Review of Financial Economics* 13:301–320.

Lee, J., K. J. Lee, and N. J. Nagarajan. 2014. Birds of a feather: Value implications of political alignment between top management and directors. *Journal of Financial Economics* 112:232–250.

Lewis-Beck, M. S., and A. Skalaban. 1989. Citizen forecasting: Can voters see into the future? *British Journal of Political Science* 19:146–153.

Love, B. C., Ł. Kopeć, and O. Guest. 2015. Optimism bias in fans and sports reporters. *Plos one* 10:e0137685.

Peterson, E., S. Goel, and S. Iyengar. 2021. Partisan selective exposure in online news consumption: evidence from the 2016 presidential campaign. *Political Science Research and Methods* 9:242–258.

Ross, L., D. Greene, and P. House. 1977. The "false consensus effect": An egocentric bias in social perception and attribution processes. *Journal of experimental social psychology* 13:279–301.

Schwartz-Ziv, M. 2017. Gender and board activeness: The role of a critical mass. *Journal of Financial and Quantitative Analysis* 52:751–780.

Surowiecki, J. 2005. *The wisdom of crowds.* Vintage.

Yockey, J. W. 2024. Recalibrating Public University Governing Boards. *The Georgetown Journal of Law & Public Policy* 22:773.