Misinformed or Uninformed? The Prevalence and Consequences of Certainty in Political Misperceptions

Brian Guay

Postdoctoral Associate
Sloan School of Management & Department of Political Science, MIT

August 29, 2022

Abstract

There is widespread concern today that much of the public is misinformed, holding factually inaccurate beliefs that they believe to be correct. However, few studies to date have measured the certainty with which political beliefs are held. Those that do use methods that cannot meaningfully distinguish actual beliefs from uncertain guesses and have led to differing conclusions about the prevalence of certainty in inaccurate beliefs. In this paper I introduce a new measure of certainty that provides meaningful context to existing self-report measures of certainty and adjusts for differential item functioning (i.e., differences in how respondents use response scales). I show that past work exaggerates the degree to which the public is misinformed and that inaccurate beliefs are four times more likely to represent uncertain guesses than actual beliefs. I then provide the first empirical test of the widespread expectation that certain beliefs are more likely to resist correction and influence attitudes than uncertain ones. The findings presented here have implications for how researchers interpret findings from the growing body of research that documents and attempts to correct misperceptions.

A common assumption across theories of political behavior holds that citizens use facts to inform their political attitudes and decision-making. As Carpini and Keeter (1996, pg. 295) put it, "Facts provide a foundation for deliberation about larger issues" and "allow individuals and groups with widely varied experiences and philosophies to have some common basis of comparison." Amid growing concerns over fake news, online echo chambers, and mainstream conspiracy theories, scholars have paid particular attention in recent years to the misperceptions, or inaccurate factual beliefs, that people hold about politics. A growing body of research finds that people hold misperceptions about everything from foreign military conflict (Nyhan and Reifler, 2010; Berinsky, 2007) and the economy (Conover et al., 1986; Lawrence and Sides, 2014) to the size of racial and ethnic minority groups (Sides and Citrin, 2007; Hopkins et al., 2019; Guay et al., 2020).

Political scientists have expressed particular concern in recent decades over the misinformed, who feel certain that their inaccurate beliefs are correct. As Kuklinski et al. (2000) put it, "If people do not hold factual beliefs at all, they are *uninformed*...but if they firmly hold beliefs that happen to be wrong, they are *mis*informed—not just in the dark, but wrongheaded." (pg.793, emphasis in the original). Kuklinski and colleagues demonstrate that incorrect responses on surveys often reflect confidently-held beliefs rather than mere guesses, with people often expressing *more* certainty in inaccurate beliefs than accurate ones. These findings raise normative concerns about the capacity of citizens to meaningfully engage in politics and hold elected representatives accountable. Specifically, scholars have hypothesized that inaccurate beliefs held with certainty are more likely to influence attitudes and resist correction than those held without certainty (Kuklinski et al., 2000; Jerit and Zhao, 2020; Hochschild and Einstein, 2015; Bode and Vraga, 2015).

However, the growing literature on misperceptions and misinformation that has emerged since largely conflates the misinformed with the uninformed, overlooking the certainty with which beliefs are held. Three recent studies identify this measurement issue, but come to different conclusions about the degree to which inaccurate beliefs are held with certainty

(Pasek et al., 2015; Graham, 2020; Lee and Matsuo, 2018; Thorson, 2015). One likely reason for these inconsistent findings in the literature is that existing methods of measuring certainty cannot meaningfully distinguish inaccurate beliefs held with certainty from incorrect guesses. Past work asks survey respondents to self-report their certainty on ordinal scales (e.g., 'very certain,' 'somewhat certain'), which require researchers to specify a threshold (e.g., 'very certain') above which to classify beliefs as certain. Since these scales are meaningful only in relative terms (e.g., 'very certain' is more certain than 'somewhat certain'), these thresholds are difficult to interpret substantively. This problem is amplified when respondents interpret the same response scale differently, mapping identical latent levels of certainty onto different response categories (i.e., differential item functioning; King et al., 2003; King and Wand, 2007; Hopkins and King, 2010). To date, each of the four published studies that have measured certainty in political beliefs have used different scales, specified different thresholds, and have not accounted for differential item functioning (Kuklinski et al., 2000; Pasek et al., 2015; Graham, 2020; Lee and Matsuo, 2018).

Our understanding of the extent to which inaccurate beliefs are held with certainty is further limited by the range of beliefs for which certainty has been measured. Past work has measured certainty in only a small fraction of the many misperceptions documented over the past two decades, with most studies measuring certainty in either a single issue (e.g., welfare or healthcare, Kuklinski et al., 2000; Pasek et al., 2015) or general knowledge items that are not expected to influence political attitudes (Lee and Matsuo 2018, but see Graham 2018). Moreover, while scholars have raised normative concerns over the consequences of being misinformed (Kuklinski et al., 2000; Jerit and Zhao, 2020; Hochschild and Einstein, 2015; Bode and Vraga, 2015), no empirical work has evaluated these empirically. Specifically, no work to date has tested the widespread hypotheses that inaccurate beliefs held with certainty are more likely to influence attitudes and resist correction. Thus, while the emerging literature on misperceptions and misinformation often claims that certainty in inaccurate political beliefs is widespread and consequential, evidence underlying both assumptions is

limited.

This paper examines the prevalence and consequences of being misinformed about a wide range of issues, using a method of measuring certainty that adds substantive meaning to ordinal response scales used to measure certainty in past work. This enables a more meaningful distinction between the uninformed and misinformed, as well as the ability to control for respondents' differential use of the same response categories. In addition to self-assessing their own level of certainty, survey respondents assessed the degree of certainty inherent in four anchoring vignettes. Each vignette expressed a different level of certainty, enabling self-reported levels of certainty to be anchored to common benchmarks (King et al., 2003; King and Wand, 2007; Wand, 2013). Using this method, I offer a more meaningful assessment of the degree to which many of the misperceptions documented over the past two decades—including the size of the immigrant population, demographic composition of the U.S., education spending, welfare, foreign aid, and stereotypes about Democrats and Republicans—reflect inaccurate beliefs or uncertain guesses. I then provide the first empirical test of the two most widely hypothesized consequences of belief certainty. First, I assess the degree to which certainty moderates the strength of the relationship between beliefs and attitudes about immigration, education policy, welfare, and members of the opposing political party. Second, I examine whether certain beliefs are more resistant to correction in an experiment by randomly assigning half of respondents to receive correct information about immigration and education policy.

The results presented here suggest that past work has exaggerated both the prevalence and consequences of being politically misinformed. I show that inaccurate beliefs are approximately four times more likely to represent uncertain guesses than firmly held beliefs, with only 12% of the sample classified as misinformed. While the most inaccurate beliefs are sometimes held with the highest levels of certainty, I show that relationship is driven by a relatively small segment of the population and may be the result of expressive responding. Moreover, I find little empirical support for two of the most widely hypothesized conse-

quences of holding beliefs with certainty. Beliefs held with certainty were no more strongly related to attitudes nor resistant to experimental interventions aimed at correcting them. In fact, in two cases *uncertain* beliefs are more strongly related to attitudes than certain ones. These findings have implications for how researchers interpret inaccurate responses to knowledge questions reported on surveys. While the growing body of research on misperceptions and misinformation often assumes these responses reflect true underlying beliefs, the results presented here suggest that they are far more likely to represent mere uncertain guesses.

Certainty in Political Beliefs

A central question in political science research concerns the extent to which citizens are knowledgeable about politics. A large body of evidence accumulated over the past half century suggests much of the public exhibits low levels of political sophistication (Converse, 1964; Campbell et al., 1960; Bartels, 1996; Carpini and Keeter, 1996). When citizens are informed, this knowledge typically extends to only a limited set of personally important issues (Converse, 1964; Key, 1961). Carpini and Keeter (1996) find that less than half of Americans could define the terms liberal and conservative, report the length of the House of Representatives term of office, or could identify their senators. Responses to more specific questions paint an even grimmer picture of political sophistication. For instance, less than 25% of respondents were able to define two first amendment rights or identify the three branches of government.

As Kuklinski et al. (2000) note, however, inaccurate responses to these types of survey questions can reflect two distinct states of knowledge: the uninformed, who are both inaccurate and aware of their lack of knowledge, and the misinformed, who are similarly inaccurate but confident they are correct. Belief certainty, then, serves as a second dimension of knowledge. While early work on political sophistication centered on a single knowledge dimension,

classifying individuals as informed or uninformed based on the *accuracy* of their beliefs, Kuklinski and colleagues introduced belief *certainty* as a second dimension of knowledge. This second dimension differentiates an individual who guesses with little certainty that, for instance, 30% of the population is foreign born from someone who confidently states the same belief as fact. Clearly neither is individual informed (the correct response is closer to 14%), but the latter is more likely to represent a belief than a random guess.

Widespread inaccurate beliefs held with certainty raise normative concerns about the ability of citizens to meaningfully engage in politics and hold elected representatives accountable. Specifically, scholars have hypothesized two primary consequences of holding inaccurate beliefs with certainty (Kuklinski et al., 2000; Jerit and Zhao, 2020; Hochschild and Einstein, 2015; Flynn et al., 2017). First, inaccurate beliefs held with higher levels of certainty should exert greater influence on political attitudes and policy preferences. As Kuklinski et al. (2000) explain, "The problem is...not that people simply lack information, but that they firmly hold the wrong information–and use it to form preferences" (pg. 792). If attitudes are constructed through a process of sampling from relevant beliefs, we should expect actual beliefs, not uncertain guesses reported on surveys, to influence attitudes. At the individual level, informing attitudes with inaccurate beliefs may result in deviations from self-beneficial preferences or 'correct' voting (Lau and Redlawsk, 1997). At the aggregate level, this can result in biased public opinion, particularly if beliefs held with certainty deviate systematically from the truth (Kuklinski et al., 2000; **?**), whether because they originate from the same external source (e.g., the media, disinformation spread online) or some internal cognitive process (Guay et al., 2020).

Second, beliefs held with greater certainty are hypothesized to be more resistant to correction than those held with less certainty (Kuklinski et al., 2000; Jerit and Zhao, 2020; Flynn et al., 2017). Normative democratic ideals hold not only that citizens will hold preferences that are rooted in factual beliefs, but that they will update those beliefs and preferences when confronted with credible information that conflicts with their prior beliefs (Achen and

Bartels, 2017; Nyhan and Reifler, 2010; Guay and Johnston, 2021). A person who, for instance, states unequivocally that 40% of the population is foreign born is generally expected to be less willing to update their belief in response to correct information than someone who holds the same belief with high levels of uncertainty.

Methodological Limitations of Existing Research

Despite these concerns, the growing body of research on misinformation and misperceptions often conflates certain beliefs with uncertain guesses. Indeed, only four studies to date have measured certainty in political beliefs. These studies arrive at different conclusions about the degree to which inaccurate beliefs are held with certainty, and mixed findings within single studies are common. For instance, while Kuklinski et al. (2000) is frequently cited as showing that inaccurate beliefs about welfare are held with high amounts of certainty, even more so than accurate ones, this relationship is apparent for only two of the six factual beliefs measured in the study. For two others, inaccurate beliefs are held with *less* certainty than accurate ones, and for the two remaining beliefs the relationship between accuracy and certainty is unclear.[1] Pasek et al. (2015) and Graham (2020) find evidence of the opposite relationship between certainty and accuracy, reporting that inaccurate beliefs about a range of issues and policies are held with *less* certainty than accurate ones. Findings from a fourth study are decidedly mixed, with inaccurate beliefs sometimes held with less certainty than accurate beliefs and other times held with more certainty (Lee and Matsuo, 2018).

One reason for this inconsistency likely lies in how beliefs and certainty are measured. First, there is wide variation in the types of beliefs measured by past work. To date, each

[1]The positive relationship between accuracy and certainty is apparent for two of the six factual beliefs measured in the study—the share of the population on welfare and the size of the average annual welfare payment. For two other facts, beliefs are often held with *less* certainty than accurate ones and for the two remaining beliefs the relationship between accuracy and certainty is more ambiguous: incorrect responses that overestimate these quantities are held with greater certainty, while incorrect responses that underestimate these quantities are held with less certainty. Moreover, the data used to draw these conclusions contains survey responses from only 273 individuals, making it still more difficult to draw inferences about the relationship between accuracy and certainty. Indeed, it is possible that, if statistical uncertainty is taken into account, some of these associations are not statistically significant.

study has measured certainty in different beliefs, and often one single issue at a time. For instance, Kuklinski et al. (2000) measure beliefs about welfare in the U.S. and Pasek et al. (2015) measure beliefs about whether specific policy provisions were included in the 2010 Affordable Care Act (e.g., requiring Americans to show an ID at hospitals, the creation of 'death panels'). Another study measures certainty expressed in responses to general political knowledge questions (e.g., identifying the party affiliations of political elites) (Pasek et al., 2015), though as Kuklinski et al. (2000, pg. 793) note, there is little reason for certainty to matter much for these types of beliefs given that they are largely orthogonal to political preferences.[2] Graham (2020) measures certainty in the widest range of beliefs to date, asking respondents to rate whether certain quantities, such as the crime rate, federal deficit, and immigrant apprehensions have increased or decreased in recent years.

Second, these studies use different methods of eliciting beliefs. Three of the four studies measure beliefs in a way that permits only a binary classification of beliefs as accurate or inaccurate, such as asking whether the Affordable Care Act creates death panels (Pasek et al., 2015), whether the crime rate has increased in recent years (Graham, 2020), or to which party the prime minister belongs (Lee and Matsuo, 2018). These types of questions differ from those used by Kuklinski et al. (2000), and much of the literature on misperceptions, which measure not only whether beliefs are inaccurate, but also the *extent* to which they are inaccurate. Such questions typically come in the form of asking respondents to estimate the size of politically-relevant quantities, such as the proportion of the population that receives welfare (Kuklinski et al., 2000), the share of the budget dedicated to foreign aid (Gilens, 2001; Scotto et al., 2017), and the size of racial and ethnic groups (Sides and Citrin, 2007; Guay et al., 2020; Wong, 2007). By accounting for the degree to which a belief is inaccurate, these questions avoid conflating beliefs that are very far from the truth (e.g., confidently overestimating the share of the population that is foreign-born by 30 percentage points) with

[2]Of the nine knowledge questions asked of respondents, only two measure policy-relevant beliefs: identifying the largest growing immigrant group and the unemployment rate in the U.K. Six others ask about political parties in the U.K and who belongs to them, and another concerns how a proposed voting system functions.

those that just miss the mark. These questions also enable a richer comparison of beliefs and attitudes, since the direction and size of the is often correlated with the preferences (e.g. Sides and Citrin, 2007; Gilens, 2001).

Third, and perhaps most importantly, past studies use methods of measuring certainty that make it difficult to draw meaningful inferences about the prevalence of the misinformed. Past work measures certainty by asking respondents to report how certain they are on ordinal scales ranging from, for instance, 'not at all sure' to 'extremely sure' (Pasek et al., 2015). Researchers then specify one category as the threshold with which to distinguish between certain and uncertain beliefs, and thereby the misinformed and uninformed, respectively. However, because responses recorded on these ordinal scales have only relative meaning, it is not at all clear where this threshold should be, nor how to interpret the difference between the uninformed and misinformed after it has been drawn. For instance, 'very sure' clearly reflects higher certainty than 'somewhat sure,' but which (if any) should be considered to reflect a certain belief?

This threshold directly influences the claims researchers draw about the prevalence of the misinformed, but to date there has been little discussion or agreement about where to draw it. For instance, Kuklinski et al. (2000) classifies response categories 'very confident' and 'fairly confident' as certain and 'not at all confident' and 'a little confident' as uncertain. Pasek et al. (2015) consider responses to reflect certainty if they are 'extremely sure' or 'very' sure, but not when they are 'moderately sure,' 'slightly sure,' or 'not sure at all.' Researchers measuring other aspects of public opinion, such as issue positions, often avoid this problem by using bipolar 'Likert' scales, for which there is an apparent and intuitive demarcation between two extremes of a latent construct (e.g., 'neither agree nor disagree' on agree-disagree scales). However, past work has justifiably avoided measuring certainty in this way because certainty does not map onto bipolar scales. It is unclear, for example, what it means to be somewhat certain versus somewhat uncertain, or what it means to be neither certain nor uncertain.

The task of using ordinal response scales to meaningfully distinguish between certain and uncertain beliefs is further complicated by the likelihood that respondents use certainty response scales differently. Differential Item Functioning, or how respondents map levels of a latent construct onto response categories, is possible anytime an ordinal scale is used to quantify a latent construct (King et al., 2003; King and Wand, 2007; Wand, 2013; Aldrich and McKelvey, 1977), but poses a particular challenge for distinguishing between the certain and uncertain. For instance, two respondents may hold a belief with identical levels of certainty, but report this certainty using different response options, say 'somewhat sure' and 'very sure.' If the researcher classifies responses above 'somewhat sure' as certain, the two individuals will be classified differently despite holding a belief with identical levels of certainty.

Taken together, the limited number of studies that measure certainty in political beliefs, and the limited range of issues, conflicting conclusions, and methodological constraints of those that do, suggest that less is known about the prevalence and consequences of being misinformed than previously thought. Thus, there appears to be substantial value in studies that measure certainty across a wide range of domains, employ methods that enable both the measurement of the extent to which beliefs are accurate and meaningful empirical distinctions between the informed, uninformed, and misinformed, and examine the consequences of holding beliefs with certainty.

The central aim of this paper is to take up this task. I measure the certainty with which commonly measured misperceptions are held, using a method that adds substantive meaning to ordinal response scales often used to measure certainty. This enables both a more meaningful distinction between the informed, uninformed and misinformed, as well as the ability to control for respondents' differential use of the same response categories. I then provide the first empirical test of two of the most frequently hypothesized consequences of belief certainty.

Design

I conducted a survey on a sample of 2,744 adults aged 18 years or older living in the U.S. in April 2020. Respondents were recruited using Lucid, a platform that connects researchers to a pool of survey respondents from multiple online panels. Recent work finds that samples drawn from Lucid closely match the demographic and political characteristics of the U.S. population, are less professionalized and politically sophisticated than other commonly used non-probability samples, and replicate experimental findings conducted on other samples (Coppock and McClellan, 2019). Quota sampling was used to obtain a sample that closely matched the Census Bureau's 2016 American Community survey on race and ethnicity, gender, age, geographic region.[3] Respondents were screened on several quality benchmarks, including two attention checks, which resulted in a final analytical sample of 2,028 respondents.[4]

Measuring Misperceptions

Three criteria were used to select the beliefs that were measured on the survey. First, I selected beliefs that are reported on interval scales, which enable me to measure the amount of error in a belief, rather than merely whether it is right or wrong. These questions take the form of asking respondents to estimate a quantity, such as the share of the population belonging to a certain group or budget allocated toward a certain program (Lawrence and Sides, 2014; Gilens, 2001; Sides and Citrin, 2007; Hopkins et al., 2019; Guay et al., 2020). Second, I included beliefs for which prior work on misperceptions has documented an association with political attitudes, including many that studies have attempted to correct. Finally, beliefs were selected that covered a wide range of domains, including the size of salient demographic groups, government spending, public policy, and partisan stereotypes.

Respondents estimated the percent of the budget that goes to foreign aid and the share

[3]This study was approved by the Duke University Institutional Review Board.

[4]The attention checks and additional exclusions based on data quality are discussed in greater detail in the Appendix (pgs. 3-4).

of funding for K-12 public schools that comes from the federal government, both of which prior work finds Americans overestimate dramatically (Gilens, 2001; Carpini and Keeter, 1996; Scotto et al., 2017; Henderson et al., 2019). Respondents were also asked to estimate the proportion of Americans on welfare, one of the two factual beliefs for which Kuklinski et al. (2000) find inaccurate beliefs are held with greater certainty than accurate ones. Given that beliefs about what is considered welfare vary widely in the population (Mettler, 2011), respondents were specifically asked to estimate the share of the population that receives food stamps.[5] Four additional quantities related to the size of racial, ethnic, and religious groups in the U.S., which are among the most cited instances of political misperceptions: the share of the population that is foreign-born (Hopkins et al., 2019; Sides and Citrin, 2007; Citrin and Sides, 2008), Black, White, and Christian (Alba et al., 2005; Nadeau et al., 1993; Duffy, 2018). Finally, respondents were asked to estimate two partisan stereotypes documented by Ahler and Sood (2018): the percentage of Democrats who are atheist/agnostic and Republicans who make over $250,000 a year. Each of the estimated quantities is reported in Table 1, along with the actual size of the quantity and its mean/median estimate.

For each question, respondents were instructed to enter a number between 0 and 100 in a text box that was proceeded by a '%' symbol. Given the potential for online survey respondents to look up correct answers (Clifford and Jerit, 2016), respondents were explicitly instructed not to look up the correct response. Specifically, respondents were told that we were just interested in their best guess and that there was no need to look up the answer to any question.[6] The full wording of all survey items is included in the Appendix.

[5] 36 million received SNAP and 6.5 million received WIC in 2019, totaling 42.5 million, about 13% of the population. To avoid the potential that respondents' estimates were influenced by emerging financial assistance programs during the Covid-19 pandemic, respondents were asked to estimate the percentage of the U.S. population that received food stamps in 2019.

[6] The instructions read: "Next, we're going to ask you some questions about politics and the U.S. population. We're just interested in your best guess for each question. There is no need to look up the answer to any question. After each question you'll be asked how sure you are about your answer."

Measuring Certainty

Respondents reported their level of certainty immediately after estimating each quantity. Following prior work, certainty was measured on an ordinal scale (e.g., Kuklinski et al., 2000). Specifically, respondents were asked "How sure are you about this" and chose between six options: extremely sure, very sure, moderately sure, a little sure, not very sure, not at all sure. As discussed above, these types of ordinal scales make it difficult for researchers to draw meaningful distinctions between the uninformed and misinformed and are likely used differently by respondents in a way that past studies do not observe and therefore cannot control for (i.e., differential item functioning).

To provide substantive information with which to interpret respondents' certainty responses and control for differential item functioning, respondents also responded to four anchoring vignettes. Anchoring vignettes document otherwise unobserved variation in how respondents use a particular response scale by asking them to report the level of a construct, in this case certainty, exhibited in hypothetical scenarios on the same response scale. In other words, anchoring vignettes mitigate the problem of simultaneous variation in both the latent level of a construct and use of the response scale on which the construct is measured by holding the latter constant across respondents. Anchoring vignettes are often discussed in the context of standardizing responses to cross-cultural surveys, since differential item functioning is especially likely to occur in surveys of respondents from multiple countries, though as Wand (2013, pg. 249) notes "the problem of inter-personally incomparable survey responses may exist anytime a question uses ordinal categories that are subjectively defined." Compared to alternative methods of measuring certainty employed in research on finance and decision-making (e.g., Enke and Graeber, 2019; Haran et al., 2010), anchoring vignettes provide additional context while introducing relatively little additional complexity for respondents or cost for researchers. Conversely, anchoring vignettes require only a few additional questions that differ from other common survey questions only in that they measure characteristics of others rather than characteristics of oneself (i.e., second-order vs.

first-order beliefs).

Unlike many constructs of interest in social science research, certainty in factual beliefs can be easily articulated in tangible and precise terms: the probability that a belief is correct. This greatly simplifies the task of designing vignettes describing different levels of certainty. Indeed, the most difficult facet of designing anchoring vignettes is articulating different levels of a latent construct in the vignettes such that respondents perceive the ordinal ranking of the construct in the vignettes equivalently. For instance, vignettes used to measure political efficacy (King et al., 2003; Wand, 2013) require the assumption that all respondents interpret a person who influences local policy-making by writing to their local legislator is more efficacious than another person who votes in a local election. The task of describing varying levels of belief certainty in a way that is ranked equivalently across respondents is comparatively easier given that certainty can be articulated, and is defined, as the probability of being accurate. The use of probability thresholds to describe latent levels of certainty in the vignettes also enables a more substantive interpretation of respondents' self-assessments.[7] For instance, respondents can be classified as being less than 10% or more than 90% sure that their response is correct.[8]

The design of the vignettes followed best practices prescribed by King et al. (2003) and Hopkins and King (2010), including using multiple vignettes, having all respondents complete each of the vignettes, and placing vignettes prior to self-assessments in order to train respondents on how to use the response categories.[9] All respondents answered questions

[7]Indeed, subjective certainty is sometimes elicited as a probability of being correct in other fields (e.g. Haran et al., 2010) This approach is intentionally avoided here due to concerns that misestimation error would co-vary across estimates of a quantity (e.g., the percentage of the U.S. population that is foreign born) and estimates of the probability that estimate is correct due to the domain-general cognitive errors people make when estimating any proportion or probability (Guay et al., 2020; Gonzalez and Wu, 1999)

[8]Of course, there is still room for between-subjects variation in interpretation the magnitude of these thresholds, even if it is unlikely that the ordering is subject to variation. But this challenge is faced by the existing method of using ordinal response scales to measure certainty, which does not account for differential item function or allow for any meaningful interpretation of the thresholds used to classify beliefs as certain.

[9]The anchoring vignettes appeared immediately before respondents completed the first set of quantity estimates in order to avoid unevenly biasing self-assessments. If the vignettes appeared after eliciting prior estimates and levels of certainty, they might influence posterior estimates and certainty in ways they did not for prior estimates. If the vignettes appeared after the posterior estimates, responses to the vignettes may be affected for respondents who received correct information in the experiment. Placing the vignettes

about the degree of certainty held by four individuals estimating the percentage of the U.S. population with a four-year college degree. The instructions read: "Next you'll see descriptions of 4 people who are asked the same quiz question. Some people are very sure that their answer is correct, others are not. It's your job to report how sure you think each person is. Jordan, Alex, Jamie, and Riley are each asked to estimate how many adults in the U.S. have a 4-year college degree." The instructions were followed by four questions describing the level of certainty with which each individual made their estimate. For instance, "Jordan thinks there is about a **10%** chance her estimate is close to the correct answer. How sure is Jordan about her answer?" Following (King et al., 2003; King and Wand, 2007), gender pronouns matched each respondent's description of their gender. The individuals were described as being 10%, 25%, 75%, and 90% sure that their estimate was close to the correct answer.

$$c_i = \begin{cases} 1 & \text{if} \quad y_i < V_{i1}, \\ 2 & \text{if} \quad y_i = V_{i1}, \\ 3 & \text{if} \quad V_{i1} < y_i < V_{i2}, \\ 4 & \text{if} \quad y_i = V_{i2}, \\ 5 & \text{if} \quad V_{i2} < y_i < V_{i3}, \\ 6 & \text{if} \quad y_i = V_{i3}, \\ 7 & \text{if} \quad V_{i3} < y_i < V_{i4}, \\ 8 & \text{if} \quad y_i = V_{i4}, \\ 9 & \text{if} \quad y_i > V_{i4}, \end{cases} \tag{1}$$

To create the vignette-adjusted certainty measures, I employed the widely used c-scale, which defines each respondent's own level of certainty in terms of how they rated each vignette (King et al., 2003; King and Wand, 2007). Equation 1 illustrates the translation of survey responses (y_i) recorded on the original ordinal scale to vignette-adjusted c-values (c_i)

in the beginning of the survey has the additional benefit of training respondents to use the response scales (Hopkins and King, 2010).

for each respondent i, where each of J vignettes is represented as V_{ij}. For instance, when a respondent's assessment of their own certainty (y_i) is lower than the response they assigned to the first vignette (V_{i1}), c_i is equal to 1. The range of the unique values of c_i is determined by the of the number of vignettes ($2J+1$). In this case there are 4 vignettes, so c_i can range from 1 to 9. For the sake of interpretability, the c_i is collapsed into a shorter 5 category scale, given here in terms of the level of certainty inherent in each of the J vignettes: $\leq 10\%$, $> 10\%$ & $\leq 25\%$, $> 25\%$ & $< 75\%$, $\geq 75\%$ & $< 90\%$, and $\geq 90\%$. In the case of ties and mis-ordered vignettes, we use a censored ordered probit model (King et al., 2003; King and Wand, 2007) to calculate the probability of each c_i using respondent's self-assessments, vignette responses, and demographic characteristics. While the use of anchoring vignettes, or any alternative measure of certainty, does not eliminate the need to specify a threshold with which to define beliefs as certain, it does provide context with which to interpret that threshold and control for differential item functioning. Just as past studies specify a response option at which to define beliefs as certain (e.g., 'very sure'), we must specify a latent level represented in a vignette to define beliefs as certain. Here I classify beliefs as certain when a respondent describes their belief using a response category greater than or equal to that which they used to describe the vignette expressing a 75% level of certainty.[10]

Evaluating Hypothesized Consequences of Certainty

To evaluate whether certain beliefs are more likely to influence attitudes than uncertain ones, I measured attitudes related to six of the beliefs. Respondents were asked whether immigration should be increased or decreased in the U.S., whether immigrants pose a threat to U.S. culture, and whether federal spending for public schools, foreign aid, and food stamps should be increased or decreased. I follow Ahler and Sood (2018) in measuring affect toward the opposing political party with a combination of feeling thermometers and social distance

[10] In a robustness check described below, I use a far lower threshold (25%) to define beliefs as certain.

items (e.g., how you would feel if a close family member married a Democrat/Republican).[11]

To evaluate whether certain beliefs are more likely to resist correction than uncertain ones, respondents were randomly assigned to receive or not receive information about one of two facts: the share of the U.S. population that is foreign-born and the share of public school funding that comes from the federal government. The first is among the most widely documented misperceptions in the literature (Sides and Citrin, 2007; Citrin and Sides, 2008; Hopkins et al., 2019) and has proved difficult to correct (Hopkins et al., 2019; Grigorieff et al., 2020), while the second is a less polarized and less difficult to correct (Henderson et al., 2019). There were 1,045 respondents in the immigration condition (523 treatment, 522 control) and 983 respondents in the education condition (469 treatment, 514 control). Respondents in the treatment group were told "The U.S. Census Bureau recently reported that about 14% of people living in the United States are immigrants born outside of the country" and "The U.S. Department of Education recently reported the share of funding for public schools that comes from the federal government. Approximately 8% of a local school's budget typically comes from the federal government. The rest comes from state and local governments" in the immigration and education conditions, respectively. Respondents in the control group received nearly identically phrased information, but without the correct quantities.[12] Across all conditions, the correct information was provided in the form of a question asking whether respondents had heard about the report in an effort to avoid backfire or demand effects.

Following the experiment, respondents were again asked to estimate the share of the population that is foreign born and the percentage of public school funding that comes from the federal government. To reduce demand effects, respondents first answered a battery of demographic questions and political interest items. To further reduce demand effects, these

[11]Independents who did not report leaning toward either political party were randomly assigned to answer these questions about either Democrats or Republicans, but are excluded from the main analysis.

[12]Specifically, the control groups were told "The U.S. Census Bureau recently reported the number of people living in the United States who are immigrants born outside of the country" and "The U.S. Department of Education recently reported the share of funding for public schools that typically comes from the federal government (the rest comes from the state and local governments)."

estimates were asked after three estimates that were not asked prior to the treatment: the share of the population that is Black, Christian, and aged 65 years or older. The attitude questions described above were also asked after the experiment, which allows me to measure not only the effect of the treatment on beliefs, but also on related attitudes.[13]

Results

I begin by reporting the accuracy with which respondents reported each of the 9 beliefs. Table 1 reports the mean and median response for each belief, along with the correct response for each. On average, respondents made large and systematically skewed errors for each of the estimated quantities. For instance, respondents overestimated the share of the federal budget spent on foreign aid by an average of 20.6 percentage points, the share of Americans receiving food stamps by 13.9 percentage points, and the size of the Black population by 16.3 percentage points. Estimates of the two quantities that half of respondents received information about later in the survey were similarly skewed: respondents overestimated the size of the immigrant population by 16.3 percentage points and the share of public school funding that comes from the federal government by 23.8 percentage points. These deviations are in line with prior research documenting political misperceptions among the American public (Gilens, 2001; Sides and Citrin, 2007; Lawrence and Sides, 2014; Hopkins et al., 2019; Scotto et al., 2017; Ahler and Sood, 2018) and follow the familiar pattern of systematically overestimating smaller quantities and underestimating larger ones documented by Guay et al. (2020) and Landy et al. (2018).

Across all beliefs, respondents reported relatively low levels of certainty. Figure 2.1 illustrates the distribution of certainty with which each belief is held using both the original 6-point certainty scale ranging from 'not at all sure' to 'extremely sure' (top panel) and

[13]The order of the estimation questions was chosen to minimize question ordering effects and maximize the space between prior and posterior estimates for the two quantities for which correct information was later provided in the experiment. Specifically, respondents were first asked about the quantity they were assigned to receive (or not receive) correct information about so that estimates and certainty questions that followed did not bias these estimates. The remaining 7 estimates were asked in random order.

Table 1: Political Beliefs

Quantity	Actual	Mean	SD	Median
% Budget Spent on Foreign Aid	1.0	21.6	17.7	20.0
% Republicans Making > $250k	2.0	35.5	27.0	30.0
% School Funding from Fed Govt.	8.0	31.8	21.3	25.0
% Democrats who are Atheist	9.0	26.4	23.4	20.0
% African American	12.0	28.9	14.3	29.0
% Receiving Food Stamps	13.0	26.9	19.0	20.0
% Foreign Born	14.0	30.3	20.3	25.0
% White	65.0	53.2	15.8	55.0
% Christian	70.0	51.7	19.5	53.0

Mean, median, and actual size for each of the 9 quantities estimated by respondents. Post-treatment estimates of corrected quantities are excluded.

the vignette-adjusted c-scale (bottom panel). On both scales, it is evident that very few respondents hold beliefs with high levels of certainty. For instance, on average only 4% reported being 'extremely sure' and only 9% reported being 'very sure.' Over half of responses were split evenly between 'slightly' and 'moderately' sure, 23% were associated with being 'not very' sure and 9% were associated with being 'not at all sure.' The vignette-adjusted certainty measures lend some context to these response options. Just one out of ten beliefs, on average, is held with certainty levels corresponding to 90% or more. In total, only one out of four beliefs are held with 75% certainty or more, and the majority of beliefs were held with certainty levels corresponding to 25% certain or less.

Having separately considered the accuracy of beliefs and the certainty with they are held, I now examine their joint distribution. Figure 2.2 illustrates the relationship between the amount of error and certainty in each belief, where error is operationalized as the absolute distance between each estimate and the actual size of the quantity being estimated ($|estimate - actualsize|$). For six of the nine beliefs, error and certainty are positively correlated, sometimes with large differences in absolute error across certainty. However, these differences are driven largely by the relatively small subset of respondents who hold very inaccurate beliefs. For instance, beliefs about partisan stereotypes and the size of the immigrant population that were wrong by less than 50 percentage points were held with

Figure 1: Distribution of Belief Certainty

(a) Unadjusted Certainty Responses

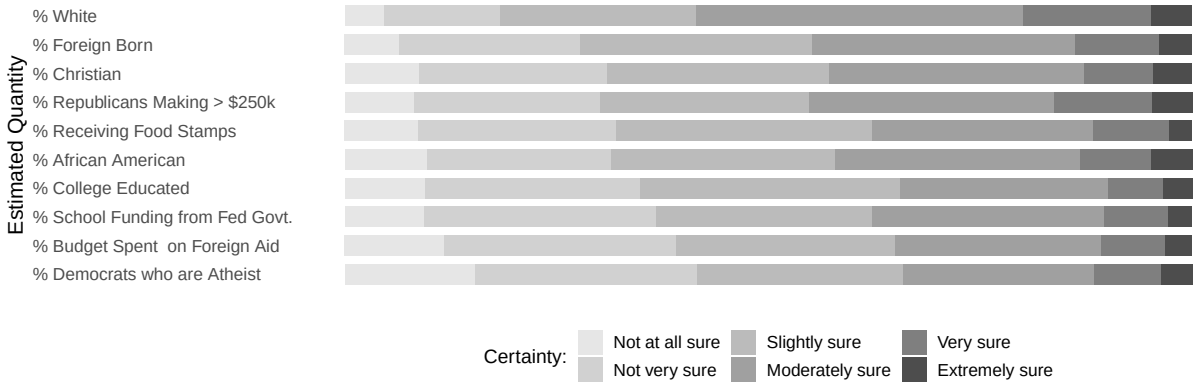

(b) Anchor Adjusted Certainty Responses (C-Scale)

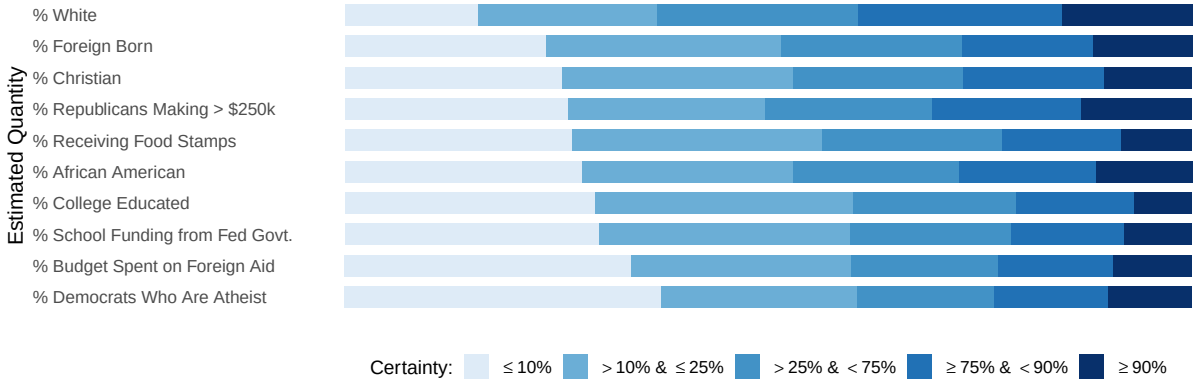

Distribution of self-reported certainty in each belief on the original response scale (top panel) and c-scale (bottom panel). The width of each bar represents the proportion of beliefs held with each level of certainty.

consistently low levels of certainty. In each case, this group represents the vast majority of respondents—75%, 88%, and 92% of respondents for beliefs about the share of Republicans who are wealthy, Democrats who are atheist, and people living in the U.S. who are foreign born, respectively. For the remaining respondents, whose beliefs are off by 50 percentage points or more, we observe far higher levels of certainty.

Certainty and error are also positively correlated for beliefs about the number of Americans who receive food stamps, the amount of public school funding that comes from the

Figure 2: Relationship Between Belief Accuracy and Certainty

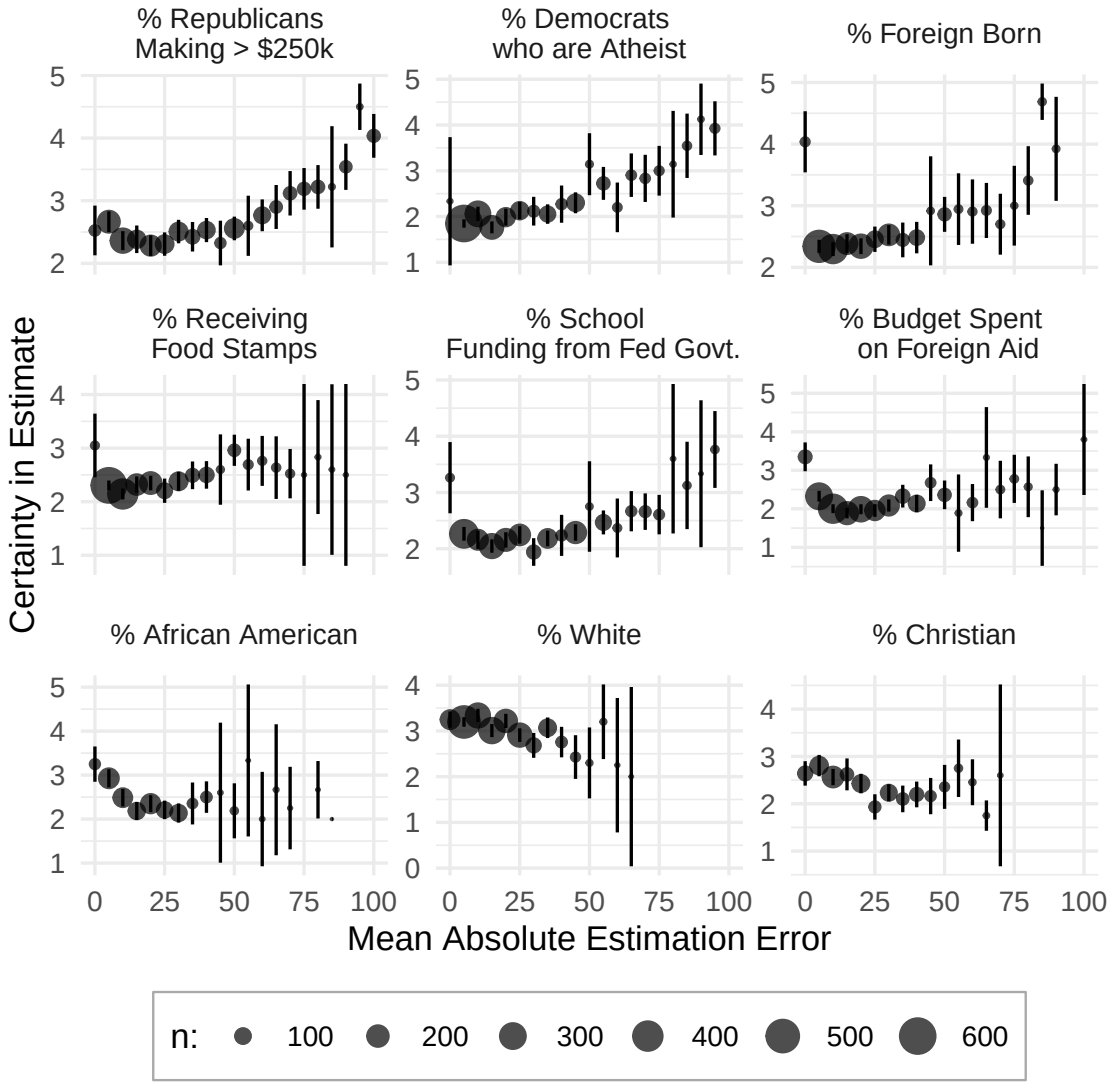

Mean certainty (c-values) across binned means of absolute estimation error, with associated 95% confidence intervals. The size of each point represents the number of respondents in the binned group.

federal government, and the share of the budget spent on foreign aid, though to a lesser extent. For the remaining three beliefs, all of which concern the demographic composition of the U.S. population, certainty is negatively correlated with error. However, these associations are relatively small: in each case the beliefs held with the most certainty have less 10 percentage points less error than those held with the least certainty, on average.

Interestingly, the beliefs for which there is a positive relationship between error and certainty are those with the clearest implications for policy preferences and attitudes. This is especially the case for beliefs about partisan stereotypes and immigration, which are the most polarizing of the nine issues. As discussed below, this suggests that in some cases respondents may express certainty as means of signaling partisanship (i.e., expressive responding).

Prevalence and Correlates of the Misinformed

Next, I examine the prevalence of the informed, uninformed, and misinformed, as well as the demographic and political correlates of each. To do so, each belief is categorized as either accurate or inaccurate and certain or uncertain. Past work measuring misperceptions with quantity estimates typically define accurate beliefs as those within 10 percentage points of the value being estimated, though some use a more conservative 5 percentage point threshold (e.g., Gilens, 2001). I therefore define beliefs as accurate when the absolute error is less than or equal to 10 percentage points, and replicate the analyses below using a 5 percentage point threshold as a robustness check in the Appendix (pg. 7). As discussed above, beliefs are classified as certain when they are held with a certainty level of 75% or greater.

Figure 2.3 reports the proportion of respondents in each of four groups: 1) the misinformed (inaccurate and certain), 2) uninformed (inaccurate and uncertain), and the informed who hold accurate beliefs with certainty (3) and without certainty (4). Across all beliefs, an average of just 12% of respondents are misinformed, 47% are uninformed, and 41% are informed. This means that people with inaccurate beliefs are 3.9 times more likely to be uncertain than certain. When a stricter threshold is used to define beliefs as certain (absolute error $\leq 5\%$) individuals we observe a nearly identical pattern—respondents are 4 times more likely to uninformed than misinformed on average. Similarly, respondents who hold accurate beliefs are 4.6 times more likely to be uncertain than certain.[14]

[14]One concern might be that 75% threshold is too high, and that this leads to misinformed being more prevalent than uninformed. As a robustness check, a very low threshold was used to define certain beliefs (beliefs $>$ 25% certainty): 23% are misinformed, 36% are uninformed, .16% are informed & certain, and 24%

Figure 3: Prevalence of the Informed, Uninformed, and Misinformed

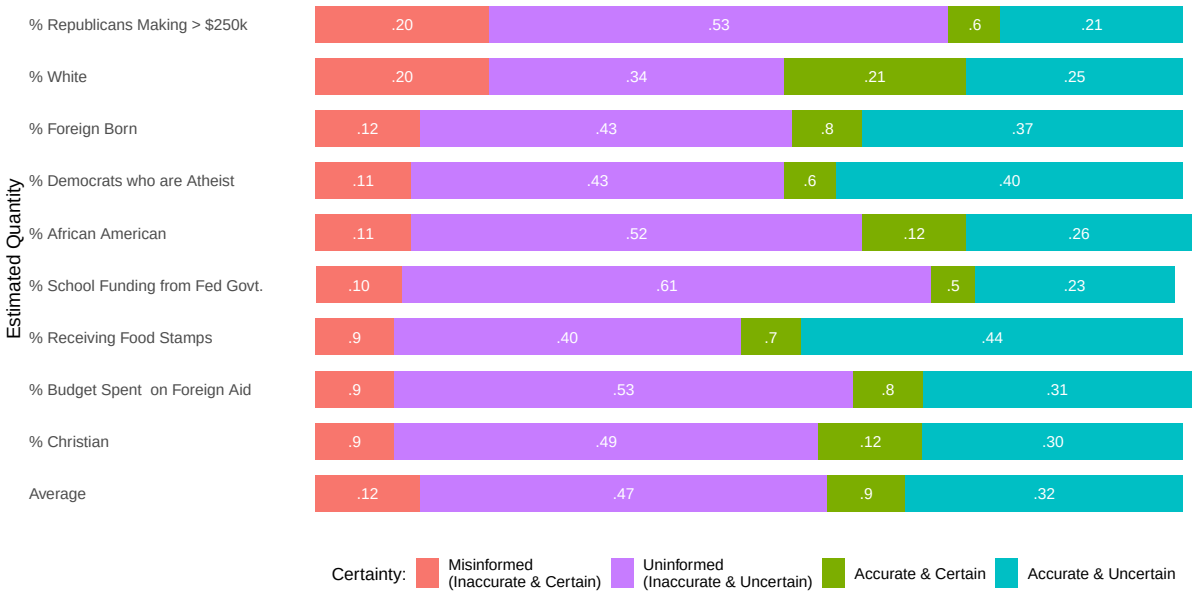

Proportion of respondents categorized as misinformed (inaccurate & certain), uninformed (inaccurate & uncertain), informed & certain, and informed & uncertain for each belief.

To understand the demographic and political characteristics associated with being misinformed, uninformed, and informed, I modeled the likelihood of belonging to each of these four groups (misinformed, uninformed, informed and certain, informed and uncertain) using multinomial logistic regression. Since each respondent reported nine beliefs, standard errors were clustered at the respondent level. I included series of standard demographic and political covariates, including education, political knowledge, political interest, sex, age, and party identification. Additionally, I include covariates describing the frequency with which respondents engage in political discussions and the degree to which their discussion partners tend to have similar political views to their own.

Figure 2.4 reports the difference in the probability of belonging to each group for a one standard deviation increase in each respondent characteristic (with the exception of gender, which increases from male (0) to female (1)). Respondents who are male, less educated, and

are informed & uncertain. So even with a dramatically lower certainty threshold, people with inaccurate beliefs are 1.6 times more likely to be uninformed than misinformed, and people with accurate beliefs are 1.7 times more likely to be uncertain than certain. See Appendix pg. 8.

Figure 4: Demographic and Political Correlates of the Misinformed

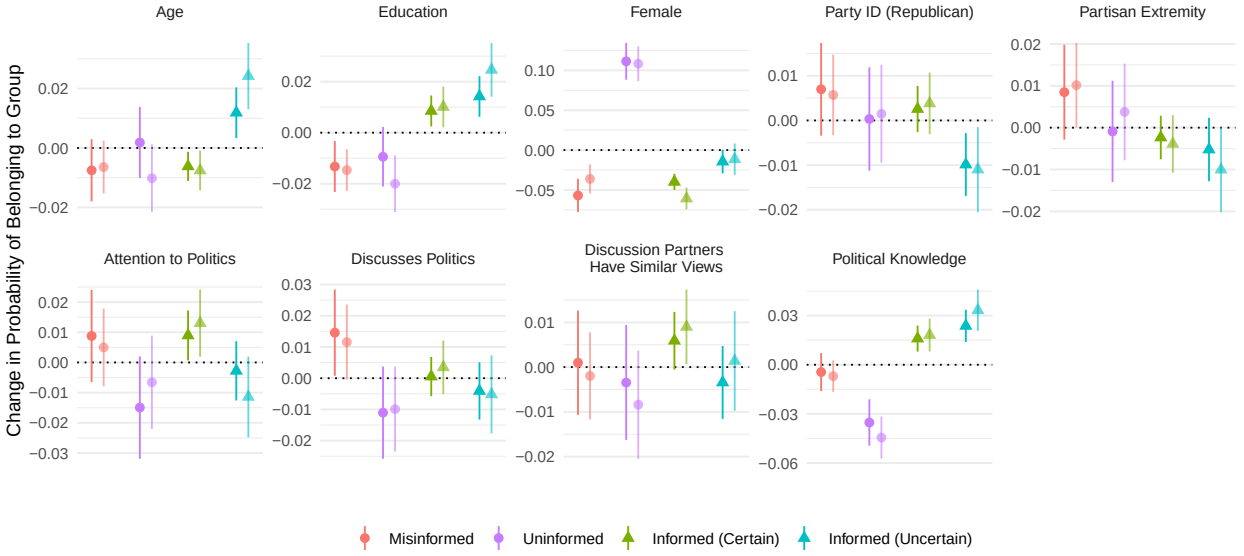

Change in the predicted probability of being misinformed, uninformed, informed & certain, or informed & uncertain for a one standard deviation increase in each covariate (e.g., age). Membership into each category was modeled using multinomial logistic regression, with standard errors clustered at the respondent level to account for the hierarchical structure of the data (multiple beliefs measured for each respondent). Vertical lines represent 95% confidence intervals.

discuss politics frequently are more likely to be misinformed. There are small and marginally significant ($p < .10$) positive associations between being misinformed and being younger, Republican, having more extreme party affiliation, and paying more attention to politics. Female and less politically knowledgeable respondents are more likely to be uninformed.

The Consequences of Holding Beliefs with Certainty

One remaining question concerns the downstream consequences of certainty in accurate beliefs. To understand whether certain beliefs are more closely related to attitudes than uncertain ones, I assess the moderating role of certainty on the relationship between beliefs and attitudes about immigration, education, foreign aid, welfare, and feelings toward members of the opposing political party. I regress each attitude on the interaction between

certainty and beliefs, as well as controls for gender, age, education, race, ideology, and political knowledge.[15] Attitudes are coded such that positive values represent the attitude we expect with higher estimates (opposition to immigration, negative attitudes toward members of the out-party, and preferences for restricted spending on education, foreign aid, and welfare).

Figure 5: Moderating Role of Certainty on Relationship Between Beliefs and Attitudes

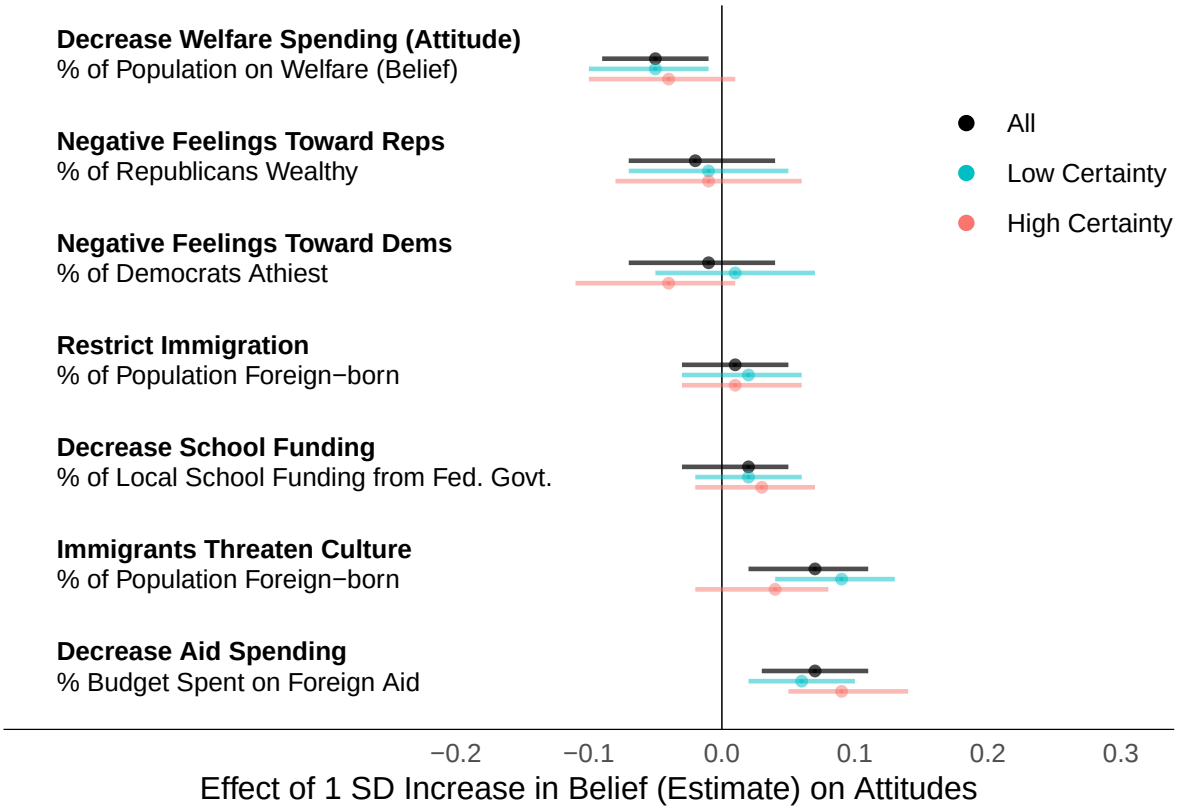

Predicted levels of attitudes (in standard deviations) for a one standard deviation increase in overestimation error. Predicted attitudes are plotted separately for low and high levels of certainty (5th and 95th percentiles of certainty associated with each estimate). Horizontal lines represent 95% confidence intervals.

Figure 2.5 reports the main effect of each belief on each attitude, as well as main effects of beliefs on attitudes for low and high levels of certainty. Overall, we observe no evidence supporting the hypothesis that certain beliefs are more closely related to attitudes than un-

[15]For attitudes related to immigration and education, I use posterior estimates (those elicited after the experiment) and exclude respondents who were randomly assigned to receive correct information about each quantity beforehand.

certain ones. While we observe a statistically significant main effect of beliefs on attitudes about welfare, foreign aid, and the threat immigrants pose to American culture, these relationships are not any stronger when beliefs are held with certainty. In fact, in one case we observe evidence in the opposite direction than hypothesized by past work: Republicans' beliefs about the share of Democrats who are atheist are *less* strongly associated with negative attitudes toward Democrats when they are held with certainty. We observe a similar pattern, though only marginally significant, between beliefs about the size of the immigrant population and perceptions of immigrants as threatening American culture, which is weaker for beliefs held with greater certainty. Thus, we observe no evidence to support the widespread expectation that certain beliefs are more strongly related to attitudes than uncertain ones.

Finally, I test the hypothesis that certain beliefs are more resistant to correct information than uncertain ones. Figure 2.6 reports the main effect of receiving correct information about the size of the immigrant population and the share of public school funding that comes from the federal government. Panel A reports the effect of the treatment on belief-updating (prior belief - posterior belief). In both cases, the treatment improved the accuracy of respondents' beliefs, decreasing estimates of the size of the immigrant population by an average of 7 percentage points and the share of school funding that comes from the federal government by 14 percentage points. Contrary to expectations, however, we observe no moderating effect of certainty on this relationship—prior beliefs held with higher levels of certainty were no more resistant to change than those held with lower levels of certainty.

Despite substantially updating their beliefs in response to correct information, attitudes remained stable for both immigration and education. In line with prior work (e.g., Hopkins et al. 2018), we observe no significant main effect of the treatment on attitudes for either issue. For both immigration attitudes we also observe no heterogeneous treatment effects across certainty. We do observe a small and marginally significant ($p < .10$) heterogeneous treatment effect for attitudes about education, but in the oppose direction than expected. The effect of learning that 8% of local school funding comes from the federal government

Figure 6: Effect of Correct Information on Belief-Updating and Attitudes

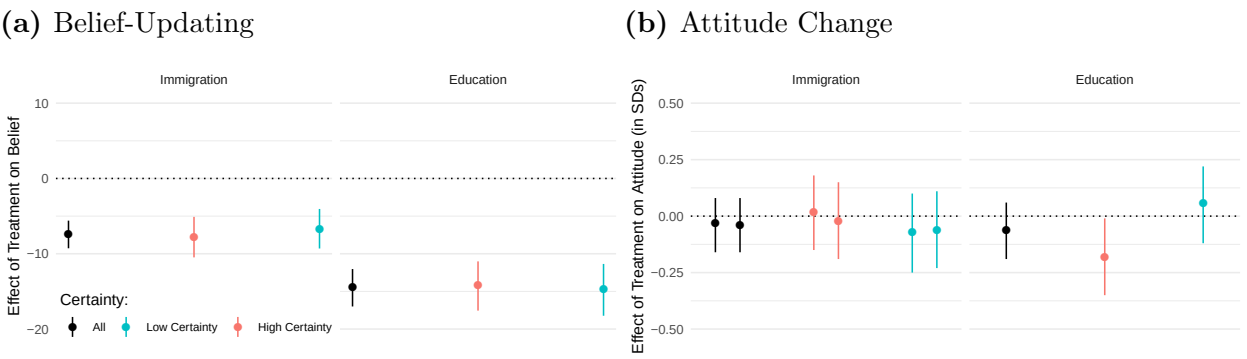

Effect of the information treatment on belief updating ($estimate_1 - estimate_2$, Panel A) and attitudes (Panel B). Each panel reports the main effect of the treatment as well as heterogeneous treatment effects estimated by modeling each outcome with an interaction between the treatment and certainty in the prior belief ($estimate_1$). The two immigration attitudes are presented separately in Panel B: preferences for restricted immigration (left) and perceptions of immigrants as threatening American culture (right).

on support for federal education spending was *stronger* for certain respondents. In other words, while past work has hypothesized that certain beliefs are more resistant to correct information (Kuklinski et al., 2000; Hochschild and Einstein, 2015; Jerit and Zhao, 2020). I find no support for this claim, and in one case find that correcting certain beliefs is more likely to results in attitude change.

Discussion

Recent work paints a grim picture of the public's capacity to meaningfully engage in politics. Concerns over the prevalence and consequences of being misinformed have sparked a growing body of work on misperceptions, but this literature has largely ignored the degree to which inaccurate beliefs are held with certainty. The handful of studies that do measure certainty examine a narrow range of misperceptions, use methods of measuring certainty that are difficult to interpret, and come to conflicting conclusions about the relationship between accuracy and certainty. The aim of this paper is to provide a clearer understanding about the prevalence, correlates, and consequences of being misinformed. The findings presented

here suggest a more optimistic and normatively positive view of the public's knowledge of politics. I find that while inaccurate beliefs related to salient political issues, partisan stereotypes, and the size of demographic groups are widespread, they are typically held with very low levels of certainty. Indeed, on average only 12% of respondents in this study were classified as misinformed.

These findings have implications not only for how we interpret inaccurate beliefs reported on surveys, but also for the growing number of studies that attempt to change attitudes by correcting inaccurate beliefs. This body of research assumes a causal model in which beliefs influence attitudes. For instance, studies that find overestimating the size of the immigrant population is associated with holding anti-immigration attitudes theorize that the former influences the latter (Citrin and Sides, 2008; Sides and Citrin, 2007; Herda, 2010), which leads to the expectation that correcting inaccurate beliefs will result in attitude change (e.g., Hopkins et al., 2019; Nyhan and Reifler, 2010). However, both of these expectations rely on the underlying assumption that beliefs are held with certainty. Indeed, it is difficult to imagine the mechanism by which highly uncertain beliefs would exert influence over attitudes. Anti-immigration attitudes cannot be rooted in beliefs that immigrants comprise 40% of the population when a person is is entirely unsure of the size of the immigrant population. In fact, it is more likely that this person may tend to provide an estimate that is too high *because* she dislikes immigrants and prefers less immigration. Moreover, when the causal arrow points in this direction, expectations that correcting beliefs will change attitudes should be low. At best, correct information provides an additional belief to compete with those in which an individual's attitude was originally grounded. Indeed, it is possible that unobserved variation in belief certainty helps to explain why informational treatments often succeed in correcting beliefs, but fail in changing attitudes (Hopkins et al., 2019; Nyhan and Reifler, 2010; Weeks, 2018; Chan et al., 2017).

The findings presented here also have implications for how scholars conceptualize the role of belief certainty. Early concerns over the misinformed were driven by the hypothesized

consequences of holding inaccurate beliefs with certainty—specifically that certain beliefs are more likely to influence attitudes and resist corrections than uncertain ones. In this study I provide the first empirical test of these hypotheses and find no empirical support for either. Certain respondents are no more likely to have attitudes associated with their beliefs, nor are they less likely to incorporate correct information into their beliefs and attitudes.

One potential explanation for this is that some respondents express certainty as a means of expressing partisan identity (i.e. partisan cheer-leading). Recent work that has considered the role that expressive responding plays in shaping misperceptions (Berinsky, 2017), but no work to date has evaluated its role in shaping expressions of belief certainty. The findings presented here offer some suggestive evidence. Indeed, the three most polarizing beliefs measured in this study were those for which respondents expressed the most certainty in highly inaccurate responses. One potential explanation for this behavior is that some respondents were expressively responding both by reporting beliefs about the opposing political party and immigration *and* by expressing certainty in those beliefs. Future work should leverage incentive compatible designs to understand the degree to which respondents expressively respond to questions measuring subjective belief certainty.

References

Achen, C. H. and L. M. Bartels (2017). *Democracy for realists: Why elections do not produce responsive government*, Volume 4. Princeton University Press.

Ahler, D. J. and G. Sood (2018, July). The Parties in Our Heads: Misperceptions about Party Composition and Their Consequences. *The Journal of Politics 80*(3), 964–981.

Alba, R., R. G. Rumbaut, and K. Marotz (2005). A distorted nation: Perceptions of racial/ethnic group sizes and attitudes toward immigrants and other minorities. *Social Forces 84*(2), 901–919.

Aldrich, J. H. and R. D. McKelvey (1977). A Method of Scaling with Applications to the 1968 and 1972 Presidential Elections. *The American Political Science Review 71*(1), 111–130.

Bartels, L. M. (1996). Uninformed votes: Information effects in presidential elections. *American Journal of Political Science*, 194–230.

Berinsky, A. J. (2007). Assuming the costs of war: Events, elites, and american public support for military conflict. *The Journal of Politics 69*(4), 975–997.

Berinsky, A. J. (2017, April). Rumors and Health Care Reform: Experiments in Political Misinformation. *British Journal of Political Science 47*(2), 241–262.

Bode, L. and E. K. Vraga (2015). In related news, that was wrong: The correction of misinformation through related stories functionality in social media. *Journal of Communication 65*(4), 619–638.

Campbell, A., P. E. Converse, W. E. Miller, and D. E. Stokes (1960). The American Voter. *Ann Arbor, MI: University of Michigan Press*.

Carpini, M. X. D. and S. Keeter (1996). *What Americans Know about Politics and Why it Matters*. Yale University Press.

Chan, M.-p. S., C. R. Jones, K. Hall Jamieson, and D. Albarracín (2017). Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological science 28*(11), 1531–1546.

Citrin, J. and J. Sides (2008). Immigration and the imagined community in europe and the united states. *Political Studies 56*(1), 33–56.

Clifford, S. and J. Jerit (2016). Cheating on political knowledge questions in online surveys: An assessment of the problem and solutions. *Public Opinion Quarterly 80*(4), 858–887.

Conover, P. J., S. Feldman, and K. Knight (1986). Judging inflation and unemployment: The origins of retrospective evaluations. *The Journal of Politics 48*(3), 565–588.

Converse, P. (1964). The Nature of Belief Systems in Mass Publics. In *Ideology and Discontent*, pp. 206–61. New York: Free Press.

Coppock, A. and O. A. McClellan (2019). Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents. *Research & Politics 6*(1), 2053168018822174.

Duffy, B. (2018). *The Perils of Perception: Why We're Wrong About Nearly Everything*. Atlantic Books.

Enke, B. and T. Graeber (2019). Cognitive uncertainty. Technical report, National Bureau of Economic Research.

Flynn, D., B. Nyhan, and J. Reifler (2017). The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics. *Political Psychology 38*, 127–150.

Gilens, M. (2001). Political ignorance and collective policy preferences. *American Political Science Review 95*(2), 379–396.

Gonzalez, R. and G. Wu (1999). On the shape of the probability weighting function. *Cognitive psychology 38*(1), 129–166.

Graham, M. H. (2020, March). Self-Awareness of Political Knowledge. *Political Behavior 42*(1), 305–326.

Grigorieff, A., C. Roth, and D. Ubfal (2020). Does Information Change Attitudes Toward Immigrants? *Demography 57*(3), 1117–1143.

Guay, B. and C. Johnston (2021). Ideological asymmetries and the determinants of politically motivated reasoning. *American Journal of Political Science*, 1–60.

Guay, B., Tyler Marghetis, Cara Wong, and David Landy (2020). Bayesian Reasoning and Demographic Misperceptions.

Haran, U., D. A. Moore, and C. K. Morewedge (2010, December). A simple remedy for overprecision in judgment. *Judgment and Decision Making; Tallahassee 5*(7), 467.

Henderson, M., David Houston, Paul Peterson, and Martin West (2019). Public Support Grows for Higher Teacher Pay and Expanded School Choice. Technical report.

Herda, D. (2010). How Many Immigrants? Foreign-Born Population Innumeracy in Europe. *Public Opinion Quarterly 74*(4), 674–695.

Hochschild, J. L. and K. L. Einstein (2015). *Do facts matter?: Information and misinformation in American politics*, Volume 13. University of Oklahoma Press.

Hopkins, D. J. and G. King (2010, June). Improving Anchoring Vignettes: Designing Surveys to Correct Interpersonal Incomparability. *Public Opinion Quarterly 74*(2), 201–222.

Hopkins, D. J., J. Sides, and J. Citrin (2019). The Muted Consequences of Correct Information About Immigration. *The Journal of Politics 81*(1), 315–320.

Jerit, J. and Y. Zhao (2020). Political misinformation. *Annual Review of Political Science 23*, 77–94.

Key, V. O. (1961). Public Opinion and American Democracy.

King, G., C. J. L. Murray, J. A. Salomon, and A. Tandon (2003). Enhancing the Validity and Cross-Cultural Comparability of Measurement in Survey Research. *The American Political Science Review 97*(4), 567–583.

King, G. and J. Wand (2007). Comparing Incomparable Survey Responses: Evaluating and Selecting Anchoring Vignettes. *Political Analysis 15*(1), 46–66.

Kuklinski, J. H., P. J. Quirk, J. Jerit, D. Schwieder, and R. F. Rich (2000). Misinformation and the currency of democratic citizenship. *Journal of Politics 62*(3), 790–816.

Landy, D., B. Guay, and T. Marghetis (2018). Bias and ignorance in demographic perception. *Psychonomic bulletin & review 25*(5), 1606–1618.

Lau, R. R. and D. P. Redlawsk (1997). Voting correctly. *American Political Science Review*, 585–598.

Lawrence, E. D. and J. Sides (2014). The consequences of political innumeracy. *Research & Politics 1*(2), 2053168014545414.

Lee, S. and A. Matsuo (2018, February). Decomposing political knowledge: What is confidence in knowledge and why it matters. *Electoral Studies 51*, 1–13.

Mettler, S. (2011). *The Submerged State: How Invisible Government Policies Undermine American Democracy*. University of Chicago Press.

Nadeau, R., R. G. Niemi, and J. Levine (1993). Innumeracy about minority populations. *Public Opinion Quarterly 57*(3), 332–347.

Nyhan, B. and J. Reifler (2010, June). When Corrections Fail: The Persistence of Political Misperceptions. *Political Behavior 32*(2), 303–330.

Pasek, J., G. Sood, and J. A. Krosnick (2015, August). Misinformed About the Affordable Care Act? Leveraging Certainty to Assess the Prevalence of Misperceptions. *Journal of Communication 65*(4), 660–673.

Scotto, T. J., J. Reifler, and J. vanHeerde Hudson (2017). We Spend How Much? Misperceptions, Innumeracy, and Support for the Foreign Aid in the United States and Great Britain. *Journal of Experimental Political Science 4*(2), 119–128. Publisher: Cambridge University Press.

Sides, J. and J. Citrin (2007). European opinion about immigration: The role of identities, interests and information. *British journal of political science 37*(3), 477–504.

Thorson, E. (2015). Identifying and correcting policy misperceptions. *Unpublished Paper, George Washington University. Available at http://www. americanpressinstitute. org/wp-content/uploads/2015/04/Project-2-Thorson-2015-Identifying-Political-Misperceptions-UPDATED-4-24. pdf*.

Wand, J. (2013). Credible Comparisons Using Interpersonally Incomparable Data: Nonparametric Scales with Anchoring Vignettes. *American Journal of Political Science 57*(1), 249–262.

Weeks, B. E. (2018). Media and political misperceptions. *Misinformation and mass audiences*, 140–156.

Wong, C. J. (2007). "little" and "big" pictures in our heads: Race, local context, and innumeracy about racial groups in the united states. *Public Opinion Quarterly 71*(3), 392–412.

Appendix

Survey Items

Anchoring Vignettes

- INSTRUCTIONS: Next you'll see descriptions of 4 people who are asked the same quiz question. Some people are very sure that their answer is correct, others are not. It's your job to report how sure you think each person is. After you've read these instructions, please click the 'next' button below:

- Jordan, Alex, Jamie, and Riley are each asked to estimate how many adults in the U.S. have a 4-year college degree. Jordan thinks there is about a 10% chance [his/her] estimate is close to the correct answer. How sure is Jordan about her answer? (Extremely sure, Very sure, Moderately sure, A little sure, Not very sure, Not at all sure)

- Alex thinks there is about a 25% chance [his/her] estimate is close to the correct answer. How sure is Alex about her answer? (Extremely sure, Very sure, Moderately sure, A little sure, Not very sure, Not at all sure)

- Jamie thinks there is about a 75% chance [his/her] estimate is close to the correct answer. How sure is Jamie about her answer? (Extremely sure, Very sure, Moderately sure, A little sure, Not very sure, Not at all sure)

- Riley thinks there is about a 90% chance [his/her] estimate is close to the correct answer. How sure is Riley about her answer? (Extremely sure, Very sure, Moderately sure, A little sure, Not very sure, Not at all sure)

Misperceptions

- Next, we're going to ask you some questions about politics and the U.S. population. We're just interested in your best guess for each question. Please do not look up the answer to any question. After each question you'll be asked how sure you are about your answer.

- What percent of people living in the United States are immigrants?

- Funding for public schools in the U.S. comes from a combination of federal, state, and local government. What percent of funding for public schools in your local school district do you think comes from the federal government?

- What percent of the U.S. population do you think received food stamps in 2019?

- What percent of the U.S. population do you think is White/Caucasian?

- What percent of Republican Party supporters do you think make more than $250,000 a year?

- What percent of Democratic Party supporters do you think are atheist or agnostic?
- What percent of the U.S. population is Black?
- What percent of the U.S. population has a 4-year college degree?
- What percent of the U.S. population do you think is Christian?

Attitudes and Policy Preferences

- Do you think the number of immigrants from foreign countries who are permitted to come to the United States to live should be.... (Increased a lot, Increased a little, Kept about the same, Decreased a little, Decrease a lot)
- How likely is it that current and future immigration will threaten the American way of life? (Extremely likely, Very likely, Somewhat likely, Somewhat unlikely, Very unlikely, Extremely unlikely)
- How important is the issue of increasing or decreasing immigration to you personally? (Extremely important, Very important, Moderately important, Slightly important, Not at all important)
- Do you think that federal funding for public schools should be... (Increased a lot, Increased a little, Kept about the same, Decreased a little, Decrease a lot)
- Do you think that federal spending on foreign aid to other countries should be... (Increased a lot, Increased a little, Kept about the same, Decreased a little, Decrease a lot)
- Do you think funding for food stamps should be increased or decreased?... (Increased a lot, Increased a little, Kept about the same, Decreased a little, Decrease a lot)

Attention Checks and Diagnostics

2,724 respondents entered the survey, and 2,477 (90.9%) completed it. There was an approximately even distribution of respondents in the treatment and control groups within each of the randomly assigned issue conditions: 1,351 respondents in the education condition (645 in control, 706 in treatment) and 1,373 in the immigration condition (692 in control, 681 in treatment).

I also included two pre-treatment attention checks. First, respondents were told "Some people fail to read surveys carefully, and simply click through. To show you do not do this, simply ignore the response options below and mark 'probably true'." This attention check captures whether respondents read the questions carefully, which is especially important in this study. 172 respondents (6.97%) failed this attention check were excluded from the analysis.

The second attention check asked respondents how many times a week (0 days, 1 day, or more than 1 day) they engaged in the following activities: eaten dinner, gone geocaching,

flown a helicopter, run a marathon, used a computer. We make the assumption that respondents will not have engaged in more than one of the three low-incidence activities (flown a helicopter, run a marathon, and gone geocaching) within the same week and flag respondents who reported doing so as shirkers. 238 respondents (9.61%) reported engaging in more than one of these activities. We also excluded 15 additional respondents who straight-lined on the same grid (this relies on the assumption that people have eaten dinner or used a computer at least once in the past week).

Respondents were also excluded if they sped through the survey. The mean completion time was 11.76 minutes (median = 11 minutes). Respondents were considered too fast if they completed the survey in less than one half the median completion time (5.5 minutes). Additionally, respondents were excluded if their completion time was greater than that of 95% of respondents (32.25 minutes). This threshold was chosen to still include some respondents who may have paused the survey and returned to it. 149 respondents (6.02%) completed the survey too quickly and 124 (5.01%) completed the survey too slowly.

In all, 2,028 (81.91%) of respondents who completed the survey passed all data quality checks (469 in education control, 514 in education treatment, 523 in immigration control, and 522 in immigration treatment).

Anchoring Vignette Diagnostics

Distribution of Vignette Responses

Respondents were provided with 4 anchoring vignettes, each describing certainty at the 10%, 25%, 75%, or 90% levels. Each vignette had six response options: extremely (6), very (5), moderately (4), a little (3), not very (2), not at all sure (1). Figure 1 reports the distribution of ordinal responses ascribed to each vignette. A strong positive relationship between the certainty inherent in each vignette and ordinal responses can be observed. The modal response for the 10%, 25%, 75% and 90% vignettes are 'Not very sure', 'A little sure', 'Very sure', and 'Extremely sure', respectively. As expected, and indeed required for the effective use of anchoring vignettes, there is some variation in the ordinal response used to describe each vignette. Interestingly, the second most common ordinal response in each case is the one directly preceding the modal response.

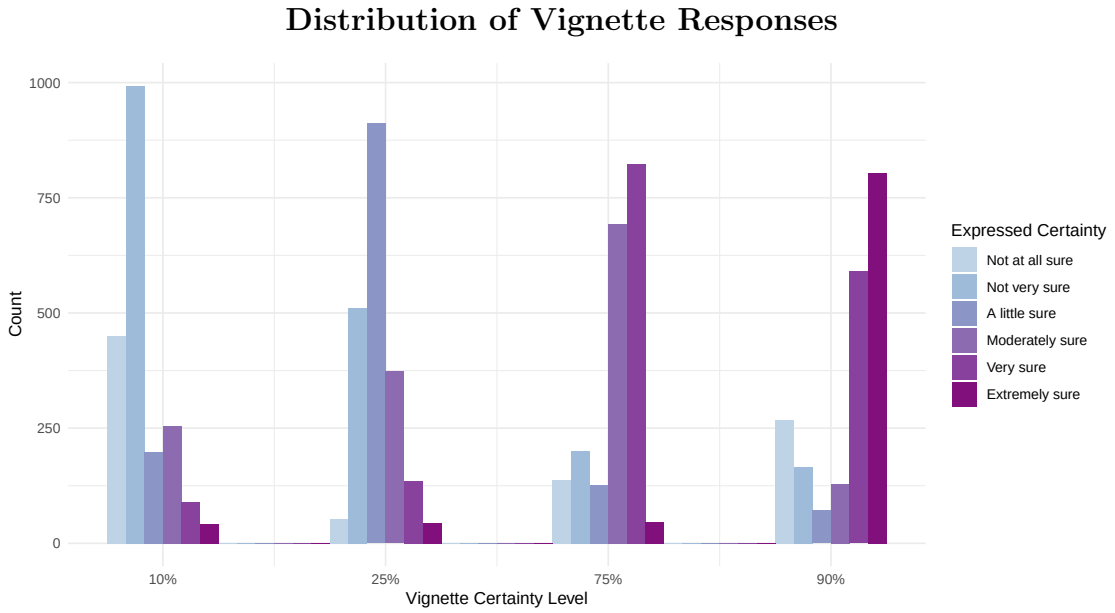

Figure 7: The distribution of ordinal responses ascribed to each vignette.

Ranking Violations

Ranking violations occur when respondents rank the vignettes out of the order intended by the researcher (i.e., when the respondents' ordinal ranking of the vignettes does not match the ordinal ranking intended by the researcher). As King et al. (2003) note, ranking violations are to be expected when using anchoring vignettes and care should be taken to design vignettes in a way that minimizes these violations. One reason for using probabilities to describe the certainty inherent to each vignette is that ranking violations should be rare. For most respondents, it should be clear that a person who is 25% certain, for instance, is more certain than someone who is only 10% certain.

Indeed, ranking violations were rare in this study. Excluding the 5% of respondents who reported the same level of certainty for each of the four vignettes (i.e., straight-lining), 90% of respondents reported that the 25% vignette was more or as certain as the 10% vignette; 80% reported that the 75% vignette was more or as certain as the 25% vignette; and 85% reported that the 90% vignette was more or as certain as the 75% vignette. Figure 2 illustrates the ten most frequent response patterns for the anchoring vignettes, where the x axis represents the level of certainty inherent in each vignette and the y axis represents which of the six ordinal response options respondents used to describe them. In total, 56% of respondents used one of these ten vignette-certainty pairings.

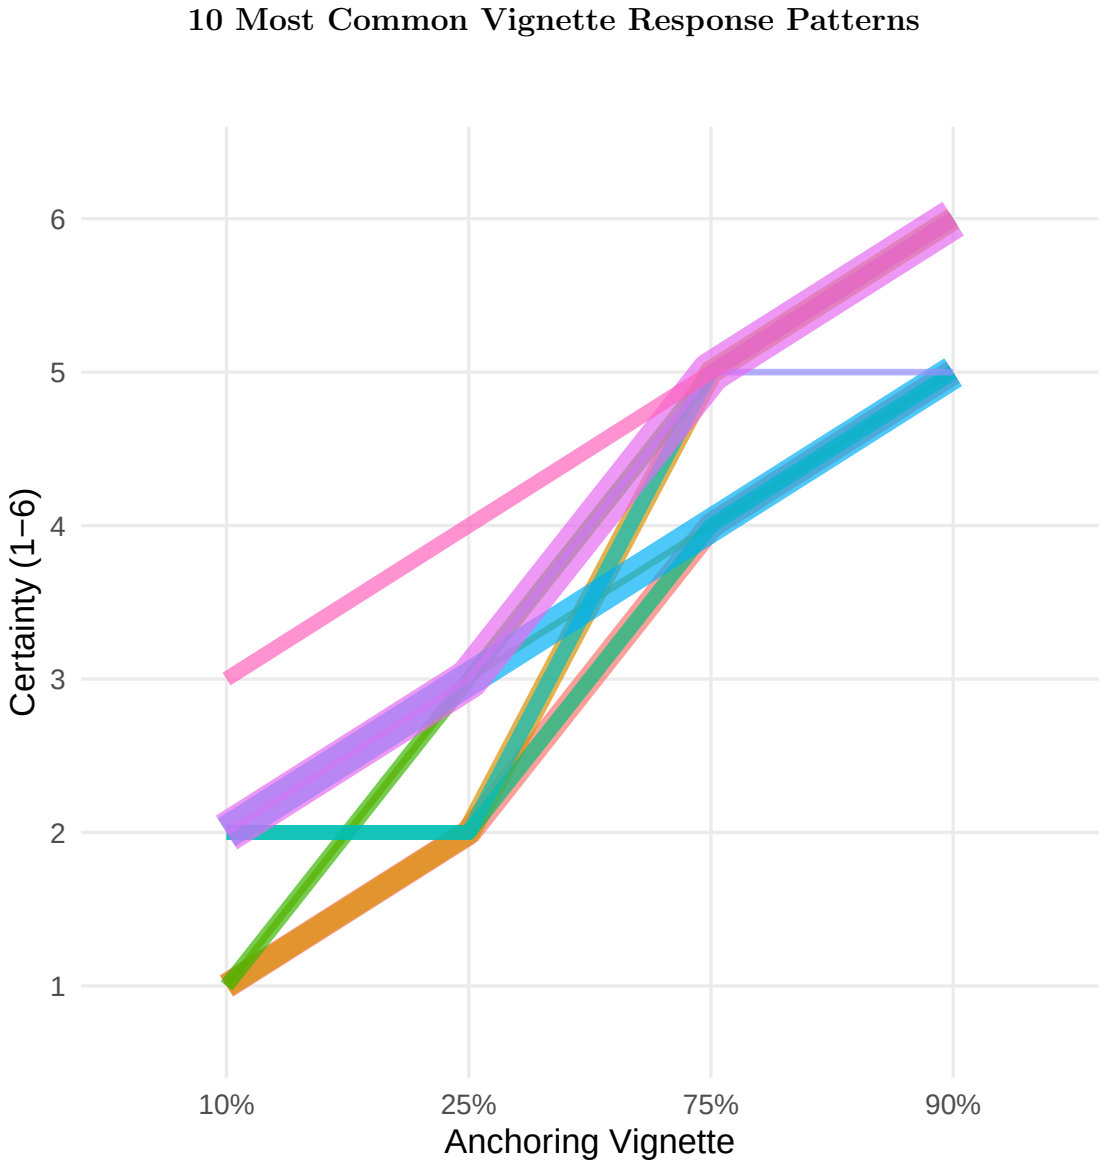

Figure 8: The ten most common mappings of anchoring vignettes to the ordinal certainty scale. The width of the line represents the relative portion of respondents who used each mapping.

Prevalence of the Misinformed vs. Uninformed Robustness Checks

Figure 3 in the main text reports the proportion of respondents who are misinformed, uninformed, informed & certain, and informed & uncertain for each of nine beliefs. Creating these four categories requires specifying two thresholds: one that distinguishes between accurate and inaccurate beliefs and one that distinguishes between certain and uncertain beliefs. As explained in the main text, I use a conservative threshold with which to define beliefs as certainty: beliefs that differ from the truth by ≤ 10 percentage points are considered accurate. Figure 3 reports the results of this analysis using a less conservative threshold: beliefs that differ from the truth by ≤ 5 percentage points are considered accurate. Naturally, we observe a far greater proportion of respondents who are uninformed. This more conservative threshold does not alter the results. While in the main text of the paper I report that respondents are 3.9 times more likely to be uninformed than misinformed, here respondents are 4.1 times more likely to be uninformed than misinformed.

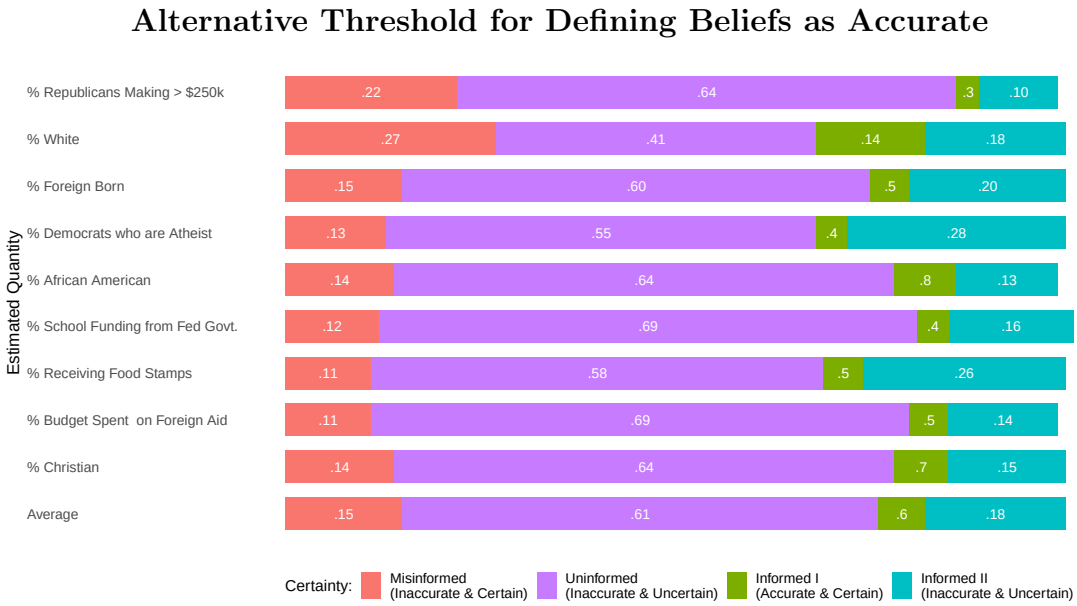

Figure 9: Proportion of respondents categorized as misinformed (inaccurate & certain), uninformed (inaccurate & uncertain), informed & certain, and informed & uncertain for each belief using a more conservative threshold with which to define beliefs as accurate.

As discussed in the main text, beliefs are classified as certain when they are held with a certainty level of 75% or greater. Figure 4 reports the results of this analysis with a far more conservative threshold: beliefs are classified as certain when they are held with a certainty level of greater than 25%. Even when this dramatically lower certainty threshold is used, people with inaccurate beliefs are 1.6 times more likely to be uninformed than misinformed, and people with accurate beliefs are 1.6 times more likely to be uncertain than certain.

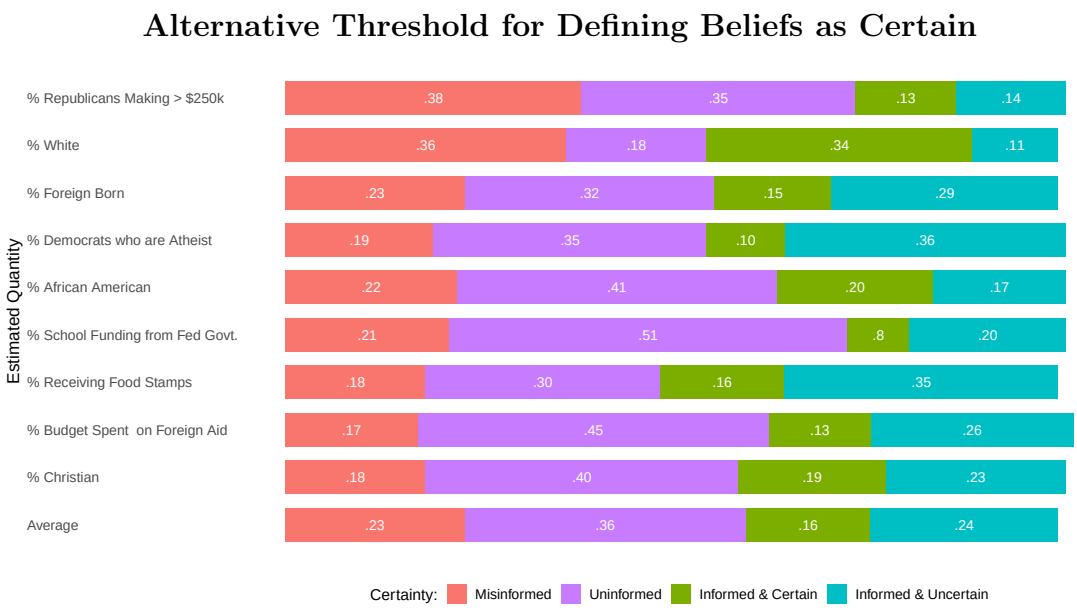

Figure 10: Proportion of respondents categorized as misinformed (inaccurate & certain), uninformed (inaccurate & uncertain), informed & certain, and informed & uncertain for each belief using a more conservative threshold with which to define beliefs as certain.