Abstract

Mismeasuring Misperceptions:
How Surveys Distort the Nature of Partisan Belief Differences

Matthew Henry Graham
2020

Survey measures suggest that partisanship has a pervasive influence on public perceptions of reality. Even when asked purely factual questions, Democrats and Republicans offer systematically different assessments of the state of the economy. When asked about rumors and politicized controversies, many members of the public readily endorse false claims that denigrate political opponents. These patterns raise serious questions about the public's fitness for democratic citizenship.

This volume shows that contemporary survey research systematically overestimates the size of partisan divides and the prevalence of political misperceptions. Members of the public fall far short of the ideal of the perfectly informed citizen. However, our second-best world is beset by a great deal more ignorance of inconvenient truths, and a great deal less outright belief in false claims, than surveys suggest at face value.

Chapter 2 sets the stage for the empirical chapters by showing that scholarship on the meaning of survey responses implicitly embraces two distinct conceptions of a "belief." The best-known argument that surveys do not generally measure beliefs, Philip Converse's non-attitudes thesis, embraces a threshold conception: a belief is a stable, pre-existing judgment. In showing that surveys do measure beliefs, Converse's critics define beliefs probabilistically: respondents tend to favor one option over the others more often than chance would predict. This raises a fundamental tension for the study of political misperceptions. In political science, the consensus definition of a misperception embraces a threshold-based conception of belief: one who misperceives is certain of something that is wrong. However, there is no evidence that surveys can identify this type of belief, and most of what evidence exists concerns attitudes rather than facts or rumors.

The next three chapters empirically examine the sense in which common measurement technologies isolate beliefs. Chapter 3 shows that the most common approach to dealing with respondent uncertainty, explicit "don't know" (DK) response options, mostly removes very uncertain responses from the sample. This leaves a substantial amount of uncertainty among respondents who answer the question.

Existing research posits that measuring respondents' level of certainty about their answer can overcome this problem. Chapter 4 examines the degree to which certainty scales capture genuine

differences in respondents' degree of belief using two indicators of whether survey-takers act as if they really believe their responses: they should be stable in their responses, and they should be willing to put their money where their mouth is. For classic knowledge questions designed to capture general political awareness, certainty scales do an excellent job of capturing variation between knowledge, defined as complete belief in the correct answer, and total ignorance, defined as assigning equal probability to all response options.

If certainty scales capture meaningful variation in knowledge and ignorance, might they also capture meaningful variation between ignorance and the "incorrect knowledge" that would result from similarly strong acceptance of false claims? In contrast to the encouraging results for political awareness questions, Chapter 5 casts doubt on respondents' claims to completely believe false statements. Incorrect answers to questions about the economy and politicized controversies consistently fall well short of the benchmarks set in the previous chapter. Even respondents who claim to be certain about incorrect answers are fairly unstable in their responses. Claims to be certain about an incorrect answer are not best-interpreted as evidence that the respondent is misinformed, but as "miseducated guesses" based on misleading heuristics.

Having dispensed with the notion that surveys measure beliefs in the threshold sense, the volume turns to the probabilistic framework's implications for measuring belief differences between Democrats and Republicans. Chapter 6 shows that prevailing survey measures systematically exaggerate partisan belief differences, by 40 percent on average. These effects are not uniform across questions: owing to partisan differences in certainty, prevailing practices sometimes *under*estimate belief differences. In addition to distorting estimates of the size of belief differences, surveys paint a false portrait of their nature. Respondents express substantially more certainty about correct answers than incorrect answers, regardless of party. Paired with the findings in Chapters 4 and 5, this shows that in general, partisan belief differences should be viewed as evidence of differential knowledge of convenient and inconvenient truths rather than as evidence of misperceptions.

None of this is to say that misperceptions and partisan divides do not raise serious concerns, where they in fact exist. Instead, the findings demonstrate that current practices in survey research provide little basis to distinguish between which false beliefs and partisan divides are real and which are artifacts of survey technology. By introducing a new conceptual and empirical basis for examining the sense in which respondents believe their answers, this volume furnishes a more accurate interpretation of existing data and a toolkit for measuring potential improvements.

Mismeasuring Misperceptions:
How Surveys Distort the Nature of Partisan Belief Differences

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Matthew Henry Graham

Dissertation Director:
Gregory Alain Huber

December 2020

Contents

To my grandparents,

who built the foundation for my opportunities,

and by their words and stories,

taught me to be grateful.

Acknowledgements

In the big picture, I'm grateful to have been able to spend so much of my life developing my own skills and ideas. Yale paid me a living wage for five and a half years of skill development and research in exchange for two years of part-time teaching. The Center for the Study of American Politics (CSAP) and the Institution for Social Policy Studies (ISPS) made the data collection possible, both directly and via generous faculty members' survey time or research accounts. I'm grateful for these opportunities and to everyone who helped me make whatever I make of it.

I'm grateful to the faculty members who saw me through the program. Greg Huber, my chair, was generous with his time and resources. He engaged seriously with my ideas, offered thoughtful advice and perspectives around every turn, and consciously worked to put me in positions to succeed. Alex Coppock was similarly essential. I was fortunate to have his support throughout, and especially over the past year when I needed help getting these ideas over the finish line. Alan Gerber displayed an impressive commitment to engaging with my work despite being a dean. The work benefitted a great deal from his insights.

Many other Yale faculty showed me support, Milan Svolik provided a model of how to ask research questions that are worth building a research agenda around. I'm fortunate to have had him as a coauthor, mentor, and role model. Kelly Rader pushed me and my classmates to elevate our standards for engaging with others' work. John Roemer's courses expanded my mind in a valuable way. My formal work will never touch his, but his teaching gave me the courage and commitment to try. Sue Stokes offered valuable advice and encouragement as I developed a very early version of these ideas into my second-year paper and first academic publication. Jacob Hacker and David Mayhew each offered valuable mentorship, not to mention excellent courses, when I was working to find my way. In the early going, Libby Wood was a welcoming presence and an attentive DGS.Others who always had their door open for me include Peter Aronow, John Henderson, Eitan Hersh, Josh Kalla, Christina Kinane, Limor Peer, Frances Rosenbluth, Ian Shapiro, and Ian Turner.

Another debt of thanks goes to all of the others who read and commented on this work during its development. Among those who read and provided valuable comments on some part of the work presented in these pages, but are not listed above, are Nicky Bell, Emily Benedicto, Scott Bokemper, Alex Held, Annabelle Hutchinson, Trevor Incerti, Jenn Jerit, Paul Merklinger, Lilla Orr, Kyle Peyton, Eli Rau, Shikhar Singh, Ira Soboleva, Gaurav Sood, and Omer Yair.

All of the ISPS and Rosenkranz/MacMillan staff deserve my gratitude, including many people

I don't know. I especially want to thank Colleen Amaro, Lani Colianna, Mary Sue Fitzsimons, Pam Greene, Tom Hallihan, Pam Lamonaca, and Cheryl Tappin for being friendly faces and easy to work with.

This is a part of my life I'll look back upon fondly. I owe that to all my people. Here where I stand in Brooklyn, I've been fortunate to have a great set of longtime friends and to have been welcomed into the Emergency Medicine family at Kings County / SUNY Downstate. I'll especially remember popping by Enbani and Ben's old place on 3rd Avenue on the way back from Grand Central after an intense couple days in New Haven, and biking over to Mo's late to bond over our shared tendency to over-commit.

My classmates are what makes me miss the program, which ended in many ways when the pandemic got real in March and I stopped going to New Haven. Tyler Bowen was a great roommate and co-host in the first two years. Lilla Orr was a great first co-author and always wonderful to be around. David Froomkin was always good for an engaging late-night conversation in years three through five. Trips to Three Sheets with Kyle Peyton provided a much-needed respite from otherwise-frenetic weeks. Shikhar Singh was always a good bet for an interesting hallway chat. Annabelle Hutchinson was great at getting me to relax a little bit. Eli Rau can 3 a.m. nerd-out with the best of them. In the later years, I enjoyed developing better friendships with Trevor Incerti, Darren Nah, Cleo O'Brien Udry, James Sudquist, Jen Wu, and Sam Zacher. A chance encounter with Nica Siegel, Stephanie Almeida, or Charlotte Hulme would always brighten my day. Stephen Moncrief was a constant in the carrels. I loved the softball team, and though that's been a collective effort over the years, it was Nate Grubman and Louis Wasser who had it going when I started. Great neighbors like Nick Frisch and Miriam Rich gave me another thing to look forward to, as did sometimes-roommates Will Ratoff and Matt Tyler. Having good people on both ends of the journey made it a lot easier to grab the two-city lifestyle by the horns.

My family was a source of strength and support in this time, and a grounding in what really matters in life. Nobody has done more to support me this time than my wife Angela, who provided the encouragement I needed to keep at it and a good reason to close the laptop every now and then. Her support and understanding were crucial day in and day out. The only people I truly couldn't have done this without are my parents, and I don't only mean that literally. They built me up, gave me skills, and kept holding me to a high standard for years when it would have been easier to throw up their hands. I'm especially grateful to my mom for giving up so much to make my life better, and to my dad for feeding me complex ideas and teaching me to write and rewrite and rewrite from a young age. I owe them everything.

Chapter 1

Introduction

> "It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so."
>
> — Unknown

Survey data are pervasive in contemporary accounts of politics in academic work and in the public sphere. A consumer of these accounts cannot help be struck by the number of things people seem to have beliefs about, and the consistency with which partisanship seems to predict what people believe. Documenting belief in the wrong answer to survey questions, and what predicts these apparent beliefs, has become a cottage industry in contemporary accounts of politics.

Sometimes, misperceptions and partisan belief differences seem to have real consequences. Misperceptions of the characteristics of out-group members seems to drive partisan hostility (Ahler and Sood 2018) and antipathy toward immigrants (Grigorieff et al. 2020). False rumors about death panels seemed to have a disruptive effect on health care reform debates at the outset of Barack Obama's presidency (Nyhan 2010, 2020). Today, partisan differences in behavioral responses to the COVID-19 pandemic (Grossman et al. 2020) are quite plausibly attributable to false and misleading claims by partisan opinion leaders. Rumors of 2016 Democratic presidential candidate Hillary Clinton's involvement in a global child sex trafficking ring motivated one armed man to drive hundreds of miles to investigate, and have metastasized into a larger QAnon conspiracy movement whose adherents defy public health measures and appear poised to win seats in Congress.

Many of these cases furnish concrete evidence of people acting on the beliefs in question. From this, we can tell that at least some people really believe false claims, and that those false beliefs can have real consequences. Yet in the survey environment, talk is cheap. There is a great deal of distance between choosing the wrong answer to a survey question and acting on one's beliefs in the real world. Consequently, it is difficult to tell whether survey measures that suggest widespread misperceptions actually support generalizing from the concrete anecdotes that fill the introductions to survey-based work on political misperceptions.

The vagaries of survey technology do not stop researchers from using surveys to generalize more broadly about the public's beliefs regarding facts and rumors. This leads to some fairly strong claims based on survey data, in both academic work and the public sphere. Bullock and Lenz (2019) note the prevalence of a "simple interpretation of survey results among journalists and pollsters: People have decided beliefs, and surveys elicit those beliefs" (326). As a consequence, leading newspapers tell their readers that "Americans believe a lot of dumb, crazy, destructive, provably wrong stuff" (Rampell 2016). Headlines boldly proclaim that "52% of Republicans actually think Trump won the popular vote" (Oliver and Wood 2016) or that "53 percent of Republicans think the unemployment rate has risen under Obama" (Ross 2015). Consumers of academic literature are told that people who say that Obama is a Muslim and that the attacks of September 11, 2001 were an inside job "truly believe" these false claims (Berinsky 2018, 211).

These interpretations fly in the face of decades of research on the nature of responses to attitudinal survey questions. Rather than directly eliciting respondents' beliefs, authoritative accounts of the survey response hold that people construct their answers on the spot based on considerations that are often only tangentially related to the topic at hand (Tourangeau et al. 2000; Zaller 1992). Put simply, when people don't have a belief to report, they're quite willing to hazard a guess.

Aware of this, survey researchers frequently deploy techniques that are designed to deal with respondents' willingness to guess. The most common invites respondents to opt out of stating a belief by saying some variant of "don't know" or "not sure." Others recommend a higher standard: only those who claim to be very confident or very certain should be said to believe their answer (Kuklinski et al. 2000; Pasek et al. 2015). Comforted that some such measure has been taken, the researcher then proceeds to treat all survey responses as though they represent whatever conception of "belief" is most convenient to the inquiry.

Little is known about whether these strategies succeed. This creates the possibility that techniques for measuring misperceptions do not isolate the sort of beliefs that observers have in mind. By extension, it is possible that existing research substantially over-states the prevalence and depth of misperceptions. Scholars of misinformation typically use the quote in the precis to highlight the pitfalls of holding a misperception. Here, it raises another question: do unknowns about measurement technology encourage researchers to make false and misleading inferences about the nature of misperceptions and partisan belief differences? Do surveys trick researchers into thinking they know what just ain't so?

This volume answers in the affirmative. After introducing a framework for thinking about the senses in which respondents do and do not believe their answers to survey questions, I show that prevailing survey technologies lead researchers to over-state the prevalence of misperceptions, misstate the sense in which people believe their answers, and as a consequence, exaggerate both the size and nature of partisan belief differences. The chapters build this case three parts. First, Chapter 2 clarifies clarifying ambiguities over the sense in which existing research shows that respondents generally believe their answers to survey questions. Next, Chapters 3 and 4 examine the properties of DK response options and certainty scales, building an objective basis to assess the sense in which these measures capture genuine variation in beliefs. Partisanship enters only after this foundation has been developed. On this basis, Chapters 5 and 6 provide empirical support for the headline findings.

The collective result is a framework that problematizes prevailing practices in survey research and, at the same time, provides objective measures that can be used in efforts to reduce or eliminate the problems it identifies. Before diving deeper into this approach and its findings, it is worth dwelling on the stakes.

The Nature of Public Ignorance

Among the essential features of democratic political systems is the ability of voters to hold politicians accountable for their policies, actions, and performance in office.[1] Pioneering research in American political behavior cast doubt on the public's ability to carry out this function effectively, noting widespread ignorance on basic elements of politics like the identities of key figures, institutional rules, and the policy positions taken by parties and politicians (Berelson et al. 1954).

Relative to the ideal of a perfectly informed citizen, ignorance of the truth and belief in falsehoods are both signs that members of the public fall short. Yet scholars have long recognized that the ideal of a perfectly informed citizen is not realistically attainable. Downs (1957) famously argued that public ignorance is rational: as the average voter cannot be expected to affect the outcome of elections and has plenty else to worry about in their life, the costs of learning about politics often outweigh the benefits.

In the face of widespread public ignorance, scholars defended the prospect of effective demo-

[1]For a defense of democracy that does not depend on these mechanisms, see Przeworski (1999).

cratic accountability by arguing that perfect knowledge is not necessary for majorities to be right most of the time. Relying on heuristics and their day-to-day observations of the world, voters may be able to form perceptions that, on average, are more accurate than not. Theories of retrospective voting hold that even attention to noisy signals like one's own pocketbook allowing even a relatively ignorant electorate to hold politicians accountable for economic performance (Key 1966; Fiorina 1981). The broader intuition that widespread ignorance mixed with a bit of knowledge is sufficient for effective accountability is known as the miracle of aggregation: as long as even a small subset of citizens had access to correct information, majorities should be expected to be right most of the time (Converse 1990; Page and Shapiro 1992).

The fundamental concern raised by misperceptions and partisan belief differences is that shortfalls in knowledge are not merely evidence of ignorance, but evidence of systematic biases in the information members of the public consume and the way they process it. Kuklinski et al. (1998, 2000) and Hofstetter et al. (1999) called attention to the prospect that many citizens may be misinformed. This raises the possibility that partisan perceptual differences are not only reflective of knowledge gaps, but of systematic exposure to false or misleading information about political reality.

When perceptual errors exist, they often appear to be related to citizens' partisan preferences. Bartels (2002) showed that Republicans were more likely to correctly answer questions about declining unemployment and inflation under President Ronald Reagan, a Republican. Similarly, Democrats were more likely to correctly answer questions about declining crime and budget deficits under President Bill Clinton, a Democrat. Subsequent research documented patterns of partisan response differences in a wide range of substantive areas, including the economy (Conover et al. 1986; Evans and Andersen 2006; Parker-Stephen 2013; Stanig 2013), war (Kull et al. 2003; Prasad et al. 2009; Jacobson 2010; but see Gaines et al. 2007), health care and other social conditions (Flynn 2016; Graham 2020), and politicized rumors and controversies (Krosnick et al. 2014; Berinsky 2018).

This generalized pattern of partisan differences suggests that ignorance may not be a benign force that can be overcome by the aggregation of diverse views. Evans and Pickup (2010) argue that partisan-tinged economic perceptions "revers[e] the causal arrow" (title): whereas the retrospective voting account assumes that correlations between perceptions and political preferences mean that citizens tend to prefer the party that they view as delivering better performance, systematic perceptual differences between Democrats and Republicans suggest that political preferences determine perceptions. To the extent that this is true, democratic accountability is likely to be weakened. Voters are unlikely to reward the other side's good performance, or punish their own side's bad performance, if they never observe it.

In addition to undermining performance-based accountability, misperceptions seem to affect what people think of government policies and of one another. Gilens (1999) argues that misperceptions of welfare beneficiaries drive opposition to welfare despite high support for aid to the poor. To gain causal leverage on these factors, a growing research program randomizes information that corrects apparent misperceptions. Based on this model, researchers have found that correcting misperceptions of immigrants boosts positive sentiment toward immigrants and increases support for liberalizing immigration policies (Grigorieff et al. 2020),[2] that correct information about the foreign aid budget reduces support for cutting foreign aid spending (Scotto et al. 2017), that correct information about the parties' views on policy reduces policy extremism (Ahler 2014), and that correcting misperceptions of the demographic characteristics of parties reduces Democrats' and Republicans' antipathy toward one another (Ahler and Sood 2018).

The leading challenge to the evidence of widespread misperceptions and partisan perceptual differences is evidence that survey respondents do not always reveal their beliefs accurately. Bullock et al. (2015) and Prior et al. (2015) show that when survey respondents are offered modest financial incentives to correctly answer survey questions about the economy and other statistics, partisan differences shrink by about one-half. Schaffner and Luks (2018) show that many respondents intentionally provided the incorrect answer to survey questions about the size of President Trump's inauguration crowd. This phenomenon has been termed *expressive responding*.[3]

A less-studied issue concerns the degree to which partisan belief differences are inflated by uncertainty. Suppose that all Democrats and Republicans think their party usually delivers better economic performance, but are otherwise ignorant; consequently, when asked how unemployment has been trending under a Republican president, each Republican assigns probability 0.51 to a decrease while each Democrat assigns probability 0.51 to an increase. When asked a survey question about the state of the economy, each respondent states their best guess: all Republicans say the economy is good, and all Democrats say the economy is bad. By failing to measure respondents' uncertainty about their beliefs, the survey inflates the partisan belief difference from 2 percent to 100 percent. This gives the appearance that massively different perceptions of economic conditions had a huge effect on the political preferences, when in fact it was a mix of political preferences and misleading

[2] Grigorieff et al. (2020) also find support for the finding by Hopkins et al. (2019) that providing only corrective information about the number of immigrants does not have much effect.

[3] There is some dispute about the mechanism behind expressive responding. Bullock et al. (2015) emphasize the prospect that expressive responding is due to cheerleading, or intentionally choosing an incorrect answer. Schaffner and Luks (2018) provide convincing evidence that this sometimes occurs. Expressive responding may also occur due to biased sampling, or the relative ease of recalling congenial information when forming a response to a survey question (Jerit and Zhao 2020). Regardless of the mechanism, concern about expressive responding boils down to the possibility that responses to factual questions do not represent the respondent's true best guess based on what they have actually perceived about the world.

survey technology that created the illusion of a massive perceptual difference.

These methodological challenges make it hard to know whether misperceptions are widespread or real, whether partisan differences are real or small, and whether informational experiments should be interpreted as correcting pre-existing misperceptions or as simply providing respondents with information about which they were previously ignorant. Consequently, getting a handle on how to measure misperceptions could have broad impacts for how researchers understand the scope and consequences of misperceptions and partisan divides.

Overview of the Volume

To set the stage for the empirical chapters, Chapter 2 links classic and contemporary arguments over whether surveys measure attitudes and beliefs. Current scholarship is characterized by two distinct conceptions of a "belief." The best-known argument that surveys do not generally measure beliefs, Converse's non-attitudes thesis, embraces a threshold conception: a belief is a pre-existing judgment that a respondent feels sufficiently certain about. In arguing that surveys do measure beliefs, Converse's critics define beliefs probabilistically: beliefs are latent probability distributions over response options, and individuals' tendency to favor some response options over others suggests the presence of meaningful, if, latent beliefs.

This raises a fundamental tension for the study of political misperceptions and partisan belief differences. The consensus measure of a political misperception — being certain of an incorrect answer (Kuklinski, Quirk, Jerit, Schwieder and Rich 2000; Pasek, Sood and Krosnick 2015; Flynn, Nyhan and Reifler 2017; Jerit and Zhao 2020) — embraces a threshold-based conception of what it means to believe an incorrect answer, but there is no evidence that surveys measure this type of belief. Survey-based estimates of partisan belief differences, which at most deal with respondent uncertainty by offering a DK response option, implicitly make even stronger assumptions about the extent to which respondents believe their answers. But in contrast to the long tradition of interrogating the nature of attitudinal survey responses, there is little evidence one way or to the other as to how well these tools work for measures of beliefs. The purpose of this volume is to begin to develop such an evidentiary basis.

Chapter 3 examines the most common approach to dealing with respondent uncertainty: explicit "don't know" (DK) response options. Allowing respondents to say DK mostly removes very uncertain responses from the sample. This leaves a substantial amount of uncertainty among respondents who answer the question. The fact that a DK option was offered does not mean that respondents who answer think they know the answer, or "believe" it is true in the threshold sense. Instead, it means that respondents were able to assign enough probability to one option or the other that they felt comfortable making a guess.

For the study of misperceptions and partisan belief differences, these findings show that DK response options do something useful: they filter out many of the least-certain responses. The remaining responses mostly reflect meaningful beliefs in the probabilistic sense. Respondents who choose to answer a question rather than say DK believe their answer more strongly than one ought to believe that a coin will land on heads or tails, but perhaps not any more strongly. The flip side of this nice property is that DK response options do not isolate people who "believe" their answer in the high-threshold sense. People who do not say DK when given the option may not be flipping mental coins, but they are also not revealing the pre-existing acceptance of false claims that their responses could imply at face value.

Chapter 4 examines an alternate strategy advocated by critics of DK response options: measuring individual-level uncertainty. Luskin et al. (2018) call the use of certainty scales a "24 carat gold standard" for identifying respondents who are misinformed. In the real world, however, the gold standard for identifying beliefs is behavior: talk is cheap, but one who really believes something reveals it through their actions. To examine the degree to which certainty scales capture genuine differences in respondents' degree of belief, the chapter makes use of two expectations for how people should behave in the survey environment if their responses represent things they really believe: they should be stable in their responses, and they should be willing to put their money where their mouth is.

Applying this strategy to classic political knowledge measures designed to capture general awareness (e.g., which party controls the House of Representatives?), Chapter 4 finds that certainty scales do an excellent job of capturing variation in the strength of respondents' belief in their answer. Respondents who claim to be more certain consistently state the same response when asked again in a follow-up survey, and consistently reveal strong belief in their answer through costly discrete choices. A third strategy suggests that for these questions, the variation captured by certainty scales can fairly be characterized as variation between knowledge and ignorance. Respondents who indicate greater certainty are also consistently more accurate in their responses.

If certainty scales capture meaningful variation in knowledge and ignorance, might they also capture meaningful variation between ignorance and the "incorrect knowledge" that would result from believing false claims? Chapter 5 applies the same strategies to the measurement of misperceptions. In contrast to the encouraging results for political awareness questions, the results offer substantial reason to doubt respondents' claims to be certain about false statements. Incorrect answers to factual questions about the economy and politicized controversies consistently fall well short of the benchmarks set in the previous chapter. Even respondents who claim to be certain about incorrect answers are fairly unstable in their responses. When financial stakes are attached to the measure of stability, the partisan valence of the false statement does not alter these results.

These results imply that the typical respondents who chooses the wrong answer and claims to be certain about it is not misinformed about the fact itself; instead, such respondents should be viewed as making "miseducated guesses" based on misleading heuristics. In the survey environment, claims to be certain that President Obama never released his birth certificate are no more robust than claims to be certain that the wrong party controls the House of Representatives. Of course, it is possible that these results do not generalize to every misperception about which researchers have ever asked. For researchers who suspect that their favorite question format or survey item captures more genuine variation in belief in the lies, Chapter 5 provides a template for assuming the burden of proof.

Chapter 6 examines the implications for measuring partisan belief differences. It begins by showing analytically that the prevailing practice of not measuring respondents' uncertainty about their answers is not neutral vis-a-vis respondents' uncertainty about their beliefs; instead, standard research practices treat respondents exactly as if they believed their answers with complete certainty. This introduces *certainty bias* into estimates of partisan belief differences. Though incorporating individual-level uncertainty into estimates of belief differences may sound straightforward, the chapter shows that doing so requires careful attention to the alignment between survey measures, the probability distributions they imply for individuals' beliefs, and the estimators for calculating partisan belief differences. The lack of such attention is to blame for the fact that widely-recognized concerns about respondents' uncertainty remain "deeply underappreciated" in empirical research on partisan belief differences (Bullock and Lenz 2019, 337).

By treating all respondents who express a belief as if they are completely certain of it, prevailing survey measures systematically exaggerate belief differences, by 40 percent on average. This varies considerably across questions; in a few cases, taking responses at face value actually *underestimates* partisan divides. These problems are systematically related to question content. Partisan divides over the state of the economy are inflated by about twice as much as partisan divides over politicized controversies. All of these results are robust to the use of measures that infer respondents' beliefs through costly choices, demonstrating that the effect of uncertainty is empirically distinct from the effect of expressive responding.

In addition to distorting the size of belief differences, surveys paint a false portrait of their nature. Respondents express substantially more certainty about correct answers than incorrect answers, regardless of party. Paired with the findings in Chapters 4 and 5, this helps demonstrate that in general, partisan differences belief differences should be viewed as evidence of differential knowledge of convenient and inconvenient truths rather than as evidence of misperceptions.

None of this is to say that misperceptions and partisan divides do not raise serious concerns, where they in fact exist. Instead, the point is that as they are generally conducted, surveys provide a little basis to distinguish between which false beliefs and partisan divides are real and which are artifacts of survey technology. Through their tendency to see the same problems around every corner, surveys both exaggerate observers' sense of the problems and dull observers' ability to detect where they are most real and most severe. Given public opinion surveys' growing centrality in portraits of politics from the academy to public-facing journalism, this is no mean set of flaws.

As is true of any collection of studies, this volume has limits on its scope. It does not examine nearly every survey question or measurement technology that has been used in the study of misperceptions and partisan belief differences. Most of the data come from convenience samples. While it is certainly possible that some other question or measure does a better job, the findings highlight that the burden of proof for stronger interpretations lies with the researcher. The approach taken in this volume is a roadmap for assuming that burden.

Chapter 2

The Two Meanings of "Beliefs"

Abstract. Scholars have produced extensive empirical evidence about the extent to which respondents believe their answers to survey questions, but its interpretation has been hampered by latent disagreement over the meaning of "beliefs." This chapter considers the role that two conceptions of belief have played in prior work on factual and attitudinal beliefs. The *threshold* view defines beliefs as statements respondent are fairly certain about. The *probabilistic* view defines beliefs as personal probabilities that respondents assign to each response option. In the literature on attitudes, scholarship has converged on the probabilistic view as the only defensible sense in which survey responses represent "true attitudes." The chapter concludes by drawing out this conclusion's implications for measuring factual beliefs, which will be tested empirically in the coming chapters.

Scholars have produced extensive empirical evidence about the extent to which respondents believe their answers to survey questions, but the substantive interpretation of this evidence has been hampered by a latent disagreement over the meaning of "beliefs." Uses of the term in public opinion research can be usefully characterized as embodying two definitions. The *threshold view* defines a belief as a statement that one is sufficiently certain about, or to which they assign a sufficiently high probability of being true. In this conception, there is a sharp distinction between those who have a belief and those who do not. The *probabilistic view* views beliefs as latent probabilities respondents assign to each response option. In this conception, beliefs are imagined probability distributions that respondents construct from the considerations they call to mind upon reading survey questions.

This chapter grounds this distinction in conceptual accounts of beliefs, then uses it to clarify the substantive interpretation of empirical evidence in two important areas of the literature: the debate over whether responses to factual questions reflect respondents' factual beliefs and the response stability debate touched off by Philip Converse's "non-attitudes" thesis (1964; 1970). In both cases, the empirical evidence suggests that while respondents usually do have probabilistic beliefs that are relevant to survey questions, they often do not "believe" their answers in the threshold sense. The simple framework presented here provides a basis for understanding and comparing the conceptual targets of different empirical methods. In doing so, it provides a conceptual foundation for the empirical studies that follow.

Two Definitions of Beliefs

This section lays out a conceptual framework for understanding how scholars talk about, and interpret evidence about, the extent to which survey respondents believe their answers. Like any stylization, it does not capture all perspectives perfectly and does not make every distinction one could make. Instead, it focuses on one particular distinction that illuminates important conceptual differences between empirical strategies and interpretations of the resulting evidence.

The Two Definitions

The probabilistic view of belief is best exemplified by Bayesian statistics[1] and is common in formal models of behavior. Beliefs are probability distributions over a set of states of the world or attributes of an object. Probability mass is allocated to these world states in proportion to the likelihood that each point is the true value. For example, one's beliefs about the numerical value of the unemployment rate are a probability distribution over all of the values of the unemployment rate that one considers possible. Some beliefs are very certain (represented by a low-variance distribution that assigns most or all probability to one response option)[2] and others are very uncertain (represented by a distribution that assigns equal probability to all possibilities).

The threshold view of beliefs is common in philosophy and in ordinary dictionaries. By belief, philosophers generally mean "the attitude we have, roughly, whenever we take something to be the case or regard it as true" (Schwitzgibel 2015). The Cambridge English Dictionary offers a similar definition of belief: "the feeling of being certain that something exists or is true."[3] According to this view, people believe a multitude of statements or propositions. Whatever set of statements one "believes" constitute one's beliefs.

[1]The Bayesian perspective is an especially good analogy for the probabilistic view because Bayesians use the word "beliefs" to refer to a probability distribution. As elaborated below, both conceptions are Bayesian in that they treat belief as a property of the belief-holder.

[2]Though the amount of probability mass at the mode of a distribution does not correspond 1:1 to the variance, the relationship is close enough that in conceptual discussions, it is reasonable to loosely use term "high certainty" or "low uncertainty" to mean "there is a lot of probability mass in one region of the imagined belief distribution."

[3]The Oxford English Dictionary offers two definitions: "An acceptance that something exists or is true, especially one without proof" and "(belief in) Trust, faith, or confidence in (someone or something)." Merriam-Webster's offers three: "a state or habit of mind in which trust or confidence is placed in some person or thing," "something that is accepted, considered to be true, or held as an opinion," and "conviction of the truth of some statement or the reality of some being or phenomenon especially when based on examination of evidence."

Figure 2.1: The two uses of "beliefs"

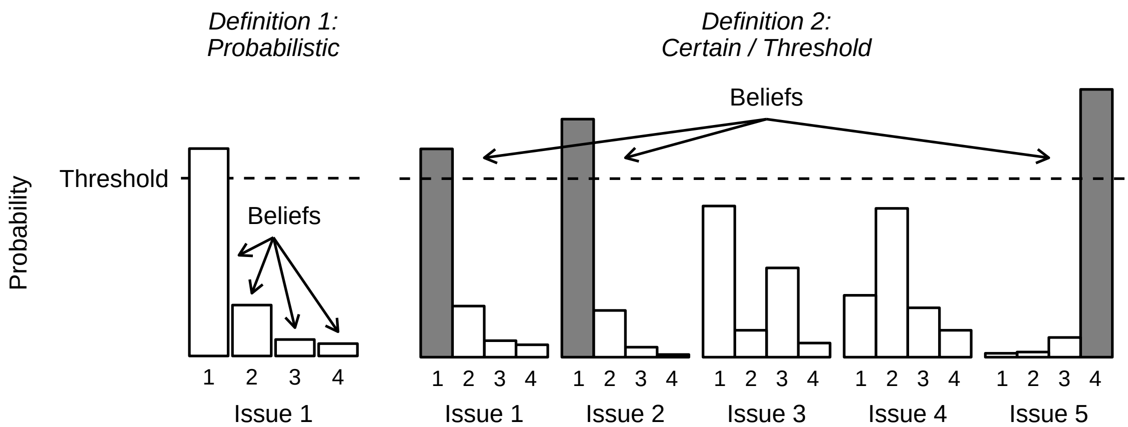

What amount of probability or certainty is sufficient to justify the claim that one "believes" a statement? Borrowing a term from philosophers who view beliefs as varying in their degree,[4] the *threshold* level of probability or certainty determines what can be called a belief (Kaplan 1996; Wedgwood 2012). Below the threshold, one is insufficiently confident in the statement at hand for it to be fairly called a belief; above the threshold, one believes the statement. This sufficiency requirement introduces vagueness into the threshold conception of belief; philosophers generally regard the choice of a precise threshold as arbitrary (Stalnaker 1984; Foley 1992; Weatherson 2005). This vagueness notwithstanding, in public opinion research, empirical practice and implied definitions of "belief" commonly embody the threshold notion. The threshold becomes whatever standard the researcher considers sufficient to say that the statement implied by the respondent's answer choice is a true attitude, genuine belief, real opinion, or other combinations of these words. "Don't know" (DK) response options, opinion filters, and definitions of misperceptions that focus on respondents who are sufficiently certain of their answer all embody the threshold notion of belief.

Figure 2.1 visualizes the two definitions. The leftmost histogram exemplifies the probability distribution view: "beliefs" refers to the probability distribution over the exhaustive set of propositions $\{1, 2, 3, 4\}$. The other five histograms exemplify the high certainty / threshold view. On issues 1, 2, and 5, one proposition is endorsed with sufficient certainty to be called a belief; collectively, these are referred to as beliefs. On issues 3 and 4, no single proposition is endorsed with sufficient

[4]Among philosophers, the mainstream view of belief does not "imply any uncertainty ... as it sometimes does in ordinary English usage" (Schwitzgibel 2015). This may be traceable to analytic philosophy's roots in conceptions of human reasoning that emerged centuries before the notion of probability (Oaksford and Chater 2007, chapter 1). Most social scientists and more mathematically-minded philosophers (e.g. Huber 2009) prefer to take for granted that beliefs vary in their degree. In the framework here, definitions that do not admit uncertainty can be thought of as a special case of the threshold conception, where one must be completely certain in order to believe.

certainty. To the threshold thinker, the probability distributions on issues 3 and 4 are not beliefs; to the probabilistic thinker, they are beliefs.

Another way to look at the threshold is as a minimum standard for generalization about the probability distribution. In Figure 2.1's issue 5, so much probability mass is assigned to option 4 that the whole distribution can be fairly summarized by that single interval. Consequently, on a survey that allowed those same four response options, it would be possible for the respondent to choose one answer that is a good summary of their probabilistic beliefs. On issue 4, the probability mass is insufficiently concentrated for any one interval to be a fair representation. To suggest that the respondent "believes" any one of the four options would not do justice to the amount of probability assigned to the other three.

Beliefs and Survey Responses

Standard practice in public opinion research is to ask respondents to select a single answer from some structured set of response options. Out of a legitimate desire to encourage respondents to focus their responses on the same construct, survey researchers almost always place some form of structure on responses. This is most evident in multiple choice formats, but even open-ended questions like the standard American National Election Study (ANES) question, "what is John Roberts' job or political office?," instruct respondents to name a profession.

This feature forces survey researchers to confront the possibility that respondents' answer choices can imply statements that they never held in their mind before entering the survey environment. Decades of research on survey response behavior — including temporal instability, question order effects, and the tendency to express attitudes and beliefs about obscure and fictitious issues (Bishop et al. 1980; Converse 1964, 1970; Schuman and Presser 1981; Oliver and Wood 2014) — led scholars to abandon the view that survey questions extract "decided beliefs" or "fixed stances" akin to pulling a file from a drawer (Bullock and Lenz 2019, 326; Berinsky 2017*a*, 317). Instead, the beliefs measured in surveys are constructed on the spot. Respondents sample information from memory, integrate it into a summary judgment, and express that judgment in the researcher's chosen format. In political science research, this view of the survey response is most closely associated with Zaller (1992). In the larger literature on the survey response, similar accounts abound (e.g., Nadeau and Niemi 1995; Schuman and Presser 1981; Strack and Martin 1987; Sudman et al. 1996; Tourangeau et al. 2000).

If respondents are, as Zaller puts it, "making it up as [they] go along" (76), is it ever reasonable to suggest that respondents believe their answers in the threshold sense? It depends on the content of the considerations. If considerations are sufficiently relevant to the question, it may be reasonable to suggest that respondents are fairly certain about statements or propositions that they have never thought about in precisely the survey's terms. Deduction and inference are fundamental in many accounts of belief and knowledge, from age-old ideas like Aristotle's syllogistic logic (Smith 2017) and Bayes' theorem (Joyce 2003) to contemporary accounts in psychology (Eagly and Chaiken 1993) and philosophy (Langley 1933; Wedgwood 2012).

To make concrete the connection between inferences and belief in survey answers, consider the unemployment rate, which the Bureau of Labor Statistics (BLS) estimates was 3.8 percent in February 2019. A respondent does not know precisely this fact might be fairly sure that the economy has steadily improved since the recession ten years ago, that the unemployment rate is historically low, and that 5 percent or less is very low.[5] Presented in a survey with the opportunity to state "unemployment is between 3 and 5 percent" or "unemployment is lower than it was ten years ago," a reasonable respondent could be fairly certain about their answer. But if asked to state the value to a hundredth of a percentage point, the response format is too specific to elicit something one could reasonably deduce, and be highly certain about, from the set of beliefs just given.[6]

In the probabilistic sense, single-choice responses are reflections of beliefs. Confronted with a question, respondents go through the same process considering what they know. The difference is that rather than report their complete probabilistic beliefs—the entire probability distribution—respondents are asked to report the one-number summary, or choose one of several response options, that best-reflects their underlying beliefs. More concretely, imagine that in a series of parallel universes, the hypothetical respondent just described was asked many times to state the unemployment rate to the hundredth of a percentage point. The resultant series of uncertain guesses—3.78, 5.91, 4.07—could reasonably be modeled as draws from a probability distribution over the values of the unemployment rate that the respondent finds plausible based on what they know. The respondent's answer reflects their probabilistic beliefs because they used real, topic-relevant thoughts to form a best guess about the question.

Facts and Attitudes

Before applying this framework to existing research on both factual and attitudinal beliefs, it is worth being explicit as to why it is valid to apply the same conceptual framework to both cases.

Both the probabilistic and threshold views of belief are consistent with the Bayesian tradi-

[5] This is a stylized description of the open-ended respondent comments analyzed in Graham (2020).

[6] Provided, of course, that the respondent has not downloaded the Current Population Survey themselves and calculated the rate to a higher level of precision than appears in BLS press releases.

tion of imagining probability as a partly-subjective property of the belief holder. To continue the unemployment example, the statement that "there is an 80 percent chance that the unemployment rate is between 3 and 5 percent" is meaningless in the frequentist paradigm—to the frequentist, the unemployment rate is a fixed quantity in the world, and probabilities are properties of the estimation procedure. To speak of the probability that some statement is true, one must admit some subjective element like a Bayesian prior distribution (Greenland et al. 2016; Vick 2002). Even in the realm of facts, then, a respondent's personal probability that a factual statement is correct is subjective.

When it comes to survey questions about attitudes, people also combine considerations to make their choice, but with an eye to what is desirable or morally justified. In this realm, there is also a notion of subjective certainty: people can be very sure of what they think about an issue, but they can also be uncertain, due either to ambivalence or a lack of knowledge of the issues at hand. Consider, for instance, a proposal to withhold federal funding from women's health providers that offer abortion services. A person who supports abortion and other women's health services would likely be very certain of their opposition to the proposal. A person who thinks abortion is wrong but supports other women's health services would be more conflicted—some aspects of the proposal are wrong and others are right, creating uncertainty as to whether support or opposition is more consistent with their values. For a simpler example, one could be uncertain as to if or when abortion is morally justified.

This dimension of peoples' beliefs—subjective certainty and uncertainty—is the focus of the two-definition framework just described. Though there are real differences between matters of fact and matters of opinion, the probabilistic versus threshold distinction has clarifying power in both realms of survey research. This chapter and the next straddle the fact-attitude divide, while the subsequent chapters focus exclusively on matters of fact.

Factual Beliefs

Researchers often interpret incorrect responses to individual survey questions are often treated as measures of false beliefs or misperceptions. Challenges to this view include evidence that people are often uncertain about their incorrect responses (Pasek et al. 2015; Graham 2020), evidence that people give more-correct, less-biased answers in the presence of monetary incentives (Bullock et al. 2015; Prior et al. 2015), and evidence that respondents choose survey responses that imply belief in made-up controversies (Oliver and Wood 2014). Concluding a discussion of this slice of the literature, Flynn, Nyhan and Reifler (2017) observe:

> [T]he notion of a "true belief" is ill- defined given that models of survey response (e.g. Zaller 1992) suggest that answers are typically constructed on the spot. It is therefore not clear to us that the alternate approaches described above would represent an improvement over the status quo, although we encourage future research evaluating the merits of different measurement techniques. (140)

The two-definition framework adds clarity to this picture by showing that leading methods for evaluating respondents' belief in their answers differentially embody the two definitions of beliefs. DK responses and certainty scales have mainly been used to interrogate whether people "believe" their answers in the threshold sense, while accuracy incentives are designed to encourage people to report answers that better-reflect their probabilistic beliefs.

Threshold approaches

Whether the subject is facts or attitudes, recording responses like, "don't know," "no opinion," or "not sure" (DK for short) remains the dominant approach to separating respondents who believe their answers from those who do not. Scholarship on the appropriate use of DK responses, and especially competing positions on the extent to which DK should be encouraged or discouraged, can be viewed as pitting different views of what certainty threshold respondents should and do use when deciding whether to answer a question.

For both knowledge and attitude questions, there exists a long tradition of encouraging respondents who are uncertain about their answers to respond DK. In what has become the canonical treatment of political knowledge, Delli Carpini and Keeter (1993, 1996) recommended encouraging DK responses to "avoid the unreliability introduced by guessing" (1996, 1183). This effectively defines knowledge as a true belief.[7] In a piece discussed more extensively below, Converse (1964) laments that although "large portions of an electorate do not have meaningful beliefs ... virtually none of the common modes of dealing empirically with public beliefs attempts to take it into account." (52). His solution was to encourage DK responses: "[i]nstead of browbeating our respondents into giving opinions they did not feel they possessed ... we explicitly invited respondents who had no opinions on a particular issue to report that fact directly" (Converse 1974, 651).

In both cases, the goal of encouraging DK is to make respondents use their personal threshold for what counts as believing or knowing to decide whether or not to answer the question. If the respondent is fairly sure of what they think about the question, they are encouraged to answer; if they are not certain, they are encouraged to say DK rather than hazard a guess. For many years, the ANES adopted this approach for both attitudinal and knowledge questions.

[7] To philosophers, "true belief" is at best a bare minimum conception of knowledge. See Chapter 4.

A challenge to encouraging DK paired a probabilistic conception of knowledge with the educational testing literature's goal of developing unbiased knowledge scores. Mondak (2000, 2001) argued that encouraging DK responses introduces systematic measurement error. Because people differ in their propensity to guess—in other words, because people use different thresholds to determine whether they are sure enough to answer a question—encouraging DK causes peoples' propensity to guess to affect the number of questions they answer correctly. Discouraging or eliminating DK response options reduces or eliminates the systematic measurement error (the influence of the propensity to guess) at the cost of additional unsystematic measurement error (whether guesses are correct or incorrect). Delli Carpini and Keeter (1993) recognized this issue but considered guessing-induced response instability to be "the more serious problem" (1183).

Subsequent research suggests that while encouraging DK hides partial knowledge, respondents who choose DK under "neutral" approaches that neither discourage nor discourage DK are (on most questions) no more likely to answer correctly than chance (Luskin and Bullock 2011; Sturgis et al. 2008). In the probabilistic sense, this is evidence that when DK is *not* encouraged, respondents who choose DK do not have meaningful beliefs—it can be assumed that their would-be answers can be thought of as random, equiprobable draws from the response options, as though their probability distribution over the choices is uniform.

In response to this and other evidence that allowing DK responses does not completely filter out guessing behavior, researchers have turned to other means. Kuklinski et al. (1998, 2000) voiced this concern via their influential distinction between misinformation and ignorance, where misinformation is defined as an incorrect survey response stated with a high level of certainty and ignorance is an incorrect response stated with a low level of certainty (see also Hofstetter et al. 1999; Mondak and Davis 2001, Table 1; Delli Carpini and Keeter 1996, 95). To distinguish the misinformed from the ignorant, Kuklinski and colleagues recommend using a higher threshold: only count respondents as being misinformed if they claim to be certain of their answer. This approach has gained a substantial amount of traction in recent years (Flynn 2016; Graham 2020; Lee and Matsuo 2018; Marietta et al. 2015; Marietta and Barker 2019; Pasek et al. 2015; Peterson and Iyengar 2020). Luskin et al. (2018) call it the "24 carat gold standard" for measuring misperceptions.

Probabilistic approaches

Another line of research posits that survey respondents may answer questions using some basis other than their beliefs about the fact in question. In addition to whatever payoff survey-takers reap from answering truthfully, they may find exerting effort unpleasant, register "expressive responses" that make their partisan team look good (Bullock et al. 2015), or "troll" the survey by intentionally choosing absurd responses (Lopez and Hillygus 2018).

One possible solution, proposed by Prior et al. (2015) and formalized and contemporaneously published by Bullock et al. (2015), is to increase respondents' motivation to respond truthfully through verbal encouragement or financial rewards. To maintain focus on the two-definition framework, this discussion looks past possible threats to inference[8] and focuses on the conception of beliefs the strategy embodies.

In their appendix, Bullock et al. (2015) formalize the pay-for-correct approach using an explicitly probabilistic notion of belief. The expected utility (EU) of response j (r_j) is:

$$\mathrm{EU}(r_j) = h(r_j) + I \times p(r_j) + c \times e(T, r_j)$$

where $h(r_j)$ represents the honesty value of truth-telling, I is the incentive amount, and $p(r_j)$ is the probability that r_j is true. The expressive benefit, $c \times e(T, r_j)$, is modeled as a constant times an expressive function of the respondent's partisanship (T) and the response. The authors let $h(r_j)$ equal $p(r_j)$, giving a rewritten formulation

$$\mathrm{EU}(r_j) = (1 + I) \times p(r_j) + c \times e(T, r_j)$$

that makes clearer the incentive's theoretical role. By increasing the influence respondents' probabilistic beliefs have on their survey response, the pay-for-correct approach aims to outweigh expressive utility and random response error (which arguably must be present but has an EU of zero). To the extent that responses change in the presence of the incentive, the conclusion is that the incentive outweighed expressive utility. Were the unincentivized responses purely a function of the respondent's probability distribution over the response options, $p(r_j)$, incentives would not change the response.

Both Prior et al. (2015) and Bullock et al. (2015) find that incentives reduce partisan differences. Prior et al. (2015) also find an increase in correct answers. The upshot from both findings is that in the probabilistic sense, "[s]urvey responses to factual questions may not accurately reflect beliefs" (Bullock et al. 2015, 4). Crucially, this is not a claim that incentives provide any evidence that respondents believe their answers in the threshold sense. Even when incentive designs succeed

[8] The three most important threats are the possibility that incentives will induce information search (Clifford and Jerit 2016), will encourage respondents to try to guess what the researcher believes rather than use their own beliefs, or will be inadequate to outweigh the expressive benefits (for example, Berinsky (2018) notes that his subtle pipeline treatment only eliminated one-fifth of false claims about voting).

in motivating responses to answer sincerely, it does not follow that $p(r_j)$ is large—just that it is larger than $p(r_{-j})$. An absence of expressive responding implies only that the response reflects the respondent's sincere best guess based on their probabilistic beliefs.

Just as Flynn et al. (2017) hypothesize, defining beliefs clarifies the value of alternative approaches to measuring beliefs. DK encouragement and certainty scales are targeted at the threshold conception while incentivized surveys aim to elicit a more accurate reflection of probabilistic beliefs. These approaches complement each other, and are not substitutes for one another, because they approach beliefs in different ways.

Non-Attitudes

The debate over the meaning of response instability—respondents' tendency, on some questions more so than others, to answer the same question differently at different points in time—also implicitly pits the two views of beliefs. In the sense of the two-definition framework, this section shows that Philip Converse's influential non-attitudes thesis is (1) correct in the sense that in many cases, people do not threshold-believe their answer choices, but is (2) incorrect in that in the large majority of cases, survey respondents seem to have probabilistic beliefs that they can apply to the question.

Converse (1964) touched off the response instability debate with his "black and white" model of survey responses. The model imagines that the population can be divided into two groups: one that is absolutely certain of its responses and another whose responses have no meaning whatsoever. "There is first a 'hard core' of opinion on a given issue, which is well crystallized and perfectly stable over time. For the remainder of the population, response sequences over time are statistically random" (49). By statistical randomness, Converse does not mean only that the survey response is stochastic, but that the probability distribution over the response options is uniform: the model posits that these people "are responding to the items as though flipping a coin" (49), which is equivalent to a uniform distribution when questions are binary. In a contemporaneously written essay, Converse labelled this second class of responses "non-attitudes" (Converse 1970).

Taken literally, the black and white model implies that respondents either have beliefs in the threshold sense—a hard core of crystallized, perfectly stable opinion—or are completely lacking in a basis to choose between the response options, lacking meaningful beliefs even in the probabilistic sense.

The most common counter-perspective to the non-attitudes thesis takes aim at the notion that respondents lack beliefs in the probabilistic sense. Empirically, it begins by imagining survey responses as drawn from a probability distribution. Borrowing notation from Ansolabehere et al. (2008), let subject i's observed response be $W_i = X_i + e_i$. The standard assumption is that e_i is normally distributed with variance σ_e^2 is exactly equivalent to an assumption that responses are drawn from a probability distribution that peaks at its mean, X_i.

Based on this model, researchers have cast serious doubt on the idea that survey responses are "random noise." Achen (1975) shows that estimating σ_e^2 suggests a stronger correlation between the theoretical means, X_i, over time than would be apparent if the raw data were treated as error-free. Erikson (1979) finds that related items are often correlated; this common variance could not appear if responses were random noise. Ansolabehere et al. (2008) take advantage of this common variance to create item scales that are more temporally stable than the individual items that constitute the scales. Were respondents drawing from uniform distributions—on binary items, responding as if flipping a coin—these results could not obtain. This constitutes decisive evidence that respondents have meaningful attitudes in the probabilistic sense of belief. Ansolabehere et al. (2008) conclude that responses are not "purely random noise" and reflect "much more content than Converse asserted" (226).

If the non-attitudes thesis is read to imply that all respondents lack meaningful attitudes in the probabilistic sense, it is clearly wrong. Yet it is also possible—and in my view, plainly more consistent with Converse's intent—to read the black and white model as a sometimes-too-bold statement of the point that not all responses are meaningful in the threshold sense. Converse (1964) asserts that only one item, the power and housing item, would pass a goodness-of-fit test for the black and white model. In a footnote, Converse (1964) writes that for the other questions, "a distribution of the population continuously across the total range of response probabilities is entirely compatible with the data" (72). Rather than a literal claim, a close reading of Converse's work makes clear that the black and white model was intended as a bold statement of a then-underappreciated point. Converse (1970) writes:

> This fact of 'near-fit' may be conceptually more important than meets the eye at first glance ... the fact that one set of these data fits the black-and-white model very well, and the other sets of conceptually comparable items only miss a fit with the model in modest degree, suggests that we should not abandon the black-and-white model completely in imagining the processes which underlie the responses to other items in the battery. (175)

This intended purpose of the black and white model often has largely been lost in debates over it, past and present. Describing supporters and detractors of his work, Converse (2000) wrote that "[w]hat both sides had in common was a basic incomprehension of the role of limiting cases in inquiry" (338).

How could one use the black and white model as an imperfect guide, conceptually important but not to be taken as a literal description? Converse's later writing uses different terminology to define response probabilities as the arbiter of a meaningful response. Converse (2000) argues that "item responses are probabilistic over a range or 'latitude of acceptance,' ... [which] are broader or narrower according to variations in what respondents bring to the items. Indeed, when I think of 'attitude crystallization,' the construct refers to the variable breadth of these latitudes" (339). A narrow latitude of acceptance counts as a crystallized attitude—in the two-definition framework, a threshold belief—because most of the probability mass lies on one response option.

Viewed this way, a mild version of the non-attitude thesis holds that not all respondents threshold-believe their answers. This thinking is evident in the 1964 piece. Based on the sort of engagement with open-ended comments that Zaller (1992) later systematized as the analysis of considerations, Converse (1964) argues that "extreme instability is associated with absence of information, or at least of interest, and that item reliability is adequate for people with pre-existing concern about any given matter. The substantive conclusion . . . is simply that large portions of an electorate do not have meaningful beliefs" (51). Because "different controversies excite different people to the point of real opinion formation," one may have "crystallized opinions" on a different set of issues than one's neighbor (53). This language oozes threshold thinking: one has "meaningful beliefs" only to the extent one's belief distribution is sufficiently "crystallized" to produce "adequate" stability.

The measurement error school's tradition of labeling X_i the "true attitude" could be read to refute even this milder version of the non-attitudes thesis. To a thresholdist, "true attitude" could sound like an implication that there exists an absolutely certain, fixed point that the right question could always reveal. Yet labels aside, X_i is just the mean of a probability distribution. Erikson (1979) is especially plainspoken on this point:

> [T]he non-attitude holders' probabilities of a 'pro' response (their mean responses) can actually be considered their 'true' positions. For example, the true attitudes of non-opinion holders on 'power and housing' are assumed to be a 0.586 probability of a 'pro' response. Thus, the term 'non-attitude' is technically a misnomer in the sense that by definition, every respondent has a theoretical mean (true) position. (100)

This is important. When items are binary, the response probabilities that Erikson defines as the true attitudes are *exactly* the response probabilities that Converse wants to use as an arbiter of whether respondents believe their answers. To Converse, the probability distribution over responses determines whether a survey response represents a true attitude. To the measurement error school, the probability *is* the true attitude. The key difference is that where the measurement error school blames the items for failing to elicit this imagined value, Converse wants survey researchers to use response probabilities—or some function thereof, like stability or DK responses—to determine whether responses to individual questions are good reflections of what respondents believe.

In this way, the debate between "non-attitudes" and "measurement error" can be productively recast. For most people on most items, there are real, underlying attitudinal beliefs in the probabilistic sense: people bring considerations to the question that allows them to answer in a more consistent manner than flipping a coin. This is what the measurement error school's evidence shows, and is what Converse admits through his repeated statements, including in his influential 1964 and 1970 pieces, that only one item fits the black and white model. Yet on many items, respondents cannot distinguish between the response options with sufficient certainty that they can reasonably be said to "believe" their answers.

Here, the measurement error school again agrees with Converse, explicitly admonishing researchers not to trust unstable items because they are poor reflections of the respondent's mean attitude (Achen 1975; Ansolabehere et al. 2008).[9] Even if one agrees that the mean of the respondent's imagined probability distribution should be called a "true attitude," survey items are only good measures of true attitudes when the respondent is highly certain about them. If the probability ("true attitude") is 0.586 the available responses, 0 and 1, will always be bad estimates. If the probability ("true attitude") is 0.9, the response will be fairly stable and the usual response, 1 will be a reasonable representation of the underlying belief distribution.

Implications

The literature on attitudes thus sidesteps the fundamental challenge raised by Flynn, Nyhan, and Reifler by embracing a definition of "true attitude" that is consistent with the nature of the data. Survey questions cannot be expected to very often capture Converse's idea of a true attitude, but they can be counted on to measure something meaningful that can reasonably be thought of as a latent probability distribution over the response options. Canonical work on the survey response holds that these probability distributions emerge through an on-the-spot process of integrating

[9] Achen writes: "If a researcher has no idea of the size of these errors, as is often the case, he has little choice but to assume that they are small and then proceed, treating the observed responses as direct measures of the underlying attitudes. The present results remind us that the errors involved may be substantial" (1975, 1231). Similarly, Ansolabehere et al. (2008) advise "caution in drawing inferences from analyses of surveys that focus on single items introduced to capture issue preferences" (228).

whatever pertinent information the respondent is able to call to mind (Nadeau and Niemi 1995; Strack and Martin 1987; Sudman et al. 1996; Tourangeau et al. 2000; Zaller 1992).

The probabilistic conception of belief thus implies a specific sense in which respondents can be generally said to believe their answers: in the same way that one believes any other inference. The notion that beliefs can be deduced from other beliefs is widespread in conceptual treatments. Psychologists who study belief formation allow that "beliefs may be formed by logical deduction from existing beliefs" (Eagly and Chaiken 1993, p. 128). Deduction is taken for granted in philosophical debates over the nature of belief (Foley 1992; Wedgwood 2012) and appears in political scientists' accounts of how people form (false) beliefs about politically-relevant facts (Kuklinski et al. 2000; Flynn et al. 2017; see discussion in Chapter 5). One "believes" the product of an on-the-spot inference in the same manner that a Bayesian believes the posterior that they construct from integrating various pieces of information.

This general description of survey responses does not rule out the prospect that some respondents have a crystallized belief about the question. Instead, Tourangeau et al. (2000) find that

> [T]here is a continuum corresponding to how well articulated a respondent's attitude is. At the more articulated end, the respondent has a preformed opinion just waiting to be offered to the interviewer; at the less articulated end, the respondent has no opinion whatever. Between these extremes, he or she may have a loosely related set of ideas to use in constructing an opinion or even a moderately well-informed viewpoint to draw on. (Tourangeau et al. 2000, 12)

Despite this view of a continuum, Tourangeau et al. (2000) argue that survey responses should be viewed as a product of approximately the same set of processes, regardless of whether the respondent can be thought of as having a stored, pre-existing belief or judgment. "[T]here is little evidence that respondents either retrieve either an existing judgment or more specific beliefs but never both" (18). In an inferential process, a pre-existing judgment can be thought of as an especially precise consideration that the respondent can draw upon. Calling such a pre-existing judgment to mind allows the respondent to form a fairly certain belief about which response option comes closest to their beliefs or attitudes. Chapters 4 and 5 examine the success with which certainty scales capture variation along this continuum.

As noted above, research on factual beliefs about politics uses two main threshold-based strategies to try to isolate which responses reflect misperceptions are believed with a high level of certainty and, consequently, plausibly reflect a pre-existing belief in something that is false — in Tourangeau and colleagues' words, a "preformed opinion just waiting to be offered" (12). The "gold standard" for identifying such beliefs is a certainty scale: researchers ask respondents how sure they are about their answer and only count those who claim to be certain as believing their answer. The more common approach is the “don’t know” (DK) response option, which invites respondents who are uncertain to opt out of providing a response.

To date, little research has examined the success of either strategy. The measurement error school’s conclusion that surveys measure probabilistic beliefs suggests that the DK strategy is unlikely to isolate respondents who are certain about their answer to one option or another. But if respondents’ claim to be certain about their answers can be taken at face value, certainty scales may succeed at isolating a group of respondents who “believe” their answers in a threshold sense.

Chapter 3

"We Don't Know" Means "They're Not Sure"

Abstract. The most familiar approach to handling respondent uncertainty about survey questions, the "don't know" (DK) response, has an inconvenient feature: it records no direct information about respondents' certainty in their answers. This chapter shows that DK responses contain more information than meets the eye. Using 161 survey questions from eleven surveys conducted between 1993 and 2019, the analysis demonstrates that the percentage of respondents saying DK can explain about half the variation in average certainty among respondents who provided an answer (i.e., those who did *not* say DK). The findings provide evidence that DK response options can be viewed as a threshold-based strategy that leaves substantial uncertainty among respondents who provide an answer.

Survey researchers know that not all survey responses are equally meaningful. Though responses sometimes reflect an attitude or belief that the respondent had formed prior to the survey, many responses are constructed on the spot (Zaller 1992; Tourangeau et al. 2000). Tools exist for measuring where responses fall along the attitude/non-attitude or belief/guess spectrum, but limited budgets usually prevent researchers from collecting the data needed to assess response stability (Converse 1964) or certainty (Kuklinski et al. 2000; Pasek, Sood and Krosnick 2015). Inconveniently, the most common approach to measuring respondent uncertainty—recording "don't know" (DK), "no opinion," or "not sure" responses—does not provide direct information as to the certainty of respondents who offered an opinion.

This chapter introduces a way of thinking about this unmeasured uncertainty: the "we don't know" means "they're not sure" heuristic. Even when people who offer an opinion are not given the opportunity to express uncertainty about their responses, the percentage of respondents who say DK offers a clue as to what they would have said. A high rate of DK responses indicates low certainty not just among those who said DK, but also among those who express an opinion.

The results demonstrate this relationship in two sets of survey evidence: 73 questions with certainty follow-ups asked in nine surveys from the 1993-2000 American National Election Studies (ANES) and 88 questions in two original surveys conducted in 2019. Across a wide range of questions, the percentage of respondents saying DK can predict about 50 percent of the variation in average certainty among those who answered the question. The two studies used different survey modes (in-person and telephone versus online) and covered a wide range of political topics.

An approximation that explains 50 percent of the variance of interest is good enough to be useful but leaves room for improvement. A key source of inaccuracy in the heuristic is question-to-question differences in the shape of the distribution of respondent certainty. Some types of questions tend to produce "know it or don't" distributions that are fairly bimodal, while others produce flatter distributions that are indicative of a higher incidence of guessing. When properly channeled, this liability becomes an asset. Because similar questions tend to produce similar certainty distributions, the percentage DK is a substantially better predictor of differences in average certainty between questions of the same type. For example, in a 12-question battery on favorability toward former President Barack Obama, President Donald Trump, and ten contenders for the 2020 Democratic presidential nomination, the percentage of DK responses predicts more than 90 percent of the variation in average certainty among respondents who expressed an opinion.

Despite its remarkable accuracy for some comparisons, the "we don't know" means "they're not sure" heuristic is still only an approximation. Yet given limited survey space and the impossibility of adding certainty measures to surveys conducted in the past, consumers of public opinion data often need to make the most of what information is available. The "we don't know" means "they're not sure" heuristic unlocks a great deal of information as to how strongly survey respondents believe their answers.

For the remaining chapters, the findings presented in this chapter have three key implications. First, DK response options can be thought of as a measurement technology that is rooted in the threshold conception of belief, as defined in Chapter 2. Second, this threshold-based technology does do something useful: DK response options eliminate a subset of the least-certain responses, which subsequent chapters will show also to be the least-stable responses. Third and conversely, DK response options do not isolate a group of respondents who think they know the answer to the question: plenty of heterogeneity in certainty remains after allowing respondents to say DK, both across questions and across respondents. In other words, because DK response options set a low threshold for answering the question, allowing a DK option does not mean that the respondents who answered the question were highly certain about their answer.

Intuition

This section explains why a high rate of DK responses would predict lower certainty among those who answered the question. The explanation rests on two pillars: a hypothesized individual-level mechanism, that the decision to say DK can be thought of as a function of the respondent's certainty level, and a basic feature of mass media societies, that people who live in them hear about many of the same events, issues, and people.

The first pillar is the notion that some threshold level of certainty determines whether respondents say DK or offer a substantive response (i.e., state or select an answer rather than say DK). Think of respondents as engaging in an inferential process in which they call to mind relevant information and attempt to determine which response option comes closest to their views (Zaller 1992; Tourangeau et al. 2000). If the respondent can choose a response with sufficient certainty, they do so; if not, they say DK. In this sense, the subject's response *and* the decision to say DK can be viewed as functions of a latent probabilistic belief the respondent forms after reading the question. If the respondent is sufficiently certain in one answer—for example, if he or she assigns at least a 60 percent probability to one of two response options—then the respondent states that answer. But if the respondent's certainty level is below that threshold, he or she says DK.

Formally speaking, the threshold model of DK responding posits that some personal certainty threshold, τ_i, determines whether a respondent will answer a question. Respondent i answers question j if their certainty level, c_{ij}, exceeds τ_i; otherwise, i says "don't know."[1]

The second pillar is that mass media allow people to hear about about a similar set of public figures, events, and issues. The average American has heard more about Trump than about former Alaska Senator and fringe presidential candidate Mike Gravel; more about abortion than regulations on derivatives trading; and more about trends in the value of the stock market than the percentage of Americans who are immigrants.[2] When asked their opinion of the president, most who respond will be fairly sure of where they stand, but because political engagement varies substantially, a few won't know. By contrast, when asked the same about Montana Governor Steve Bullock, many will say DK and most who offer an opinion will be working from just a couple of details.

These two notions—that people living in the same society hear about a lot of the same things, and that a latent certainty threshold is a reasonable model of the decision to say DK—jointly provide an intuition as to why DK responses signal that other people are guessing. When average certainty

[1]Here, a subtle but important implication of the i subscript is that the threshold may vary across people. Appendix A uses simulations to show that the prediction of a "we don't know" means "they're not sure" relationship is robust to differences in thresholds across individuals.

[2]Study 2 includes questions on each topic in this paragraph.

Figure 3.1: Illustration of why "we don't know" means "they're not sure."

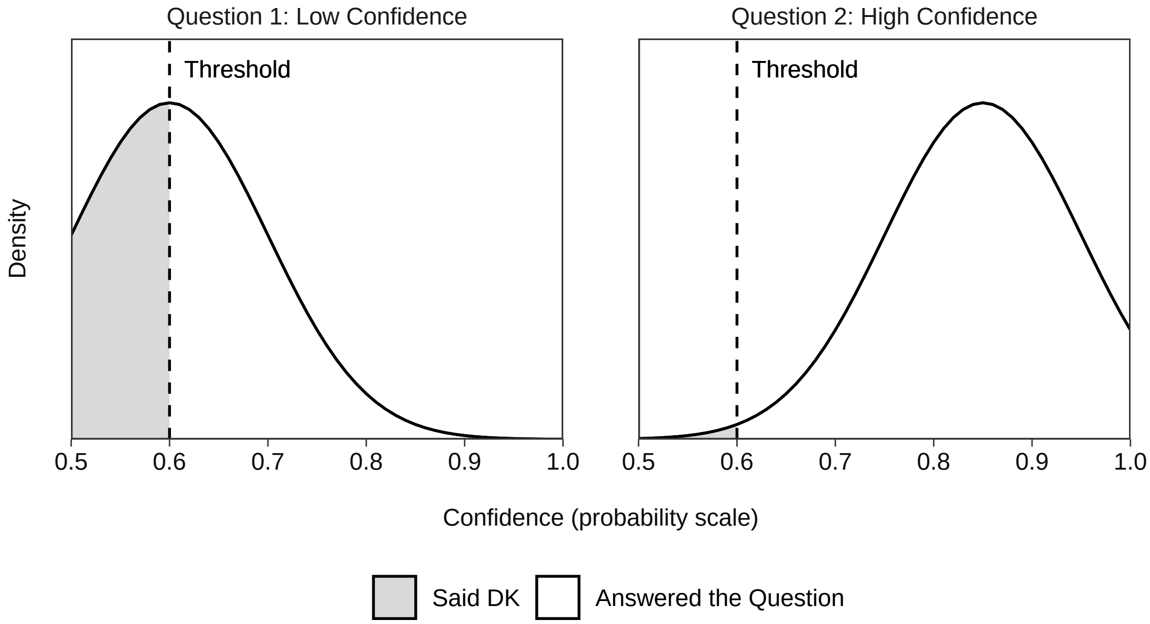

is very high, almost everyone will answer the question. Few peoples' certainty will fall below the threshold for saying DK, but given well-known (and more commonly emphasized) differences in media consumption (Prior 2005), we should not be surprised by scattered DKs about even those things that "everyone" knows. On the flip side, when average certainty is very low, many more peoples' certainty will fall below the threshold, resulting in more DK responses.

To make this more concrete, imagine two questions, each with two response options. For ease of illustration,[3] suppose that everyone has a threshold of 60 percent, answering if at least 60 percent certain and saying DK otherwise. In Figure 3.1, the left panel plots the certainty distribution for a question on a little-known topic. It elicits a lot of DK responses, represented by the large shaded area to the left of the threshold. Despite the low level of public knowledge, plenty of people remain willing to answer, and a few know a lot about the topic. The right panel represents a question that a lot of people know things about. Despite the high level of public knowledge, a few people do not know enough to form an opinion.

If this stylization approximates the real world, we should expect the percentage of DK responses to be a predictor of average certainty among people who do answer. More precisely, we should expect the probability mass to the left of the DK threshold to be predictive of the conditional mean to the right of the threshold.

[3]Appendix A.1 shows that the expectation of a "we don't know" means "they're not sure" relationship does not depend on the assumption of a homogenous threshold.

Relationship to Existing Research

The threshold model is consistent with existing conceptions DK responses. Supporters of encouraging DK responses want to make it clear to respondents that they should not answer unless they really know what they think about the question (Converse 1974; Delli Carpini and Keeter 1993). In other words, they want to raise the threshold high enough, maybe to 80 or 90 percent certainty, that all responses can reasonably be interpreted as true attitudes. On the other side, critics of allowing or encouraging DK responses note that some people are more likely than others to guess when they are uncertain, distorting the observed distribution of knowledge and opinion (Mondak 2000; Mondak and Anderson 2004; Laurison 2015). A threshold probability can also arise as an implication of formal models of the survey response (Bullock et al. 2015, p. 51-53).

Schuman and Presser (1981) suggest a threshold-based model to explain their finding that education and other "general personality or social characteristics" did not predict how DK propensity changes in response to an opinion filter. They posit that the decision to say DK depends "in part on the respondent's position on the DK-propensity spectrum and in part on the height of the barrier created by question form." In this chapter, the barrier's height is the threshold.

This chapter's use of certainty to measure the "DK-propensity spectrum" follows treatments of the survey response by Zaller (1992), Tourangeau et al. (2000), and others: respondents call to mind relevant information and attempt to choose the response that most closely matches what they know or believe. On a knowledge question, one is certain to the extent one's existing knowledge supports the choice; one who knows the specific fact might be completely certain in their answer. On an attitude question, one is certain to the extent that one statement clearly aligns with one's views; ambivalence or ignorance could result in a lack of certainty. Though the nature of the inference is different — one is more like taking a test, while the other incorporates moral reasoning — respondents can be more and less certain which answer represents what they know or believe.

In existing research, the strongest support for the prediction of a "we don't know" means "they're not sure" relationship comes from Dodd and Svalastoga (1952), who found a correlation of -0.91 between the percentage of DK responses and response stability among seven similarly worded items. Smith (1985) also includes the percentage DK in a list of predictors of low response stability.[4] In light of the strong association between item-level certainty measures and response stability (Chapter 4), one might also expect the rate of DK responses to also predict lower certainty.

[4]Smith attributes this to Philip Converse, Herbert Asher, and David Dreyer, but the finding is not mentioned in any of the cited work by these authors. I speculate that this was an unpublished rule of thumb passed between colleagues.

A subtler link emerges from Luskin and Bullock's (2011) finding that encouraging DK responses hides partial knowledge, but discouraging them does not. Discouraging DK, akin to lowering the threshold, should only induce a very low-certainty group of respondents to answer. Encouraging DK, akin to raising the threshold, should induce respondents with relatively higher (if still modest) levels of certainty to say DK. In light of the strong relationship between certainty and accuracy on political knowledge questions (Chapter 4), this is precisely what the threshold model predicts.

The threshold model can also help link some of the findings that will emerge below to existing research on DK responses. Although both Study 1 and Study 2 find a clear "we don't know" means "they're not sure" relationship, two differences emerge: the relationship in Study 1 is very steep at low levels of DK, and "bottoms out" at a lower point on the certainty scale (compare Figures 3.2 and 3.3). Based on data that will be analyzed further in the coming chapters, Appendix A uses simulation to show that that this is precisely what one should expect if relative to online surveys, interviewer-administered surveys set a lower certainty threshold. When respondents who are very uncertain are more likely to provide an answer, each respondent who is nevertheless uncertain enough to say DK should predict a larger drop-off in average certainty among those who offer an opinion.

The threshold model can link these findings to other work that documents variation in DK rates across survey modes. In a comparison between questions asked online and over the phone, Atkeson and Adams (2018) found few differences between DK rates on most questions, with the exception of one battery that produced especially high DK rates in the telephone sample. In this high-DK battery, the rate of DK responses shot up an additional 18 percent with the move online. They suspect that these "are likely questions for which respondent uncertainty was very high" (18). A similar finding can be gleaned from the opinion filter experiment conducted by Bishop et al. (1980): the higher the DK rate without an opinion filter, the more substantive responses are "converted" by the presence of an opinion filter.

The threshold model suggests that the findings in Atkeson and Adams (2018) and Bishop et al. (1980) may have occurred for the same reason that "we don't know" means "they're not sure": a higher percentage DK indicates that the people who answered were less sure as well, and hence more likely to switch to DK in response to an increase in the threshold. To visualize this, imagine that an opinion filter drags the threshold across Figure 3.1. Questions that produce low rates of DK without a filter are likely to be high-certainty, meaning that raising the threshold may not "convert" many substantive responses to DKs (right panel). But questions that produce high rates of DK without an opinion filter are likely to be low-certainty, meaning that a threshold-raising filter "converts" a larger share of the substantive responses to DKs (left panel). Appendix A supports this intuition in

(1) a simulation study and (2) a direct test using the authors' data.

Analytic Strategy

The results combine visual comparisons and summary statistics. Throughout, the independent variable and x-axis is the percentage of DK responses, while the dependent variable and y-axis are average certainty among respondents who offered a substantive response.[5] Though these terms have causal undertones, the relationship is descriptive. At stake is the percentage DK's ability to *predict* question-to-question variation in average certainty among those who do not say DK.

The analysis makes use of three summary statistics. The first two are the slope and intercept from the bivariate regression $\overline{C_q} = \alpha + \beta \overline{DK}_q + \epsilon_q$, where q indexes questions, $\overline{C}_q$ is average certainty among those who answered the question, and $\overline{DK}_q$ is the percentage saying DK. The sign of β is a test of whether a relationship exists at all: if the estimate can be statistically distinguished from zero, it is reasonable to believe a higher percentage DK predicts lower certainty among respondents who answer. The magnitude of β measures the steepness of the relationship. It is scaled to be equal to the predicted difference in certainty between a question with 0 percent DK and 100 percent DK. α estimates average certainty on a question with 0 percent DK. The third statistic is R^2 from the same OLS regression. R^2 quantifies how much of the variation in average certainty can be explained by the percentage of DK responses. The more variation is explained, the more trustworthy the heuristic.

To estimate statistical uncertainty about these summary statistics, I use the block bootstrap, which is commonly used to handle dependence between observations, e.g. in time series data (Bertrand et al. 2004) and conjoint experiments (Hainmueller et al. 2015). Appendix A discusses this procedure in more detail. A key point is that statistical uncertainty is estimated based the sampling procedure, not from the selection of questions. The proper interpretation of the certainty intervals is, "in 95 percent of samples, the certainty interval should contain the true value for this fixed set of questions."

The intuition above — that the "we don't know" means "they're not sure" relationship emerges as a function of variation in individuals' certainty around their personal thresholds, as opposed to person-level tendencies to be certain or say DK — has two observable implications that can be tested within this framework. First, the relationship should exist within subgroups of respondents that might be expected to differ along these lines (e.g., education, interest in politics). Second, the

[5] More specifically, if C is certainty and DK is an indicator variable for saying DK, the x axis is $E[DK]$ and the y-axis is $E[C|DK = 0]$.

relationship should be robust to using only within-person variation in certainty, which fully accounts for any association between certainty, personal characteristics, and DK responding.

Evidence from the 1993-2000 ANES

The first test of whether "we don't know" means "they're not sure" uses data from nine ANES surveys conducted between 1993 and 2000. Across these surveys, 73 questions featured a follow-up question about the respondent's certainty level. Most asked respondents to place themselves or politicians on ideological or policy scales. The certainty scales were some version of "How certain are you of where you stand on this?" or "How certain are you of Bill Clinton's position on this?" The three response options, very certain, pretty certain, and not very certain, are respectively scored as 1, 0.5, and 0. A total of 6,593 respondents answered at least one of these questions in person or over the phone.

To examine the relationship graphically, Figure 3.2 plots the percentage of respondents who said DK (x-axis) against average certainty among those who chose a substantive response (y-axis). The left panel's x-axis is in natural units, while the right panel's x-axis is a logged scale. On the Y-axis, a value of 1 would indicate that everyone chose the highest certainty level; 0, that everyone chose the lowest certainty level; and 0.5, that the average person chose the middle level.

The plot shows that the percentage of respondents choosing DK is a predictor of the certainty levels of other respondents. The relationship is steep at very low levels of DK, then becomes increasingly gradual. At very low levels of DK—around 1 percent—the average respondent chooses either the highest certainty level or the middle certainty level. Above 10 percent DK or so, the average person chooses the middle level or below. This suggests that in interviewer-administered surveys, the first several percentage points worth of DK responses predict a lot of variation in certainty among those who offered an opinion.

The right panel's logged x-axis, which stretches out the small values and condenses the large values, suggests that the relationship closer to linear in the logged percentage of DK responses. To test this more rigorously, Appendix Table A.1 uses OLS to estimate the extent to which average certainty is predicted by the percentage of DK responses and the logged percentage of DK responses. Using both functional forms, there exists a negative, statistically significant relationship between the percentage DK and average certainty. The percentage DK explains about 30 percent of the variation in average certainty, and the logged percentage explains about 52 percent (difference: 22; 95% CI: 16, 27).

Figure 3.2: Percent DK versus average certainty among other respondents, 1993-2000 ANES.

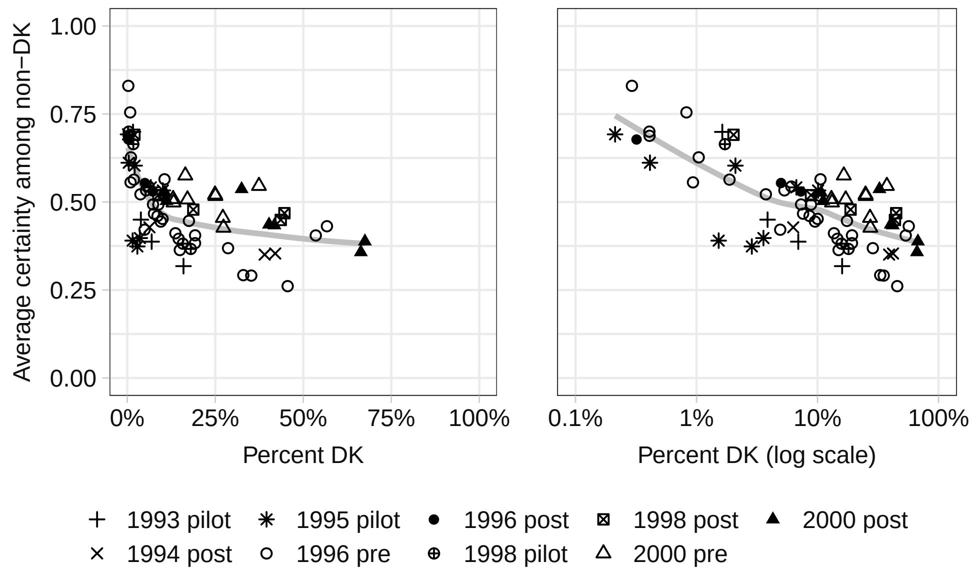

Note: For each ANES question, this figure plots the percentage of respondents who said DK (X-axis) against the average certainty level among people who answered the question (Y-axis). The left panel uses the raw percentage DK while the right panel uses a log scale; the data in the two panels are identical. Appendix A lists each question.

To check the threshold model's implication that the relationship is driven by within-respondent variation in certainty, Appendix A examines both observable implications noted above: that the relationship should be similar within groups that might be expected to vary in their political engagement, and that it should obtain even when only within-person variation in certainty is used. The subgroup analysis finds similar relationships within different levels of political interest and frequency of news consumption. The within-person analysis produces Figure 3.2 and the accompanying regression using only within-respondent variation in certainty. The results are quite similar: R^2 increases by 2.5 percentage points and the slope is hardly perturbed.

The subgroup analysis also supports the other key pillar of the intuition above: widely-shared points of reference in mass media societies. Though certainty is higher and DK rates lower among more-politically interested respondents, these same variables are highly correlated across groups at the question level. For the percentage of respondents saying DK, the correlation is 0.97 between the "not much interested" and the "somewhat interested," 0.98 between the "somewhat interested" and "very much interested," and 0.91 between the "not much interested" and "very much interested." For average certainty among respondents who offered an opinion, these correlations are 0.88, 0.94, and 0.79. Stronger relationships emerge across categories of news consumption. Cross-person differences in certainty are perfectly compatible with shared points of reference.

Evidence from Original Surveys

To examine how the "we don't know" means "they're not sure" relationship would generalize, Study 2 paired a DK option and a certainty scale for a total of 88 total questions. Respondents were recruited online by Lucid, which quota samples to Census demographic margins. The first survey, conducted in July 2019, included 3,662 respondents and had a refusal rate of 2.5 percent. Respondents were randomly assigned to one of two batteries about politically-relevant facts (36 total questions), or to a third battery that was identical to one of the other two, but without a DK option.[6] The second survey, conducted in August and September 2019, included 6,670 respondents and had a refusal rate of 2.6 percent. Respondents were assigned to one of four possible batteries of questions (52 total questions). Respondents to the first survey were excluded from recruitment for the second.

All questions in Study 2 had two response options and an explicit "don't know" option. After each question, respondents who chose an answer were asked, "how sure are you about that?" and presented with a 6-point scale corresponding to a 50, 60, 70, 80, 90, and 100 percent certainty. Respondents who said DK were asked, "what is your best guess?" and presented again with the two substantive response options, with no option to say DK and no certainty follow-up.[7] The question topics represent the following four categories:

- *Favorability ratings* of twelve political figures: President Obama, President Trump, and ten candidates for the 2020 Democratic presidential nomination.
- *Knowledge.* Two batteries of political knowledge questions covering knowledge about public figures (13 questions), civic trivia (11 questions), and statistics about social and economic conditions (12 questions).
- *Policy attitudes.* Two 12-question batteries asked about the respondent's policy attitudes. One battery consisted of the abortion and immigration scales from the Cooperative Congressional Election Survey (CCES), while the other was a grab bag designed to vary in the percentage of DK responses.
- *Policy preferences.* A 16-question battery on the content of Medicare for All and Green New Deal proposals, modelled after contemporaneous polling by the Kaiser Family Foundation.

Though these questions were not a random sample from some larger population of questions, the variety of topics provides some assurance that the findings are broadly applicable.

Figure 3.3 plots the relationship pooling across all questions. Relative to Study 1, the relationship appears tighter and closer to linear in the percentage DK. An OLS regression suggests that

[6] The survey also included two additional arms that were designed to investigate separate research questions.

[7] This strategy balances the costs of saying DK and choosing a substantive response. If one response option had more follow-up questions than the others, respondents could modify their response patterns to avoid answering extra follow-up questions.

Figure 3.3: Percent DK versus average certainty among other respondents, Study 2.

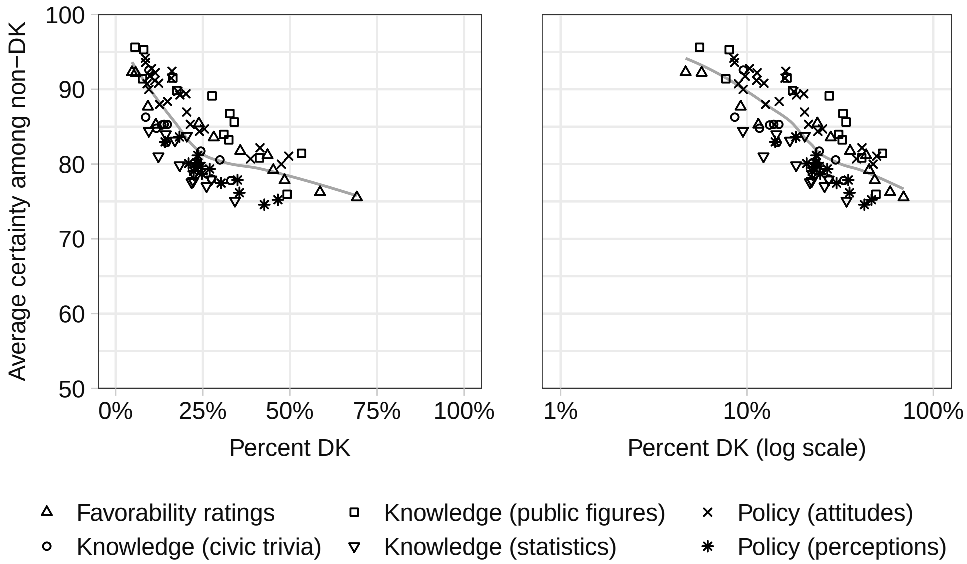

Note: For each question in Study 2, this figure plots the percentage of respondents who said DK (X-axis) against the average certainty level among people who answered the question (Y-axis). To summarize the relationship, each figure features a loess line with block bootstrapped 95 percent certainty intervals; see Appendix A for details.

a question with no DK responses would correspond to an average certainty level of 91.0, whereas a question with only DK responses would correspond to an average certainty level of 61.9 (slope: 29.1; 95% CI: 27.7, 30.6; see Figure 3.6 below). The percentage DK explains 50.0 percent of the variation in average certainty, while the logged percentage explains 59.2 percent (difference: 9.2; 95% CI: 7.5, 10.8). Appendix A shows that these relationships are quite similar across age groups, education levels, gender, income, partisanship, race, and ethnicity, as well as when only within-respondent variation in certainty is used.

Looking within question categories makes the implications more concrete. A particularly clear relationship emerged on the favorability questions, which concerned the two most recent presidents (Obama and Trump) and ten contenders for the 2020 Democratic nomination. Almost everyone offers an opinion about the president, and most who do are fairly sure of their opinion (Figure 3.4, top left). Moving down the smoothed line, one encounters candidates in rough proportion to the public attention they had received by summer 2019. The next-lowest on DK, and next-highest on certainty, is Sen. Bernie Sanders (I-VT), who gained wide public attention as the runner-up for the 2016 Democratic nomination. Obscure candidates like the already-mentioned Bullock and Gravel bring up the rear.

Relative to the full results in Figure 3.3, the relationship on the favorability questions in Figure 3.4 is strikingly tight. In fact, the percentage of respondents saying DK explains more than

Figure 3.4: Favorability questions, Study 2.

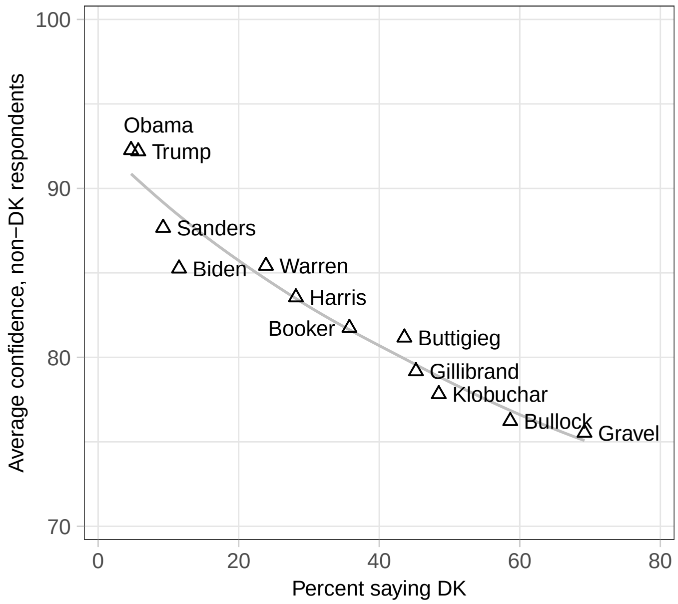

Note: For each favorability rating question in Study 2, this figure presents the same information as Figure 3.3, but with a narrower axis range and labelled data points.

90 percent of the variation in average certainty among those who said "favorable" or "unfavorable" (Appendix Table A.2). This reflects the "we don't know" means "they're not sure" heuristic's greater accuracy in comparisons between similar questions.

Strengthening the Heuristic

For a closer look at the stronger relationship among similar questions, Figure 3.5 separately plots the subcategories of knowledge and policy questions. In the left panel, the top set of questions asks respondents about their policy attitudes (e.g., "Do you support or oppose repealing the Affordable Care Act of 2010?"), while the bottom set asks about the content and consequences of policy proposals (e.g., "Do you think this would happen under a national health care plan, sometimes called Medicare For All? Doctors and hospitals would be paid less."). A fairly tight relationship emerges within both categories. The percentage DK explains about 83 percent of variation in average certainty between the attitude questions and 82 percent between the perception questions (Appendix Table A.2).

Mechanically speaking, the relationship is stronger within each group of questions because of the certainty gap between them. Across the board, respondents are consistently 5 to 10 scale points more certain in their answers to attitude questions. This is equal to 10 to 20 percent of the 50-100 scale, or roughly one-third of the variation in average certainty that can be explained by DK

Figure 3.5: Policy and knowledge questions, Study 2.

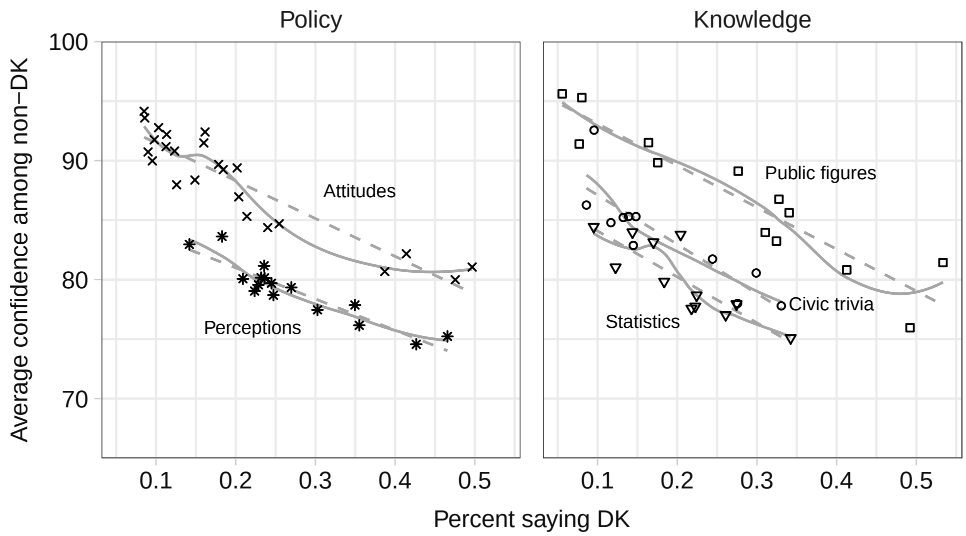

Note: For each policy and knowledge question in Study 2, this figure presents the same information as Figure 3.3, but with a narrower axis range and separate trend lines for each subcategory of question.

responses. When comparing two questions about policy attitudes, one can use the rate of DK to make a fairly precise guess about which question elicited the more certain responses. Yet barring a substantial difference in the DK rate, one should assume that people are more certain in their policy attitudes than in their policy perceptions.

The right panel of Figure 3.5 shows that a similar pattern emerges when the knowledge questions are broken into the three categories listed above: facts about public figures (e.g., the job or political office held by John Roberts), knowledge of political institutions (e.g., the length of a Senate term), and statistics about social and economic conditions (e.g., whether unemployment has been increasing or decreasing). Between any two questions within one of these categories, one can use the DK rate to make a reasonably accurate inference as to the certainty level of respondents who answered the question. Yet across categories, it is important to be cognizant that on average, people tend to be more certain in their answers about public figures than about economic indicators.

To summarize the accuracy gains when examining questions of the same type, Figure 3.6 compares three linear models of the "we don't know" means "they're not sure" heuristic. In the first column, the relationship is assumed to be the same for all types of questions, reproducing the results discussed alongside Figure 3.3. In the second, questions of different types are assumed to differ in their average certainty, but not the slope of the relationship. In the third, average certainty and slope are allowed to vary across question categories.

Figure 3.6: Regression Test for Differences in Intercepts and Slopes, Study 2

	Model		
Term	Bivariate	Variable Intercepts	Variable Intercepts & Slopes
(Intercept)	0.910 (0.906, 0.914)	0.928 (0.922, 0.935)	0.912 (0.906, 0.918)
DK	-0.291 (-0.306, -0.277)	-0.302 (-0.313, -0.291)	-0.250 (-0.269, -0.231)
Knowledge (civic trivia)		-0.036 (-0.045, -0.029)	0.000 (-0.012, 0.011)
Knowledge (public figures)		0.024 (0.016, 0.033)	0.054 (0.044, 0.065)
Knowledge (statistics)		-0.066 (-0.074, -0.058)	-0.034 (-0.047, -0.020)
Policy (attitudes)		0.015 (0.008, 0.022)	0.035 (0.028, 0.042)
Policy (perceptions)		-0.055 (-0.064, -0.046)	-0.050 (-0.063, -0.037)
DK × Knowledge (civic trivia)			-0.161 (-0.220, -0.106)
DK × Knowledge (public figures)			-0.101 (-0.139, -0.065)
DK × Knowledge (statistics)			-0.129 (-0.191, -0.071)
DK × Policy (attitudes)			-0.067 (-0.098, -0.038)
DK × Policy (perceptions)			-0.009 (-0.049, 0.031)
R^2	0.500 (0.463, 0.538)	0.891 (0.874, 0.907)	0.905 (0.889, 0.920)

Note: This table tests for differences between the question categories displayed in Figure 3.3. The first column tests across all question categories, the second column allows for mean differences between categories, and the third column allows the slope to vary between categories (see text).

The results suggest that differences in average certainty between questions of different types are responsible for almost all of the "we don't know" means "they're not sure" heuristic's greater accuracy within sets of similar questions. In the first column of Figure 3.6, R^2 is the same as reported above, 0.500. In the second column, allowing average certainty to vary across question types boosts R^2 to 0.891, an increase of 0.391. In the third column, allowing the slope to vary across categories boosts the explanatory power a bit more, to 0.905, an increase of 0.405. The additional gain in explanatory power, 0.014, is less than 5 percent of the gain from allowing the intercept to vary.

Though allowing for differences in slopes does not add much explanatory power, it does demonstrate a pattern that is too subtle to be visible in the figures: the slope of the relationship between DK responses and others' certainty is slightly steeper for knowledge questions than for policy and favorability questions. Roughly speaking, the interaction terms in the third column suggest that for every 10 percentage points of DK responses, respondents are about 1 percentage point less certain in their answers to knowledge questions than they are in their answers to favorability and policy questions. Though this difference is likely too small to worry about in practical applications, it may contain lessons for future work on estimating unmeasured certainty.

Why DK Options Leave Uncertainty Behind

Cognizance that different types of questions vary in their average certainty level allows one to apply the "we don't know" means "they're not sure" heuristic with greater accuracy. But if a DK option is offered, why would some questions have greater average certainty than others? Shouldn't respondents who are unsure say DK?

A fundamental flaw in this premise — that one either has an attitude or belief, or does not — is that different questions ask respondents to make different types of inferences. When deciding whether John Roberts is more likely to be the Chief Justice of the Supreme Court or the Secretary of Defense, most people either know or have no idea. This produces a not-quite-bimodal, "know it or don't" certainty distribution. By contrast, when deciding whether unemployment is more likely to have gone up or down over the past year, one who does not know the exact figures may still have a strong basis to guess.[8] This produces a more spread out distribution, with lower certainty in the average answer.

To make this more concrete, Figure 3.7 plots the distribution for eight questions: four about the country led by a foreign leader and four about social and economic statistics. The foreign leader

[8] For further analysis of informed guessing's role in responses to such questions, see Graham (2020).

Figure 3.7: Example of question type differences in certainty distributions.

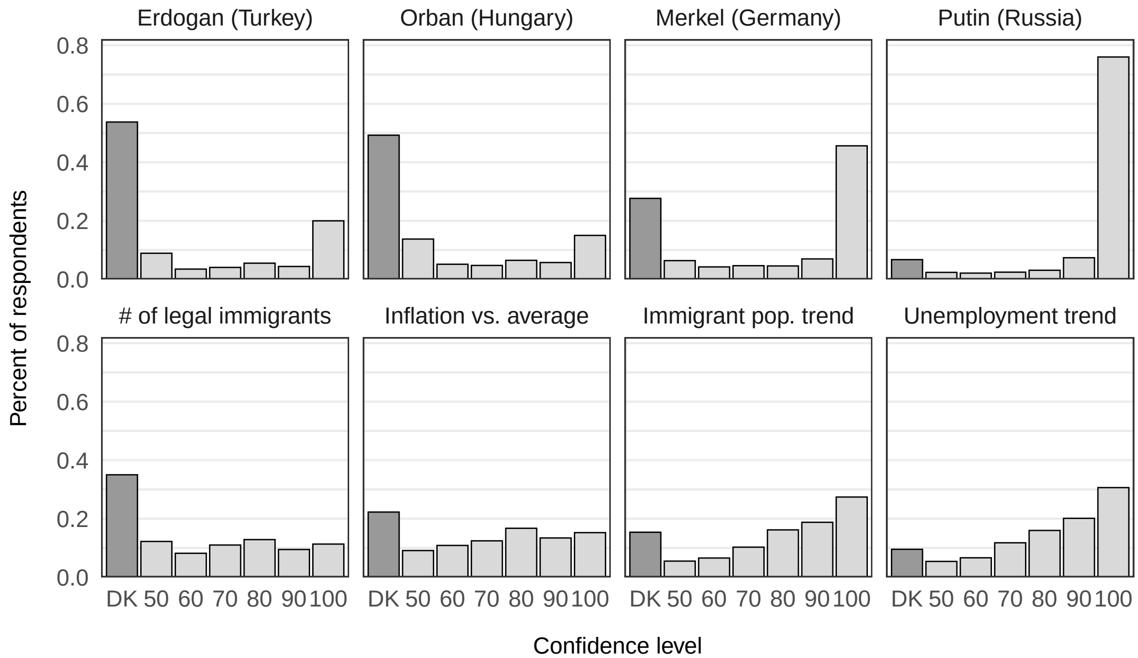

Note: For eight selected knowledge questions in Study 2, this figure plots the distribution of respondent content. The top row's questions all come from the public figures category, and the bottom row's all come from the social conditions category. Darker grey bars are the percentage of respondents saying DK, while the lighter grey bars are the certainty distribution for those who answered the question. The key pattern is that the certainty distributions look similar within each row, but that the two rows look systematically different from one another.

questions have fairly bipolar, "know-it-or-don't" certainty distributions. The social statistics have more mass in the middle, representing the greater ease of producing educated guesses on these questions. Within both categories, the percentage DK (dark grey bar) predicts of average certainty among respondents who answered (light grey bars). For example, a question about the number of immigrants who obtain legal permanent residence each year in the U.S. produces a higher percentage DK than, and lower average certainty than, a question about whether the percentage of Americans who are immigrants has been increasing or decreasing. Yet the questions about foreign leaders have higher average certainty at all levels of DK responding. For example, although the question about which country Angela Merkel leads produces a higher rate of DK responses than a question about the rate of inflation, answers to the Merkel question were chosen with greater certainty than answers to the inflation question.

The variation in certainty levels that exists within and across questions is too rich to be eliminated by a blunt instrument like DK responses, which divides respondents into only two certainty levels, "don't know enough to make a guess" and "know enough to make a guess." To provide a clearer sense of this, one of the two factual batteries featured a split ballot experiment in which a

Figure 3.8: Certainty distribution with and without a DK response option.

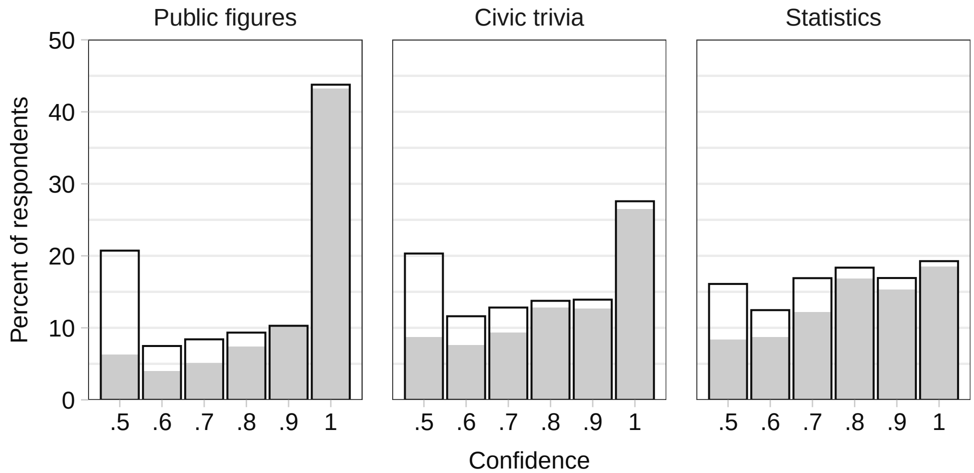

Note: For all of the knowledge questions included in form 1 of study 2, this figure plots the distribution of respondent certainty. The hollow, black-outlined bars represent the distribution when a DK responses option is not offered. The solid, grey bars represent the certainty distribution when DK is offered. The difference in height between the two sets of bars represents the certainty levels of the responses that convert to DKs when a DK option is offered.

randomly-selected set of respondents answered identical questions, but without a DK option. Figure 3.8 visualizes the differences between the two randomly-selected groups of respondents. The hollow bars with black outlines represent the certainty distribution when DK responses are not offered, while the grey bars represent the certainty distribution when DK is allowed. At low certainty levels, the grey bars "fill up" the hollow bars halfway or less, indicating that most — but not all — respondents who are completely uncertain choose DK when that option is made available. Moving across the x-axis to higher certainty levels, the grey bars increasingly fill up the hollow bars. The more certain the respondent, the less likely the respondent is to say DK when a DK option is offered.

Though DK responses work roughly as they are supposed to, removing some low-certainty respondents does not remove the substantial differences that exist across questions. On the questions about public figures, more than 40 percent of respondents claim to be completely certain of their answer, regardless of whether a DK option is offered. On the questions about statistics, which are much friendlier to guessing, the rate of certain responses is less than 20 percent.

The split ballot experiment can be used to compute a more specific estimate of how certain DK respondents would have been if they had answered the question. Let $\mathbb{E}[C \mid \cdot]$ denote the conditional average certainty level among some group of respondents. Although $\mathbb{E}[C \mid \text{said DK}]$ is not observable in the data, it can be estimated. Appendix A shows that the average DK respondent's certainty level can be expressed as

$$\frac{\mathbb{E}[C \mid \text{DK not allowed}] - \big(\mathbb{E}[C \mid \text{DK allowed}] \times (1 - \Pr[\text{said DK} \mid \text{DK allowed}])\big)}{\Pr[\text{said DK} \mid \text{DK allowed}]}, \tag{3.1}$$

all components of which are observable. Pooling across all questions, this estimator suggests that the average respondent who says DK would have stated a certainty level of 59.9 if the DK option had not been provided (95% CI: 55.9, 63.6).

Together, this estimate and Figure 3.8 provide a clearer sense of why the relationship between DK and average certainty has a tendency to "bottom out" at 70 to 75 percent certainty (Figures 3.3, 3.4, and 3.5). DK response options absorb most of what would otherwise have been the lowest-certainty responses, along with as a smattering of responses that would have been stated with low certainty but not a complete lack of it. Consequently, while DK response options cut down the rate of totally-uncertain "coin flips" by half or two-thirds, they do not do much to eliminate guesses that are offered with even slightly higher levels of certainty.

Summary

Three takeaways emerge from Study 2. First, providing a DK option does not even out the substantial variation in certainty that exists from question to question. Second, though respondents are not usually asked about their certainty in their answers, one can use the percentage to DK make informed comparisons between questions. Third, when one looks within the same family of questions, the percentage DK goes from being a helpful approximation to a remarkably accurate predictor.

Implications

This chapter developed an account of what DK response options do, then used that account to demonstrate the "we don't know" means "they're not sure" heuristic. The heuristic itself is most useful when certainty is not measured: even when DK responses are the only information about uncertainty, one can learn something about peoples' confidence *in their answers*. This serves as motivation for the remaining chapters: there is meaningful variation in belief strength that is unlikely to be filtered out by a DK response option.

For the measurement of misperceptions and partisan belief differences, three key implications emerge from this investigation. First, DK response options do something useful: they eliminate a subset of responses that respondents otherwise would have characterized, by and large, as complete guesses. This isolates a set of responses that is more likely to be based on at least some concrete thought, closer to the middle of Tourangeau et al. (2000)'s continuum. For researchers who simply want a reflection of respondents' factual perceptions, DK responses can do a lot of good.

Second, the idea of a latent certainty threshold provides a reasonable model for what DK response options do. Threshold-based conceptions of belief are not just reflected in the thinking that researchers use to justify DK response options — they are also evident in what DK responses actually do. The threshold model is a good model for how many researchers think about beliefs and what the measurement techniques deployed in that spirit actually do.

Third, because the certainty threshold for answering questions tends to be fairly low, DK response options come nowhere close to isolating respondents who "believe" their answers in the sense that they are highly certain about it. This means that DK response options are a poor tool for identifying misinformed respondents as defined by Kuklinski et al. (1998, 2000) and Jerit and Zhao (2020), or those who have misperceptions as defined by Flynn, Nyhan and Reifler (2017) and Pasek et al. (2015). The opposite of "completely uncertain" is not "completely certain." Instead, there is a wide range of variation in certainty levels. DK response options are pretty good at what they do, but cannot be expected to do everything.

Fourth, and by extension, identifying respondents who know the answer to a factual question, or completely believe their choice of an incorrect response, requires researchers to do something other than provide a DK option. The natural choice, given the discussion so far, is to measure respondents' certainty about their answers. Yet to date, research on political beliefs that makes use of such measures completely takes their measurement properties for granted; no systematic evidence exists as to whether a claim to be certain about a survey response actually reflects its face-value meaning. The next two chapters take up this question.

Chapter 4

Certainty, Knowledge, and Ignorance

Abstract. This chapter introduces a framework for evaluating the extent to which certainty scales measure genuine variation in respondents' probabilistic belief in their answer. Using political knowledge items that have traditionally been used to measure general political awareness, the chapter applies this strategy to the question of whether certainty scales can capture meaningful variation between the knowledgeable and the ignorant. The chapter finds that respondents who claim to know the answer consistently behave as if they have a correct belief: in follow-up measures, they consistently choose the same correct answer, express a high degree of certainty about it, and reveal highly certain beliefs through costly choices. This demonstrates that certainty scales can capture meaningful variation in respondents' probabilistic beliefs about their answers, and that the extent to which any particular question succeeds in doing this can be objectively assessed through measurement.

We learned in the previous two chapters that survey responses can be characterized as probabilistic beliefs, which vary along a continuum from complete belief in one response option to "no opinion whatever" (Tourangeau et al. 2000, 12). The standard approach to filtering out less-meaningful responses, "don't know" (DK) response options, targets the lower end of this continuum while leaving the variation in the middle and the high end largely untouched. Consequently, to identify respondents who think they *do* know the answer, something more is needed.

This chapter introduces a characterization and a measure for this continuum as it relates to factual questions: a continuum from complete uncertainty to complete certainty. On a two-option question, a respondent who is totally uncertain about their answer assigns probability of 0.5 to it, the equivalent of a coin flip. A respondent who thinks they know the answer assigns a probability of 1 to it, equivalent to an inference made with complete certainty on the basis of an extremely precise signal: pre-existing knowledge of the truth or "incorrect knowledge" (Hochschild and Einstein 2015,

11) of a falsehood.

Here and throughout the other chapters, *knowledge* is defined as a 100 percent certain belief in something that is true, while complete *ignorance* is defined as assigning equal probability to all of the available response options. Relative to the philosophical literature, this is weak definition of knowledge. There, discussions of knowledge feature another minimum standard: one's "true belief" must also be justified, in that the believer came to hold the belief through a valid process of learning or observation rather than by flukes or mistakes that happened to result in a true belief (Ichikawa and Steup 2018). Following Gettier (1963), many philosophers have argued that even the justified true belief standard is insufficient.

For the present inquiry, the relatively relaxed, "true belief" definition of knowledge has two advantages. First, it is empirically observable, and is consistent with the operational conception of knowledge embedded in most survey research. Whereas many have worried about how blind guessing affects the apparent level of political knowledge, researchers rarely attempt to learn about the process through which "knowledgeable" respondents acquired their belief in true information. Second, philosophers' examples of cases in which true beliefs can be unjustified are highly contrived and of questionable real-world political relevance. For example, a canonical case of an unjustified true belief is a person who correctly believes that there is a barn in a particular county, even though the person only actually saw a painted facade of the barn instead of the real barn that is located elsewhere in the county (Goldman 1976).

To examine the degree to which certainty scales capture meaningful variation in respondents' degree of knowledge or ignorance about the fact in question, I begin in this chapter by examining political awareness questions, which I define as the trivia-style questions about institutions and political figures that have traditionally been used to measure general political knowledge (Delli Carpini and Keeter 1993, 1996). Do respondents who claim to be absolutely certain that John Roberts is the Chief Justice of the Supreme Court make that claim because they really think they know the answer? And are respondents who claim to be totally uncertain about Roberts' job truly revealing a total lack of knowledge?

The evaluation strategy is predicated on the notion that one who knows or believes something should behave as if they know or believe it. Three yardsticks will measure this. First, the temporal stability of respondents' belief statements. This is predicated on the notion that one who genuinely believes something should consistently re-state the same belief. For a threshold thinker, the interpretation is that one who really believes something should always state that belief again when asked to do so again (Converse 1964); for the probabilistic thinker, the interpretation is that someone with meaningful probabilistic beliefs should make the same inference, on average, when asked to do so many times (Gilens 2012). Second, the extent to which stated beliefs align with beliefs that are revealed through costly choices. Here, the logic the same, with the amendment that beliefs and inferences should also remain consistent when they are captured using a distinct measurement technology. Third, the relationship between certainty and accuracy. This supports the interpretation that variation in peoples' degree of belief in their answers to knowledge questions is capturing variation not only in how strongly they believe their answer, but how well they know it.

Although questions that tap general political awareness are not typically the focus of work on partisan belief differences or misperceptions, two of their properties make them a nice initial case for examining the measurement properties of certainty scales. First, one should expect that a substantial number of people actually know the facts in question. By contrast, we will see in later chapters that few respondents actually know the answers to questions about relatively obscure economic statistics like the inflation rate. Beginning with test cases where some base of knowledge is reasonable to assume provides assurance that if certainty scales can identify knowledge in any situation, we should be able to identify it here.

Second, there is no reason to expect false claims about these topics to be pervasive in the public sphere. In this chapter, this opens up the relationship between certainty and accuracy as a metric for the performance of the certainty scale. If all of the available information about the facts can be assumed to be correct, it follows that the feeling of knowing the fact should only emerge from knowledge. By contrast, when misinformation is expected, a weak relationship between certainty and accuracy could occur even if the scale is doing its job.

Approach

To examine how well certainty scales separate knowledge from ignorance on political awareness questions, this chapter examines five surveys that included such questions. All of the surveys were conducted between 2017 and 2020; Appendix E contains more information on each survey.

The questions are a mix of questions recommended by classic work on measuring political knowledge (Delli Carpini and Keeter 1993, 1996), questions that appear in the American National Election Studies (ANES), and some additional topics chosen in the same spirit but are designed to assure variation across the full range of the certainty scale. Roughly speaking, these questions can be grouped into the following categories:

- **Institutional rules.** The length of a Senate term (six years), which branch of government rules on the constitutionality of laws (judicial), the threshold for a presidential veto (2/3), the name of the first ten amendments to the Constitution (the Bill of Rights), the procedure that allows the Senate to make budget changes with a simple majority vote (the Senate) and which chamber of Congress must introduce revenue bills (the House).

- **Public figures.** The job or political office held by John Roberts (Chief Justice of the Supreme Court), Mike Pence (Vice President), Kevin McCarthy (House Majority Leader in 2017), Jerome Powell (Chairman of the Federal Reserve), William Barr (Attorney General), and Charles Schumer (Senate Minority Leader).
- **Foreign leaders.** The country led by Angela Merkel (Germany), Vladimir Putin (Russia), Viktor Orban (Hungary), and Recep Erdogan (Turkey).
- **Political parties.** Which party controls the House of Representatives, which party is known as being more conservative.

Almost all of the questions proceeded in the following format.[1] To begin, respondents were presented with the question and the answer choices, just as they would be in an ordinary survey. The moment the respondent clicked their answer choice, a certainty scale appeared below the question, asking "how certain are you that your answer is correct?," "how likely is your answer to be true?," "how sure are you about that?," or some variant; Appendix E lists all of the wordings. Graham (2020) shows that this approach to measuring certainty does not systematically change the distribution of best guesses: regardless of whether a certainty scale is offered, the percentage of correct answers is about the same. This means that certainty can be measured in a manner that preserves comparability with surveys that do not measure certainty.

Though all of the surveys took this basic approach, the four surveys varied somewhat in the format of their questions and the characteristics of their certainty scales. The September 2017 Lucid survey used a five-point subjective scale similar to that used in most extant misinformation research (e.g., Kuklinski et al. 2000; Pasek et al. 2015). Each of these questions had three response options. The other four surveys used numerical scales ranging from 50 to 100, with verbal labels to assist subjects. Because each question had two response options, these measures are interpretable as the probability the respondent assigns to their correct answer. Two of the strategies, stability and revealed belief, require a second measure of belief. To measure stability, respondents were recontacted and asked identical questions weeks or months after the initial survey. Revealed belief was always measured in the same survey, and always after the stated measure of beliefs.

The validity of all of these metrics could be disrupted in some way by respondents who look up the answer: such respondents ought to be certain and correct, stable in their responses, and willing to bet that their answer is correct not because they really know the answer, but because they looked

[1]The exception is the October 2019 survey, which included a split-ballot experiment described below in the section on all-in-one versus branching scales.

up the information during a survey. To mitigate this concern, all surveys conducted in October 2019 or later included snippets of Javascript that monitored whether the survey remained visible on the respondent's screen (see Appendix E.3). All cases of possible information search are dropped from the analysis. The remaining surveys also took steps to discourage and detect information search. Prior to the beginning of the factual question, every survey included a pledge not to look up the answers. Some of surveys also included "catch" questions that respondents would be unlikely to answer correctly without looking up the answer. Appendix E lists the exact safeguards for each survey.

Accuracy

As the first test of whether certainty scales can distinguish the knowledgeable from the ignorant, this section examines the relationship between certainty and accuracy. To the extent that respondents who claim to be certain truly know the fact, such respondents should be 100 percent accurate. To the extent that respondents indicate complete uncertainty, they should be no more accurate than chance.

As a first look at the degree to which certainty predicts accuracy, Figure 4.1 plots this relationship for each of five surveys that included these questions. The x-axis plots the respondent's stated certainty level, while the y-axis plots the percentage of respondents who answered the question correctly for each of these groups. Running through each figure are two dashed lines: a horizontal line indicating the percentage of answers that would be correct completely by chance, and a 45-degree line indicating perfect correspondence between certainty and accuracy.

On the September 2017 survey, about 36.7 percent of respondents stating the minimum level of certainty provided the correct answer (95% CI: 34.3, 39.1). This is barely more accurate than chance: completely random responses to a three-option question would be expected to be correct one-third of the time. Meanwhile, respondents who claimed to be absolutely certain were correct 85 percent of the time. Despite the face implausibility that many respondents would be exposed to outright false information about such facts, many respondents appeared to fall into "traps" set by the response options. The most common claims to be certain and incorrect pegged Nancy Pelosi as the Senate Minority Leader (instead of Chuck Schumer) and the filibuster as the Senate procedure to make budget changes via a simple majority (instead of reconciliation).

The results were similar for the other four surveys, which paired two-option questions with numerical certainty scales. In all four surveys, respondents who stated a probability of 0.5 chose the

Figure 4.1: Probability of correct answer by certainty level, political awareness questions.

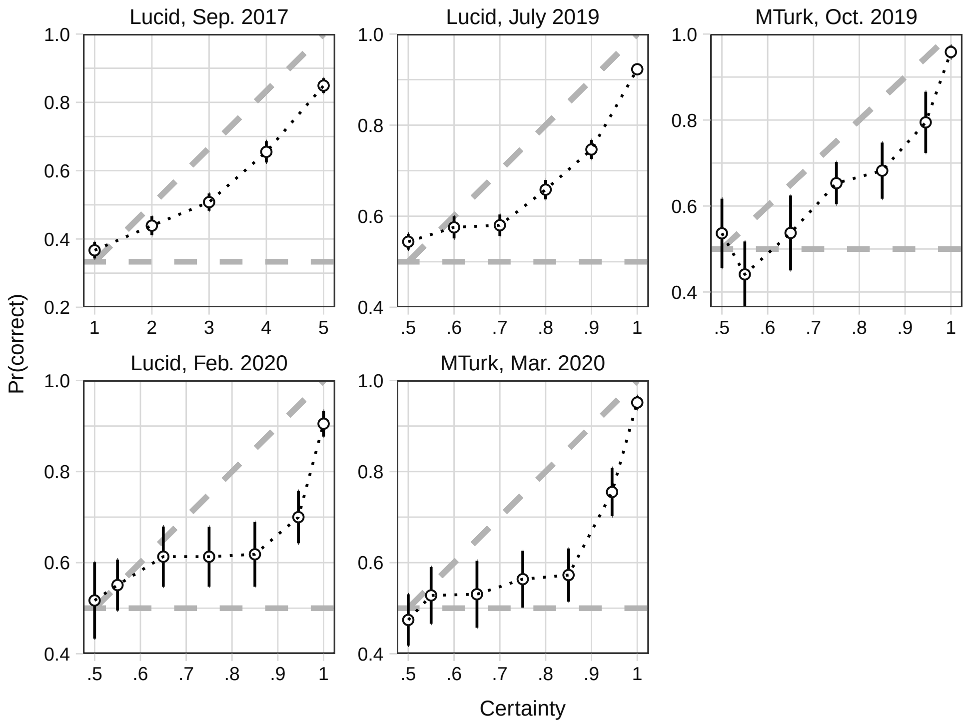

Note: For each of the five surveys that included political awareness questions, this figure displays the probability of answering correctly (y-axis) conditional on the respondent's stated certainty level (x-axis). Dots are point estimates; vertical bars, 95% confidence intervals. The dashed grey line indicates the relationship that would realize under perfect calibration, as defined in the text. The September 2017 and July 2019 surveys used discrete certainty scales; respondents could only choose the values displayed on the x-axis. The other surveys used quasi-continuous scales that allowed respondents to select any integer between 50 and 100. Here and in all subsequent figures, these measures are binned using the following groups: 0.5; [0.51, 0.59]; [0.6, 0.69]; [0.7, 0.79]; [0.8, 0.89]; [0.9, 0.99]; 1. Appendix B presents all of these estimates in tabular form.

correct answer between 49 and 54 percent of the time, while respondents who stated a probability of 1 chose the correct answer between 89 and 96 percent of the time. MTurk respondents seem to report their certainty level slightly more reliably than Lucid respondents: while the October 2019 and March 2020 surveys saw respondents who claimed to be absolutely certain answer correctly 96 and 95 percent of the time, respectively, equally certain respondents in the July 2019 and February 2020 surveys were correct 92 and 89 percent of the time.

Across the five surveys, there was some variation in the characteristics of the scales: the three conducted in 2017 and 2019 had labels on middle scale points, while the surveys conducted in 2020 did not. Figure 4.1 suggests that this design feature made a difference. In the top row, the surveys that had labelled middle scale points remain fairly close to the 45-degree line indicating perfect correspondence between certainty and accuracy. By comparison, the bottom row shows accuracy dip more substantially below the 45-degree line at the middle points on the x-axis, before rising

Table 4.1: Calibration of certainty-accuracy relationship.

Survey	Calibration
Lucid, Sep. 2017	0.884
Lucid, July 2019	0.913
MTurk, Oct. 2019	0.916
Lucid, Feb. 2020	0.894
MTurk, Mar. 2020	0.895

Note: This table displays estimates of (4.1).

again near the right side.

To summarize the degree to which certainty matches up with accuracy, I first turn to *calibration*. Researchers apply calibration statistics in a wide range of contexts, from knowledge questions and facial recognition to clinical assessments and weather forecasts (Lichtenstein and Fischhoff 1977; Koriat et al. 1980; Weber and Brewer 2003; Luna and Martín-Luengo 2012). I measure calibration using the absolute value method first proposed by Oskamp (1962), with the scale flipped so that a score of 1 indicates perfect calibration:

$$1 - \frac{1}{N}\sum_j N_j |j^* - \Pr(\text{corr}|j)| \tag{4.1}$$

where j^* is the numerical probability or proximity score that corresponds to certainty level j, $\Pr(\text{corr}|j)$ is the percentage of correct answers in level j, and N_j is the number of respondents naming level j. Intuitively, the statistic compares each dot in Figure 4.1 to the 45-degree line, then takes a weighted average.

Calibration between certainty and accuracy was high on each survey. Table 4.1 shows the calculation of expression (4.1) for each of the five surveys. The lowest calibration was 0.884 on the September 2017 Lucid survey. Among the four surveys that used probabilistic scales, calibration was slightly higher on the two surveys that labelled the middle scale points (0.913 and 0.916) than on the surveys that did not provide such labels (0.893 and 0.894).

For many readers, OLS regression is likely to be a more familiar approach to summarizing relationships like those displayed in Figure 4.1. To summarize the overall relationship between certainty and accuracy on each survey, I used OLS to estimate

$$\text{Correct}_{ij} = \alpha + \beta c_{ij} + \epsilon_{ij}, \tag{4.2}$$

where i indexes respondents, j indexes questions, Correct_{ij} is an indicator variable for choosing the correct answer, and c_{ij} is the respondent's certainty level. c is scaled to range $[1/k, 1]$, where k is the number of response options on the question. For the four surveys with numerical scales, this simply converts the 50-100 scale to a [0.5, 1] probability scale; for the September 2017 Lucid survey, this amounts to a transformation that makes the regression coefficients easier to compare.

Relative to the calibration statistic, an advantage of OLS is the ability to easily control for other sources of variation in the relationship. A crucial test for a measure of respondents' certainty levels is within-respondent predictive power. If the certainty scale captures meaningful variation in the respondent's personal feeling of knowing, certainty should predict accuracy even when controlling for each person's average level of certainty. Even if this were not the case, a relationship between certainty and accuracy could exist: it could be that generally more-confident people also tend to have greater political knowledge, even if people have little sense of which facts they themselves know or do not know.

To estimate the slope based on within-respondent variation in certainty, I used OLS to estimate the parameters in

$$\text{Correct}_{ij} \sim \sum_{i=1}^{N} \alpha_i + \beta c_{ij} + \epsilon_{ij} \tag{4.3}$$

where α_i is a respondent fixed effect and all other terms are as defined above. For good measure, I also present between-respondent slope. This is estimated by

$$\overline{\text{correct}}_i \sim \alpha + \beta \overline{c}_i + \epsilon_i \tag{4.4}$$

where the vertical bars indicate each respondent's personal average across all of the questions they answered.

The regression results confirm that certainty strongly predicts accuracy (Table 4.2). Across the five surveys, the overall slope of the relationship varies from 0.61 to 0.88, implying that a one unit increase in certainty implies a large increase in accuracy. The predicted values that emerge from combining the slopes and intercepts imply relationships that are fairly similar to the means plotted above in Figure 4.1. On the Lucid survey, a respondent who was completely uncertain between the three options is predicted to be correct 32.5 percent of the time, while one who was completely certain would be correct 80.5 percent of the time — both slight under-estimates. The estimates for the other four surveys, which use two-option questions, imply that totally uncertain respondents would be correct 44.1, 49.4, 44.1, and 38.9 percent of the time, while respondents who

Table 4.2: Regression results: certainty and accuracy.

(a) Lucid, September 2017

Term	Overall	Within	Between
α Constant	0.085^{**}		-0.001
	(0.017)		(0.026)
β Certainty	0.720^{**}	0.610^{**}	0.850^{**}
	(0.024)	(0.030)	(0.039)
Adj. R^2	0.123	0.174	0.340
Num. obs.	7176	7176	898

(b) MTurk, October 2019

Term	Overall	Within	Between
α Constant	-0.001		0.008
	(0.046)		(0.064)
β Certainty	0.842^{**}	0.853^{**}	0.830^{**}
	(0.053)	(0.071)	(0.076)
Adj. R^2	0.099	0.133	0.149
Num. obs.	2522	2522	631

(c) Lucid, July 2019

Term	Overall	Within	Between
α Constant	0.106^{**}		0.074^{**}
	(0.015)		(0.023)
β Certainty	0.776^{**}	0.740^{**}	0.811^{**}
	(0.017)	(0.022)	(0.027)
Adj. R^2	0.117	0.156	0.268
Num. obs.	19267	19267	2432

(d) Lucid, Feb. 2020

Term	Overall	Within	Between
α Constant	0.137^{**}		0.159^{*}
	(0.044)		(0.065)
β Certainty	0.607^{**}	0.552^{**}	0.573^{**}
	(0.055)	(0.076)	(0.083)
Adj. R^2	0.052	0.084	0.087
Num. obs.	2399	2399	523

(e) MTurk, March 2020

Term	Overall	Within	Between
α Constant	-0.063		-0.105
	(0.037)		(0.060)
β Certainty	0.883^{**}	0.791^{**}	0.929^{**}
	(0.043)	(0.057)	(0.072)
Adj. R^2	0.118	0.185	0.165
Num. obs.	3417	3417	926

Note: This table displays OLS estimates of the parameters in (4.2), (4.3), and (4.4). Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

were completely certain would be correct 84.1, 88.2, 74.4, and 82.0 percent of the time — again, all under-estimates.

The slight-yet-systematic inaccuracy of these regression-based predictions highlights a hazard of leaning heavily on predicted values from regressions, which is a standard practice in social science research; often, researchers show such values *instead of*, rather than in addition to, showing the data. In this case, the predicted values are too low because measurement error is concentrated at the middle scale points. OLS imposes a flat line on the data, which predicts values that are too low at the endpoints and too high in the middle. Analysis that examines the certainty-accuracy relationship based only on predicted values from linear regressions would under-state the degree to which respondents who claim to be perfectly certain and uncertain are aware of their ignorance, while also obscuring evidence of superior performance of scales with verbal labels at the midpoint.

The within-person estimates support the conclusion that certainty scales capture within-person variation in the feeling of knowing. These estimates are presented in second column in Table 4.2 plots the within-person estimates. The coefficients are similar in magnitude to, but a little smaller than, those estimated to describe the overall relationship.

Across all of these metrics, the strong relationship between certainty and accuracy suggests that certainty scales capture meaningful variation in the respondent's level of information. At least when it comes to general political awareness questions, a claim to be certain is a pretty good indication that one actually knows the answer. However, because the expectation that accuracy should correspond depends on an assumption that respondents were not exposed to false information, this strategy cannot be used to evaluate certainty scales on questions where misinformation or misperceptions are expected. The next two sections introduce two strategies that tap certainty more directly.

Response Stability

The idea that survey respondents who really believe something should choose the same response in a second survey has a long history in public opinion research. A completely crystallized attitude or belief should be completely stable, while a response that does not at all represent an attitude or belief should be no more stable than chance. This section uses response stability as a measure of whether the respondent makes the same inference when asked to do so a second time. To the degree they do, single-shot measures of probabilistic certainty can be treated as good proxies for the latent, personal probabilities that respondents assign to each choice.

Two measures of stability are used. First, *response stability*, defined as the percentage of respondents who choose the same answer when asked the same question a second time. Second, *belief stability*, defined as how strongly one believes their initial response. Formally, let $a_{it} \in \{0, 1\}$ be respondent i's answer choice at time t, and $c_{it} \in [0.5, 1]$ be the probabilistic certainty level that they assigned to this response at time t. In the initial measurement taken at $t = 1$, the respondent's certainty level is equivalent to the probability they assign to their answer: $c_i = p_{i1}^{\text{initial}}$. In the second measurement, taken at $t = 2$, the probability the respondent assigns to their initial $t = 1$ response, p_{i2}^{initial}, is

$$\begin{aligned} c_{i2} \quad & \text{if} \quad a_{i1} = a_{i2} \\ 1 - c_{i2} \quad & \text{if} \quad a_{i1} \neq a_{i2}. \end{aligned} \tag{4.5}$$

which is equivalent to the probability that the respondent assigned to their initial answer when asked to make the same inference a second time. If beliefs were measured without error, it would always be the case that $c_{i1} = p_{i2}^{\text{initial}}$.

Though response stability may seem simpler and more familiar, belief stability has an important advantage: easily interpreted expectations about good performance. If somebody really thinks that the probability that the unemployment went down is 0.75, an error-free measure would result in the respondent choosing 0.75 as the probability every single time, while an unbiased measure would result in the respondent choosing 0.75 on average. By contrast, response stability is defined in terms of the option the respondent thinks is most likely to be true. If the respondent's "true" probability is 0.75, an error-free measure of the respondent's best guess would find that unemployment decreased every single time — as would be the case if the true probability were 0.6 or 0.9. Consequently, the percentage of the time that a respondent chooses an answer should not be equal to the probability that a respondent assigns to that answer. To reflect this, plots of response stability will omit the 45-degree line that, for the other measures, indicates ideal performance. If these figures featured such a line, it would be a vertical one at 100 percent, indicating perfect stability regardless of how strongly they believed their best guess.

To examine these quantities, I will use two surveys that included both political awareness questions and a panel component. The October 2019 MTurk survey included a second wave in November 2019, while the March 2020 MTurk survey included a second wave in August.

Analysis

As a first look at the relationship between certainty and response stability, Figure 4.2 displays the results separately for both surveys. In each panel, the x-axis displays the certainty level stated in wave 1. In the top row, the y-axis displays the percentage of respondents who chose the same answer again in the second survey. In the bottom row, the y-axis displays the probability assigned to the initial response (i.e., expression 4.5). In both panels, the dashed line at 0.5 indicates the amount of response or belief stability that would be expected from completely random responding.

Certainty was highly predictive of response stability. Among respondents who claimed to be certain of their answer in the first wave, 98 percent of respondents selected the same answer in the second wave. Among those who claimed to be totally uncertain, 61 percent selected the same answer in the second wave of the October-November 2019 panel and 53 percent did so in the second wave of the March-August 2020 panel. Between these extremes, stability follows a steady upward trend, suggesting that middling levels of certainty are also capturing meaningful variation in how

Figure 4.2: Response stability and belief stability by certainty level.

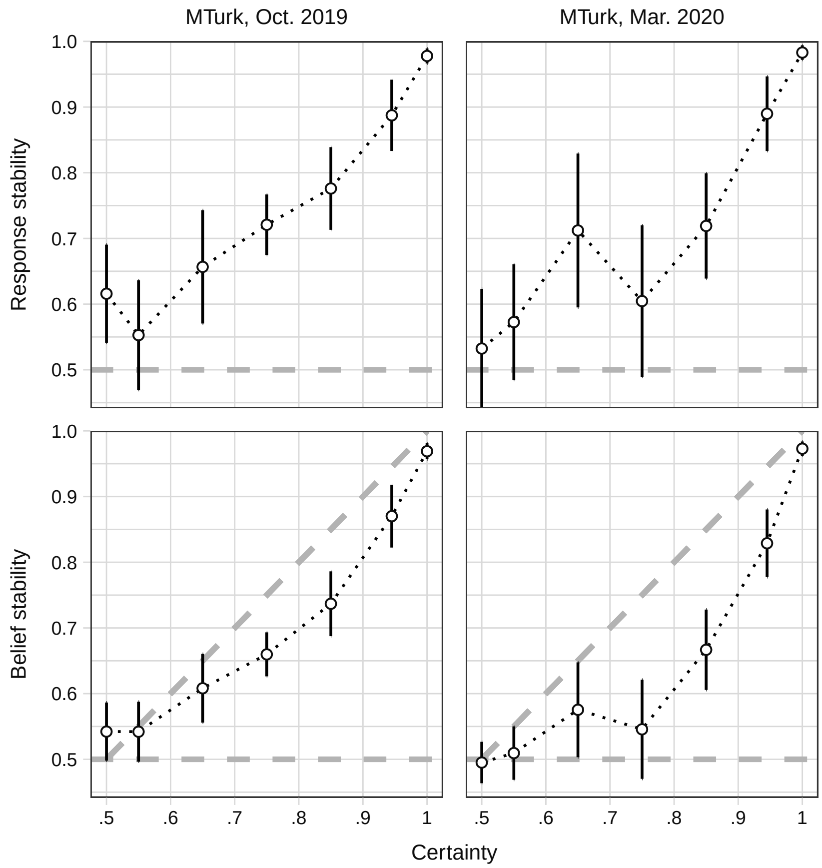

Note: Conditional on the certainty level the respondent stated in wave 1 (x-axis), this figure displays average response stability and belief stability (y-axis). Dots are point estimates; vertical bars, 95% confidence intervals. For belief stability, the dashed grey line indicates the relationship that would realize in the absence of measurement error. Appendix B presents all of these estimates in tabular form.

well respondents really know the fact.

Certainty was also highly predictive of belief stability. Among respondents who claimed to be certain in the first wave, the average respondent assigned a probability of 0.97 to that same response when asked again in the second wave of both surveys. Not only do these respondents seem to know the answer, but they consistently indicate that they are certain of it. On the other end of the spectrum, respondents who claimed to be completely uncertain assigned a probability of 0.54 and 0.50 to their initial response. In the middle, the certainty scale once again appears to capture meaningful variation, with the dip below the 45-degree line indicating that measurement error is relatively concentrated in these levels.

To summarize this relationship, I once again begin with calibration, this time substituting response stability and belief stability for the percentage of correct answers. Using both metrics,

Table 4.3: Calibration: response stability and belief stability.

	Calibration	
Survey	Response stability	Belief stability
MTurk, Oct. 2019	0.964	0.947
MTurk, Mar. 2020	0.953	0.934

Note: This table displays OLS estimates of (4.1), with response stability and belief stability substituted for accuracy. Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

calibration exceeds what was observed when accuracy was the outcome metric. Response stability saw calibration scores of 0.96 and 0.95, while belief stability saw scores of 0.95 and 0.93 (Table 4.3).

These results are particularly encouraging for the March 2020 MTurk survey. Above, we saw that on this survey, accuracy-based calibration appeared worst at middling levels of certainty (Figure 4.1). Using response stability, performance at the middle levels appears much better. This implies that many of the respondents at middling certainty levels were guided by misleading heuristics that they consistently used to inform their answer. On average, such heuristics did not result in greater accuracy, but they did result in greater response stability. The next chapter, which splits stability by whether the wave 1 response was correct or incorrect, verifies this suggestive inference.

For a further test of the relationship between certainty and stability, I turn again to OLS regression. To estimate the overall relationship between certainty and belief stability, I estimate the parameters in

$$p_{i2j}^{\text{initial}} = \alpha + \beta c_{i1j} + \epsilon_{ij}, \tag{4.6}$$

where all terms are as defined above. Values of β closer to 1 indicate greater correspondence, on average, between the beliefs stated in wave 1 and wave 2. I also estimate the within- and between-person relationships by making the same change to the dependent variable in equations (4.3) and (4.4).

The regression results demonstrate a strong correspondence between certainty and response stability (Table 4.4). The slope estimates of 0.88 and 0.93 indicate a near-perfect correspondence between the certainty level stated in wave 1 and the average belief stated in wave 2. For the October-November 2019 panel, the predicted values imply that a totally uncertain respondent would assign a probability of 0.472 to their initial response, while a completely certain respondent would assign probability 0.914 to it. For the March-August panel, these figures are 0.440 and 0.904. As with the certainty-accuracy relationship, the tendency for regression to under-estimate stability among the

Table 4.4: Regression results: certainty and belief stability.

(a) MTurk, October 2019

Term	Overall	Within	Between
α Constant	0.030		0.015
	(0.031)		(0.042)
β Certainty	0.884^{**}	0.867^{**}	0.901^{**}
	(0.036)	(0.045)	(0.051)
Adj. R^2	0.305	0.339	0.418
Num. obs.	1645	1645	412

(b) MTurk, March 2020

Term	Overall	Within	Between
α Constant	-0.024		-0.009
	(0.025)		(0.045)
β Certainty	0.928^{**}	0.937^{**}	0.903^{**}
	(0.029)	(0.036)	(0.055)
Adj. R^2	0.356	0.447	0.348
Num. obs.	1572	1572	418

Note: This table displays OLS estimates of the parameters in (4.2), (4.3), and (4.4), with belief stability substituted for accuracy. Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

most-certain respondents highlights the dangers of assuming that predicted values from a regression faithfully reflect key patterns in the data. Relative to the actual values, 0.969 and 0.972, the predicted values are three times as far from perfect belief stability.

The within-respondent results suggest that the greater stability of more-certain responses reflects meaningful within-person variation in the degree to which responses reflect knowledge or ignorance of the fact. The within-person coefficient estimates are substantively close to, and statistically indistinguishable from, the overall relationship.

Together, these results provide further assurance that respondents' claims to be certain capture a substantial amount of genuine variation in certainty. Even so, response stability has two key shortcomings. First, because response stability takes the same measure twice, one must take for granted that the construct that is stable is, in fact, certainty. To the degree that a measure of certainty also captures systematic variation in some other construct, stability may not be evidence that the scales measure certainty itself. Second, because the surveys are conducted at two different points in time, instability can occur due to genuine changes in beliefs.

Revealing Beliefs through Costly Choices

To address the shortcomings of response stability, I now turn to an alternative strategy for examining the extent to which certainty scales measure certainty. In three of the surveys, all respondents' probabilistic beliefs were measured a second time using a series of costly choices. Because respondents reveal their beliefs through their choices rather than stating them directly, this measure is referred to as the *revealed belief* measure. When drawing contrasts between this measure and the standard, directly-stated measures of certainty used in most political science research, those measures are sometimes referred to as *stated* measures of belief.

The revealed belief measure proceeded as follows. Before beginning, respondents were told that they would make a series of choices about which option gave them a better chance to win either a \$100 bonus payment or a \$100 gift card. Some respondents' choices would be randomly selected to be entered into a drawing for these prizes. Each respondent then completed a three-part training exercise that familiarized them with the structure of the task. Each revealed belief task began by asking respondents whether they would rather win a bonus if one of two statements was true. For example, would you rather win a bonus if John Roberts is the Chief Justice of the Supreme Court or if he is the Secretary of Defense? After making this choice, respondents chose between winning a bonus if their answer was correct and a 6 in 10, 7 in 10, 8 in 10, 9 in 10, and 99 in 100 chance of winning the same bonus. Appendix E.2 includes screen shots of each step of this process.

To see how a respondent should make this choice, suppose that respondent i assigns a probability of 0.75 to Roberts being Chief Justice ($p_i = 0.75$). If offered the choice between winning \$100 if Roberts is Chief Justice and a 7 in 10 chance to win \$100, i should prefer to win a bonus if Roberts is Chief Justice, because the probability offered in the lottery is less than p_i. But if offered an 8 in 10 chance, i should prefer the lottery. The probability which i "crosses over" from preferring payment for a correct answer to preferring to enter the lottery reveals i's probabilistic belief; for this reason, some researchers refer to the revealed belief measure as "crossover payment."

Tasks of this type have an important theoretical advantage: because the reward is held constant, the only difference in the expected payoff is the respondent's self-assessed probability that their answer is correct. This means that the measure is invariant to risk preferences (Ducharme and Donnell 1973; Allen 1987), unlike other methods like the quadratic scoring rule. In empirical tests, Trautmann and van de Kuilen (2015) find that lottery preference tasks outperform the quadratic scoring rule, and Holt and Smith (2016) finds that discrete choice methods like this implementation outperform methods that ask respondents to directly state their crossover probability. Hill (2017) uses this method to study Bayesian updating on politically relevant facts.

Analysis

As a first look at the relationship between certainty and revealed beliefs, Figure 4.3 displays the results by question category and survey. Its interpretation is the same as the interpretation of the bottom row of Figure 4.2, with the exception that revealed belief has been substituted for belief stability. The x-axis displays the respondent's certainty level from the initial measure of belief, while the y-axis displays the belief in the same response that is implied by the respondent's discrete choices.

Figure 4.3: Revealed belief by certainty level.

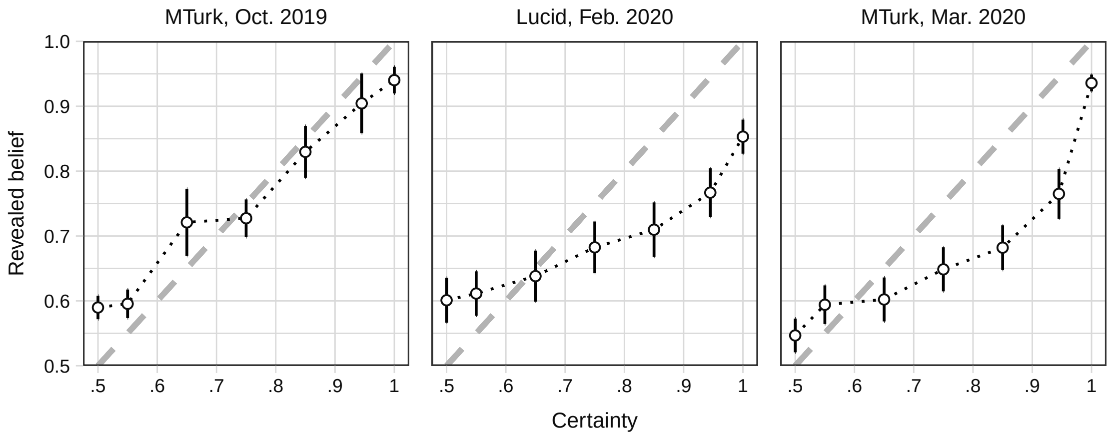

Note: Conditional on the certainty level the respondent stated (x-axis), this figure displays the average probability respondents assigned to their response using the measure of revealed beliefs (y-axis). Dots are point estimates; vertical bars, 95% confidence intervals. The dashed grey line indicates the relationship that would realize in the absence of measurement error. Appendix B presents all of these estimates in tabular form.

In the two MTurk surveys, the revealed belief results are quite similar to the belief stability results. In both surveys, respondents who stated that they were completely certain in the initial measure revealed a belief of 0.94 and in that same response. Meanwhile, respondents who claimed to be completely uncertain revealed a belief of 0.59 and 0.55.

The relationship between certainty and revealed belief was slightly weaker in the February 2020 Lucid survey. Respondents who initially claimed to be certain revealed a belief of 0.85 in their response, while respondents who initially claimed to be completely uncertain revealed a belief of 0.6. This suggests that Lucid respondents had more trouble using the revealed belief measure than did MTurk respondents.

To summarize these relationships, the same calibration statistic as before is adapted to the case of revealed beliefs. On the October 2019 MTurk survey, calibration reaches 0.95, almost identical to the figures for accuracy, response stability, and belief stability. On the March 2020 MTurk survey, the mark of 0.91 was also similar to its performance on the other metrics. The February 2020 Lucid survey's mark of 0.89 is quite similar to the relationship between certainty and accuracy.

To further probe the relationship between certainty and revealed belief, Table 4.6 conducts the same regression tests that were conducted in the previous two sections. These results have the same interpretation as Table 4.4, with the revealed belief measure substituted for belief stability. In each case, the overall and within-person relationships indicate strong correspondence between the two measures: stated belief captures meaningful variation in revealed belief. However, in each case, the coefficients are somewhat attenuated, with larger intercepts and shallower slopes.

Table 4.5: Calibration: revealed belief.

Survey	Calibration
MTurk, Oct. 2019	0.950
Lucid, Feb. 2020	0.892
MTurk, Mar. 2020	0.912

Note: This table displays OLS estimates of (4.1), with revealed belief substituted for accuracy. Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

Table 4.6: Regression results: certainty and revealed belief.

(a) MTurk, October 2019

Term	Overall	Within	Between
α Constant	0.270^{**}		0.241^{**}
	(0.019)		(0.026)
β Certainty	0.618^{**}	0.552^{**}	0.656^{**}
	(0.024)	(0.038)	(0.034)
Adj. R^2	0.315	0.521	0.339
Num. obs.	1261	1261	631

(b) Lucid, Feb. 2020

Term	Overall	Within	Between
α Constant	0.335^{**}		0.379^{**}
	(0.028)		(0.041)
β Certainty	0.477^{**}	0.563^{**}	0.419^{**}
	(0.037)	(0.040)	(0.054)
Adj. R^2	0.102	0.289	0.108
Num. obs.	2399	2399	523

(c) MTurk, March 2020

Term	Overall	Within	Between
α Constant	0.166^{**}		0.214^{**}
	(0.019)		(0.032)
β Certainty	0.700^{**}	0.740^{**}	0.638^{**}
	(0.023)	(0.025)	(0.040)
Adj. R^2	0.233	0.430	0.202
Num. obs.	3417	3417	926

Note: This table displays OLS estimates of the parameters in (4.2), (4.3), and (4.4), with revealed belief substituted for accuracy. Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

Compared with the results above, the certainty scales' general tendency to perform equally well to slightly worse on the revealed measure of belief is encouraging: even when using a distinct measure that reveals beliefs through a complex set of choices, the certainty scales predict what they ought to predict. At the same time, the instances in which the results are weaker are a reminder that, like all survey measures, these measures are affected by measurement error. It is not obvious whether this weaker relationship is due to some flaw in the stated measures, the revealed measures, or some idiosyncratic difference between the two measures.

Application: Branching Versus All-In-One Scales

In political science research that takes a specific interest in confidence or certainty, the dominant paradigm is to use a branching scale: first ask the respondent to answer the question as they would in a normal survey, then ask a follow-up question about certainty. However, sliders and Likert scales with probabilistic scale point labels are also commonly used in social science research. Here, these are referred to as *all-in-one* measures to represent the fact that at least in theory, a single scale can capture both the respondent's best guess and the probability they assign to it.

The received wisdom from survey research is that breaking up tasks through a branching approach tends to improve measurement (Krosnick and Berent 1993; Tourangeau et al. 2000). Regardless of which approach is better or worse, it is also worth understanding the extent of any quality tradeoff between branching and all-in-one scales, as survey researchers typically work with a limited time or question budget.

The framework deployed so far can be used as a set of strategies to resolve such questions empirically. To compare the measurement properties of the branching and all-in-one approaches, the October 2019 MTurk survey randomly assigned respondents to answer identical questions using either a branching scale like the ones described above, or an all-in-one scale. Respondents using the branching scale first stated their best guess, then their certainty level on a 50-100 scale with the labels "don't know," "probably [answer]," and "definitely [answer]," where [answer] was filled in with the respondent's best guess. Respondents using the all-in-one scale responded on a 0-100 scale, with five scale points: "definitely [answer 1]," "probably [answer 1]," "don't know," "probably [answer 2]," and "definitely [answer 2]." This created complete symmetry in scale point labels, regardless of the measure.

In one sense, both strategies elicit information that is completely comparable: the information measured with either method can be converted to a 0 to 1 probability that the respondent assigns to the correct answer. However, the all-in-one scale does not directly ask respondents for their best guess. For respondents who did not choose "50," the best guess was assumed to be the option to which the respondent assigned more probability. Respondents who chose 50 on the all-in-one scale were asked a follow-up question about their best guess.

Analysis

The survey was designed to allow comparisons using all four of the performance metrics considered above: the relationship between certainty and accuracy, response stability, belief stability, and revealed belief. As a first look at the two scales' relative performance on these measures, Figure 4.4 plots them against one another. In each panel, the x-axis is the respondent's stated level of certainty on the initial measure of belief. The y-axis label indicates the outcome plotted in each panel.

On all four measures, the two scales perform in a strikingly similar manner. Respondents who claim to be completely uncertain are no more accurate than chance, slightly more stable than chance, and reveal a slightly higher degree of belief in their initial response than they would by chance. Respondents who indicate complete certainty about their answer are almost always correct, stable in their best guesses and their probabilistic beliefs, and reveal a high degree of belief in their initial response. Visually, there is no clear evidence of a performance tradeoff between the two measures.

To make these comparisons more precise, the regression tests carried out above can be applied to these measures. For each of accuracy, belief stability, and revealed belief, I used OLS to estimate the parameters in

$$Y_{ij} = \alpha + \beta_1 c_{ij} + \beta_2 branch_i + \beta_3 (c_{ij} \times branch_i) + \epsilon ij \tag{4.7}$$

where Y_{ij} is the outcome measure, $branch_i$ equals 0 if the respondent is assigned to the all-in-one scale and 1 if the respondent is assigned to the branching scale, and all other terms are as defined above. Here, β_1 has the same interpretation as β above. β_2 is the predicted difference in means when $c_i = 0$. The most important coefficient estimate is the estimate of β_3. The simplest interpretation of β_3 is that a larger value indicates better performance for the branching scale, while a smaller value indicates worse performance. Impressionistically, Figure 4.4 does not show any substantial differences in non-linearities that might complicate this interpretation.

The regression tests confirm the visual impressions: neither format has a clear performance advantage over the others. In Table 4.7, the first two columns estimate the relationship separately

Figure 4.4: Stability measures, branching versus all-in-one scales.

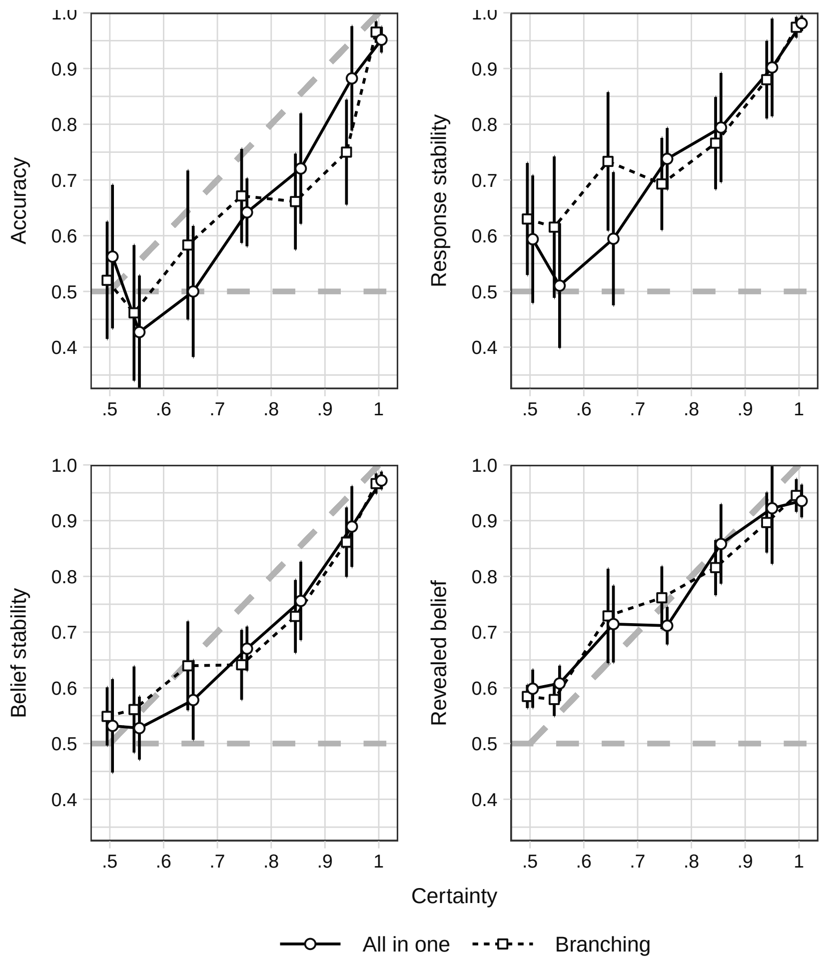

Note: This figure displays all of the outcome measures used earlier in the chapter, but split by whether the respondent was randomly assigned to the all-in-one scale (solid line) or the branching scale (dashed line). Dots are point estimates; vertical bars, 95% confidence intervals. Appendix B presents all of these estimates in tabular form.

Table 4.7: Regression results: branching versus all-in-one scales.

(a) Accuracy

	All in one	Branching	Comparison
Constant	-0.123	-0.012	-0.123
	(0.068)	(0.072)	(0.068)
Certainty	1.054^{**}	0.916^{**}	1.054^{**}
	(0.075)	(0.079)	(0.075)
Branching			0.110
			(0.099)
Certainty × Branching			-0.138
			(0.109)
Adj. R^2	0.188	0.145	0.167
Num. obs.	955	936	1891

(b) Belief stability

	All in one	Branching	Comparison
Constant	-0.024	0.056	-0.024
	(0.047)	(0.045)	(0.047)
Certainty	0.971^{**}	0.873^{**}	0.971^{**}
	(0.052)	(0.051)	(0.052)
Branching			0.079
			(0.065)
Certainty × Branching			-0.098
			(0.073)
Adj. R^2	0.433	0.330	0.377
Num. obs.	577	656	1233

(c) Revealed belief

	All in one	Branching	Comparison
Constant	0.234^{**}	0.211^{**}	0.234^{**}
	(0.028)	(0.025)	(0.028)
Certainty	0.689^{**}	0.727^{**}	0.689^{**}
	(0.037)	(0.033)	(0.037)
Branching			-0.024
			(0.037)
Certainty × Branching			0.038
			(0.049)
Adj. R^2	0.435	0.514	0.475
Num. obs.	318	312	630

Note: This table displays OLS estimates of the parameters in (4.7). Standard errors clustered by respondent. $^{*}p < 0.05$, $^{**}p < 0.01$.

for the all-in-one and branching scales, while the third column presents the estimates of (4.7). In the tests of the certainty-accuracy and certainty-belief stability relationship, the estimates of β_3 are statistically insignificant, though their negative signs imply that the branching scale does a little worse.

The only statistically detectable difference between the two measures of certainty comes when they are compared with the revealed measure of belief. Here, the estimate of β_3 is positive and statistically significant, indicating a tighter correspondence on the branching scale. Though this is an encouraging result for the branching scale, note that the revealed measure of belief was more similar in structure to the branching scale: as described above, respondents first chose a best guess before making the lottery choices, which is more similar in structure to the branching measure than the all-in-one measure.

The key implication of this comparison between all-in-one and branching scales is that both perform similarly in capturing information about where respondents lie on the spectrum between knowledge and ignorance. Indications of certainty and uncertainty are about equally meaningful, regardless of this measurement choice. The remainder of this dissertation makes exclusive use of branching scales. However, based on this test, it seems likely that researchers can assume that all-in-one scales can also safely be interpreted as capturing the same type of variation in the respondent's probabilistic belief in their answer.

Implications

This chapter applied a series of tests to examine the extent to which certainty scales identify meaningful variation along a spectrum between knowledge and ignorance, where knowledge was defined as complete belief in the correct answer and ignorance was defined as thinking the correct and incorrect answers are equally likely to be true. The results were encouraging. Certainty scales captured a substantial amount of variation in four separate measures that indicate the degree to which respondents truly know, and think they know, the answer.

For research that aims to measure survey respondents' beliefs about politically controversial facts, the success of these tests has two key implications. First, certainty scales can capture meaningful variation in how much people think they know about survey questions. This does not mean that they will be equally successful for all question types, but it does mean that a high level of success is possible.

Second, the performance of certainty scales can be objectively assessed using a range of outcome metrics. Though researchers who use certainty scales generally take for granted their ability to measure variation in respondents' degree of belief in their answers, it is not necessary to do so. Similarly, skeptics of whether people can meaningfully report variation in the probability that they assign to their answer should not simply fall back on an argument that people are generally innumerate or bad with probabilities. If certainty scales behave as they should, the survey has successfully tapped variation in respondents' degree of belief in their answers, regardless of whether those same respondents are well-versed in the axioms of probability. The latent probabilistic belief can be tapped even if respondents do not literally think that way.

Before getting too excited about certainty scales' strong performance in this chapter, it is important to be cognizant that political awareness questions are a relatively easy test. For the trivia-style facts considered here, there are good reasons to think that some respondents will know the answers, no reason to think that respondents will hear and accept false claims, and no reason to think that respondents will misrepresent their beliefs (e.g., by responding expressively [Bullock et al. 2015; Prior et al. 2015; Schaffner and Luks 2018] or by feigning uncertainty to conceal socially undesirable beliefs [Berinsky 1999]). The next chapter applies the same strategies to a harder test: the measurement of political misperceptions.

Chapter 5

Incorrect Knowledge or Miseducated Guesses?

Abstract. This chapter examines the sense in which incorrect answers indicate that respondents have been misinformed. To do so, it applies the same strategies that, in the previous chapter, show that certainty scales capture the full spectrum from knowledge to ignorance on factual questions that gauge general political awareness. In contrast to these results, the chapter finds that claims to be certain of incorrect answers to questions about politicized controversies, the state of the economy, and status quo government policies do not represent "incorrect knowledge." Though certainty is somewhat predictive of belief stability and revealed belief on these questions, it falls well short of the benchmarks set by political awareness questions. This suggests that incorrect answers represent misperceptions only in the probabilistic sense: respondents sometimes consistently draw on the same misleading heuristics to make their inference, but are not stating pre-existing beliefs. To describe incorrect answers that are believed probabilistically but not in the threshold sense, I introduce the term *miseducated guess*.

We saw in Chapter 3 that "don't know" (DK) response options fail to isolate respondents who claim to be certain about their answers. Although these particular results are novel, scholars have a long history of arguing that allowing respondents to say DK is not sufficient to distinguish the misinformed from the uninformed. This perspective originates with Kuklinski et al. (1998, 2000) and Hofstetter et al. (1999), who wrote at a time when misinformation was only beginning to be recognized as a problem in American politics. To distinguish the "guessing uninformed" (who "do not hold factual beliefs at all") from the "genuinely misinformed" (who "firmly hold beliefs that happen to be wrong"), Kuklinski and colleagues recommend asking respondents how certain they are about their answers (Kuklinski et al. 2000, 793). In the terms defined in Chapter 2, the recommended empirical approach embraces a threshold conception of belief: to be considered misinformed, one's degree of certainty must surpass a minimum level.

Kuklinski and colleagues' threshold conception has proven influential. Review pieces by leading scholars consistently define a person who holds a misperception, or who has been misinformed, as someone who is certain of the wrong answer (Flynn, Nyhan and Reifler 2017; Jerit and Zhao 2020). A growing body of empirical research applies the same standard (Jerit et al. 2020; Lee and Matsuo 2018; Marietta and Barker 2019; Pasek et al. 2015; Peterson and Iyengar 2020). Luskin et al. (2018) call methods that focus on those who claim to be certain a "24 carat gold standard" for researchers who want to identify misinformed respondents without misinforming their audience.

Despite the growing popularity of this approach, no empirical research examines the properties of this approach to measuring misperceptions. This chapter applies the strategies developed in Chapter 4 to these measures. The findings there provide some reasons for optimism: on general political awareness questions, it is possible to use certainty scales to distinguish between the knowledgeable from the ignorant. If certainty scales perform equally well for incorrect responses to more politicized items, perhaps the gold standard for measuring misperceptions can achieve its face-value purpose.

The findings problematize analysis that takes claims to be certain of falsehoods at face value. Even as respondents who provide correct answers to questions about politicized controversies and the state of the economy frequently approach the benchmarks attained in the previous chapter, respondents who answer incorrectly fall well short. In the survey environment, claims to be certain that President Obama has never released his birth certificate (he has) are no more meaningful than claims to be certain that Republicans control the House of Representatives (which is currently controlled by Democrats).

The implication of these findings is that claims to be certain and wrong do not represent the "incorrect knowledge" that many researchers think they are observing (Hochschild and Einstein 2015, 10). If incorrect answers are not incorrect knowledge, then what are they? The next section develops a theoretical case for an alternative interpretation of claims to be certain about the wrong answer: a *miseducated guess*, defined as an incorrect response that is based on a heuristic or inference that ultimately proves to be misleading.

The Two Meanings of "Misinformed"

In what sense have respondents who claim to be confident about the wrong answer been misinformed? One possibility is that they have been misinformed about the facts themselves. This is consistent with the bulk of research on misinformation in the information environment, which defines misinformation in terms of outright false, purportedly-factual claims and showcases such claims in its motivating examples (e.g., Lewandowsky et al. 2005, 2012; Thorson 2015; Swire et al. 2017).[1] Consistent with such research, respondents who are certain and wrong are said to "firmly hold the wrong information" (Kuklinski et al. 2000, 792) or "hold a false or unsupported belief about the answer" (Flynn et al. 2017, 127). At face value, these accounts — as well as the respondents' own claims to be absolutely certain about their answer — would seem to indicate that certainty scales can identify respondents who had a pre-existing belief in a specific piece of false information.

Notwithstanding the substantial degree of focus on outright false claims, authoritative definitions of misinformation also encompass information that is misleading (e.g., Lazer et al. 2018). This perspective makes occasional appearances in misinformation research: misinformation has been defined to include over-representation of successful cardiopulmonary resuscitation (CPR) on television (Diem et al. 1996), misleading suggestions in news headlines (Ecker et al. 2014), and misleading advertising (Glaeser and Ujhelyi 2010). Accounts that describe respondents who claim to be certain about incorrect answers as "misinformed" leave a role for "inferential reasoning" (Hofstetter et al. 1999, 353), "mistaken inferences" (Flynn et al. 2017, 128) and "translat[ing] ... general notions into more specific ones" (Kuklinski et al. 2000, 795). "For instance, a person who hears a [Rush] Limbaugh comment about Hillary Clinton that is laden with sarcasm might infer that Mrs. Clinton engaged in corrupt practices as an attorney, even though that was not explicitly stated" (Hofstetter et al. 1999, 360-61).

This raises the possibility that respondents who claim to be certain of the incorrect answer may be making "miseducated guesses" that are not actually a product of false information about the fact in question. No prior research provides an evidence-based evaluation of the difference between these two interpretations, either for incorrect answers in general or for incorrect answers about which the respondent claims to be certain.

How could mistaken inferential reasoning lead respondents to claim absolute certainty about factual beliefs that they did not hold before entering the survey environment? It is widely recognized that surveys do not generally elicit "decided," "fixed," or "true" beliefs that existed before the survey (Flynn, Nyhan and Reifler 2017, 140; Bullock and Lenz 2019, 326; Berinsky 2017*b*, 317). Instead, surveys induce respondents to construct beliefs on the spot by "sampling" their perceptions and integrating them into a summary judgment (Nadeau and Niemi 1995; Tourangeau et al. 2000; Strack and Martin 1987). A little-recognized implication of such accounts is that the process of "sampling" considerations can affect the extent to which respondents appear to be certain about their answers. Zaller and Feldman (1992) describe such a case:

[1] These definitions often encompass claims that are presented as true, but later turn out to be false.

> Consider respondent A. His first reaction to a guaranteed standard of living was that it was inconsistent with American ideals; he was also bothered by the unfairness of supporting those who refuse to work. Yet he worried about letting individuals get ahead on their own, saying that some people need special help and that society has an obligation to help the needy. In the second interview, however, there was no sign of this ambivalence. Respondent A gave six reasons why individuals ought to get ahead on their own, including a restatement of his feeling that job guarantees are un-American, without raising any opposing considerations. Respondent A ... went from being an ambivalent conservative on this issue to being a confident conservative. (594-596)

A similar phenomenon may affect the extent to which respondents appear certain of the wrong answers to factual questions. As we saw in the previous chapter, some respondents claim to be certain of the wrong answers to questions for which no observer would expect misinformation about the facts themselves. In a detailed analysis of the September 2017 Lucid survey, Graham (2020) notes that such claims are concentrated on questions that laid traps for respondents: plausible responses that could connect to something the respondent knows, but are wrong. For example, many respondents say that they are certain that the filibuster is the Senate procedure that allows changes to the budget via a simple majority vote (instead of reconciliation) or that Nancy Pelosi is the Senate Minority Leader (instead of Chuck Schumer).

The lack of concrete evidence as to whether respondents who answer incorrectly have been misinformed leaves a potentially fundamental gap between research on misinformation itself and research that identifies members of the public as having been misinformed. Accordingly, this chapter evaluates the approach of focusing on respondents who claim to be certain of the incorrect answer. Do such responses signify that the respondent actually possesses "incorrect knowledge," or does it identify respondents who are a bit more confident in their miseducated guesses?

Hypotheses

To evaluate the sense in which claims to be certain of the wrong answer indicate that a respondent has been misinformed, I turn to the latter two strategies developed in Chapter 4. First, in three separate sets of data, I examine the degree to which certainty scales predict response stability and belief stability. Second, I examine the degree to which certainty predicts revealed belief. The subject matter of the questions, which covers politicized controversies, the state of the economy, and government policies, will be introduced in more detail in the pertinent sections of the analysis.

The Nature of Claims to be Certain

To set expectations for this analysis, it is useful to specify some hypotheses that more explicitly relate the discussion in the previous section to what will be observed in the data. The *incorrect knowledge hypothesis* holds that claims to be certain of incorrect answers can be taken completely at face value. To the extent that this hypothesis is true, measures of certainty about incorrect responses should meet, or at least come close, to the benchmarks set by the general political awareness questions analyzed in Chapter 4: stability should steadily trend upward in certainty, culminating in close-to-perfect stability among respondents who claim to be certain. This benchmark allows the same degree of measurement error that afflicts certainty scales on political awareness questions, but no more.

The other extreme would be that incorrect answers all originate where they were traditionally assumed to originate: out of ignorance. The *total ignorance hypothesis* holds that incorrect answers should be completely unstable, as posited by the extreme form of Converse's (1964) black and white model. Claims to be certain of falsehoods would be entirely the product of measurement error, and would consequently be no more stable than responses chosen completely at random.

Between these extremes, claims to be certain of incorrect answers may indicate that respondents have meaningful beliefs in the probabilistic sense, even if they have not been misinformed about the facts themselves. The *miseducated guesses hypothesis* holds that incorrect answers vary meaningfully in the extent to which respondents believe them, but that stability never reaches the benchmarks set by the performance of political awareness questions. This interpretation allows that many incorrect responses reflect exposure to misleading information or a reliance on misleading heuristics, but problematizes the stronger interpretations that pervade prevailing accounts.

The Role of Partisanship

Partisanship has a prominent place in research on misperceptions. Define a *congenial* response as an answer that, if true, would benefit one's preferred party or harm the other party. For example, consider the false claim that former Democratic President Obama has never released his birth certificate. This is congenial to Republicans because it would benefit their side if it were true: it would cast doubt on the legitimacy of Obama's presidency and affirm claims made by a popular Republican politician, President Donald Trump. Conversely, the correct answer on a question about this topic is congenial to Democrats. Obama did in fact release his birth certificate, and was in fact eligible for the presidency.

Existing research provides two reasons to think that incorrect answers might be more strongly believed when they are congenial. First, partisans are more likely to be exposed to congenial information, both in their news consumption (Gentzkow and Shapiro 2010; Stroud 2010) and in other informational channels like the people with whom they associate (Mutz 1998; Lang and Pearson-Merkowitz 2015). People with a strong preference for pro-attitudinal information are especially likely to be exposed to false claims online (Guess et al. 2020). Second, partisans have been found to more readily accept claims that support their predispositions (Nyhan and Reifler 2010; Prasad et al. 2009; Taber and Lodge 2006).

Accordingly, the *partisan difference in interpretation* hypothesis holds that incorrect answers should be viewed differently depending on the partisanship of the respondent supplying them. In particular, incorrect answers should be more meaningful when they are supplied by the party that theory predicts will have greater exposure to, and acceptance of, false claims.

On the flip side, research also provides reasons to think that on any given question, both Democrats' and Republicans' incorrect answers might reflect misperceptions to a more similar degree. Most Americans consume balanced media diets (Guess 2020). Even when exposed to uncongenial arguments or factual information, people usually update their beliefs in the direction of the evidence (Guess and Coppock 2018; Wood and Porter 2018).

Accordingly, the alternative hypothesis is the *no partisan difference in interpretation* hypothesis. According to this hypothesis, the typical claim to be certain about an incorrect answer means about the same thing, regardless of the respondent's partisanship.

It is worth emphasizing that these hypotheses about partisanship concern the measurement properties of the certainty scales, not whether Democrats and Republicans have different beliefs or are more or less likely to believe false things.

The Instability of Incorrect Answers

For the first test of the sense in which respondents who claim to be certain of incorrect answers have been misinformed, this section uses data from two panel surveys conducted on Amazon Mechanical Turk. One of these surveys, conducted in March and August 2020, was analyzed in the previous section. The other survey was conducted in June 2019 and June 2020. Both surveys included 50-100 certainty scales of the type described in the previous section and pictured in Appendix E.

The surveys included six total questions about politicized controversies. All but one of these controversial matters has been the subject of false claims in the public sphere. The exception is a question that taps partisan perceptions of the abuse of executive power, the importance of which is discussed in more detail the next chapter. The questions were selected to have partisan balance.

Three questions had incorrect answers that are congenial to Democrats and, by the converse, correct answers that are congenial to Republicans:

1. **Clinton email.** Respondents were asked whether the following statement is true or false: "While she was Secretary of State, Hillary Clinton used a private email server to send and receive classified information." This was a key controversy in the 2016 election campaign that, after the election, continued to fuel President Trump's demands that Clinton be imprisoned. Both before and after an FBI investigation revealed that Clinton had sent classified information, she falsely claimed that she had not.[2]

2. **Obama DAPA reversal.** Respondents were provided with a one-sentence description of the Deferred Action for Parents of Americans (DAPA) program, an Obama initiative that was struck down in court. They were then asked whether the following statement is true or false: "About a year earlier, Obama said that he would be ignoring the law if he issued such an order." Obama did say directly that such a program would amount to "ignoring the law," and later falsely claimed not to have changed his position on the matter.[3]

3. **Trump-Russia collusion.** Respondents were provided a one-sentence description of the Robert Mueller's special counsel investigation into Russian interference in the 2016 presidential election. They were then asked whether the following statement is true or false: "Robert Mueller's report stated that President Trump personally conspired with Russia to influence the 2016 election." Prior to the release of the report, many left-leading opinion leaders made confident claims that Mueller would find evidence of such collusion.

The other three questions have incorrect answers that were congenial to Republicans and, conversely, correct answers that are congenial to Democrats:

4. **Obama birth certificate.** Respondents were asked whether the following statement is true or false: "President Obama has never released his birth certificate." This question taps a clearly factual element of a larger conspiracy promoted by President Trump and other right-leaning opinion leaders. Even after Obama released both his short- and later long-form birth certificates, demands that he do so continued to populate public discourse and social media.[4]

5. **Trump Article II.** Respondents were told that Article II of the Constitution describes the President's powers. They were then asked whether the following statement is true or false: "President Trump has said that Article II gives him the power to do whatever he wants." This is the only question that, to my knowledge, has not been the subject of false claims in public discourse.

6. **Trump said 'grab them.'** Respondents were asked whether the following statement is true or false: "Before becoming president, Donald Trump was tape recorded saying that he kisses women and grabs them between the legs without their consent." The release of this tape recording was the source of substantial controversy during the 2016 presidential election. After initially admitting that the tape was authentic, President Trump has since falsely claimed that he was not recorded making this statement statements.[5]

Together, items 1, 3, 5, and 6 cover many of the most prominent partisan, politicized controversies that have been subject to false claims in the public sphere in recent years. Items 2 and

[2] "FBI findings tear holes in Hillary Clinton's email defense," *PolitiFact*, July 6, 2016.

[3] "Barack Obama: Position on immigration action through executive orders 'hasn't changed'," *PolitiFact*, November 20, 2014.

[4] "Fact check: Old fabricated Obama "Kenyan birth certificate" resurfaces," Reuters, June 17, 2020.

[5] "Trump Once Said the 'Access Hollywood' Tape Was Real. Now He's Not Sure." *The New York Times*, November 28, 2017.

5 provide a benchmark of less-prominent controversies while simultaneously bringing the abuse of executive power into the study of partisan perceptions and misperceptions.

Distribution of Responses

Before diving into the measurement properties, it is worth pausing to examine the distributions of the measures. Figure 5.1 plots the distribution of certainty separately for correct and incorrect answers for each of the six questions, as well as the general political awareness questions analyzed in the previous sections. The shaded distributions with solid lines represent correct answers, while the dotted line represents the distribution among incorrect answers. Above each distribution, some further distributional statistics are printed: the percentage of correct answers, average certainty among respondents who answer correctly ($E[c|\text{corr}]$), and average respondents who answer incorrectly ($E[c|\text{incorr}]$).

Two key patterns emerge in Figure 5.1. First, we see that respondents tend not to claim to be certain of incorrect answers. While the certainty distribution among respondents who answered correctly tends to spike at probability 1 (on the right side of the x-axis), the distribution among incorrect answers is spread relatively evenly across the possible values of the certainty scale. However, this pattern is not universal. In particular, the two questions about abuse of executive power lack such a spike.

Second, the political awareness questions are strikingly similar to the other questions in terms of both the shape of the distributions and the average level of certainty. Both types of questions see strikingly similar average confidence among both correct answers (0.85 for controversies and 0.86 for political awareness) and incorrect answers (0.74 for controversies and 0.72 for political awareness). On the 0.5 to 1 scale, this means that the average correct answer is stated with 48 percent and 60 percent greater confidence.[6]

With this context in mind, the analysis now turns to the measurement properties of the certainty scales. Though claims to be certain of the incorrect answer are less common than claims to be certain of the correct answer, the focus of this chapter is how we should interpret claims to be more and less certain of incorrect answers, conditional on respondents having selected such responses.

[6](0.854 - 0.739) / (0.739 - 0.5) = 0.480; (0.857 - 0.724) / (0.724 - 0.5) = 0.598

Figure 5.1: Belief distribution, politicized controversies.

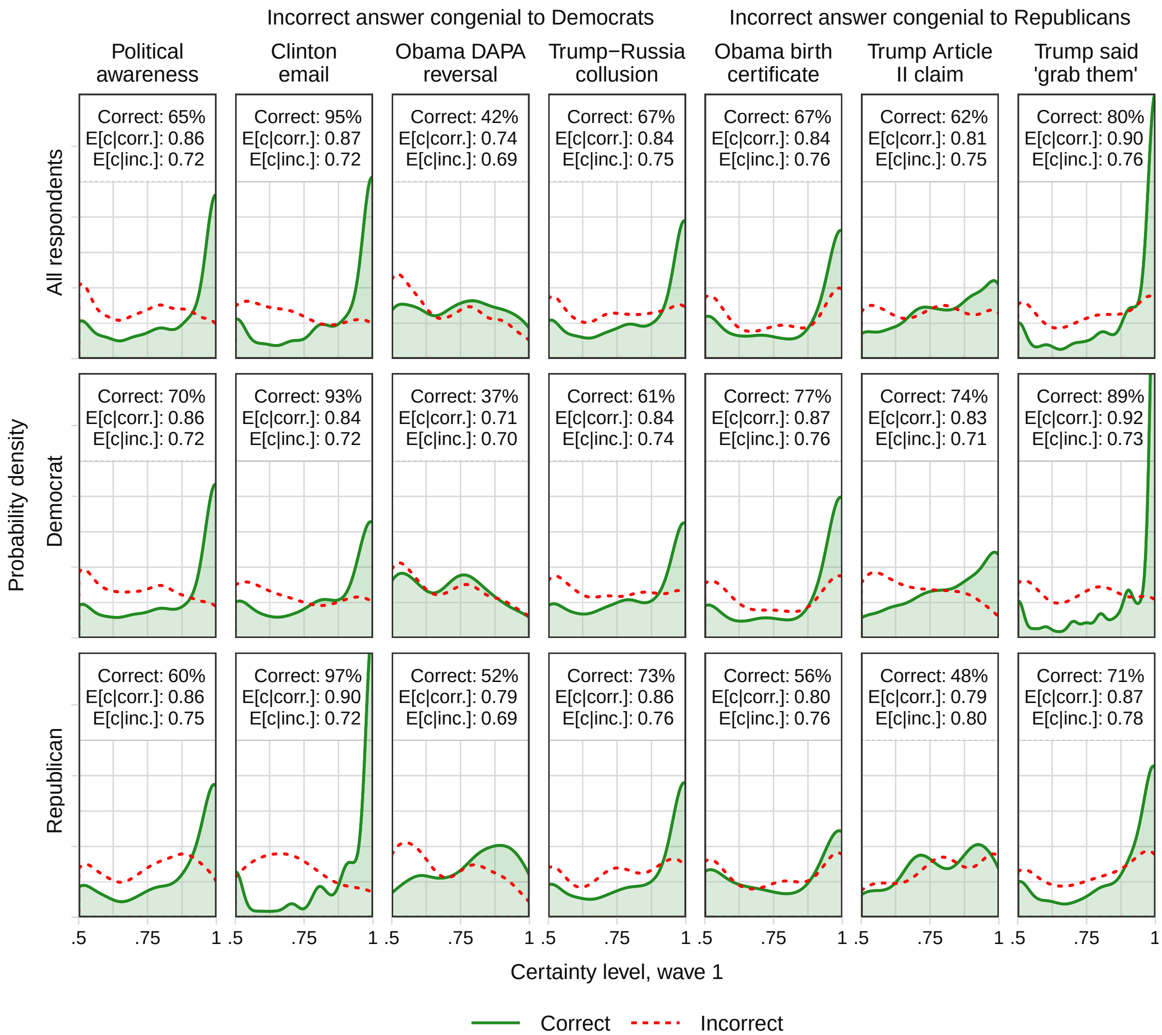

Note: This figure plots the certainty distribution for each question analyzed in this section. The green, shaded distribution with a solid outline represents correct answers, while the red, dashed line represents incorrect answers. Printed above the figure are the percentage of respondents who answered correctly, average certainty among those who answered correctly, and average certainty among those who answered incorrectly.

Belief Stability

To begin the analysis of belief stability among respondents who answer incorrectly, Figure 5.2 plots the results for each of the six controversy questions, as well as the political awareness questions analyzed in Chapter 4. The x-axis and y-axes have the same interpretation as the belief stability figures there. The x-axis plots the respondent's wave 1 certainty level on an 0.5 to 1 probability scale, while the y-axis plots the probability the respondent assigned to their wave 1 response when asked the same question again several months later in wave 2.

The key difference between Figure 5.2 and the belief stability analysis in Chapter 4 is that responses have now been separated by whether the respondent's wave 1 answer was correct. For the study of misperceptions, the key quantity of interest is stability among incorrect answers, which is represented by the hollow red squares connected by dotted lines. Meanwhile, the solid green squares connected by solid lines plot stability among respondents who answered correctly in wave 1.

The first row of Figure 5.2 plots belief stability among all respondents. The leftmost panel plots the same political awareness questions that were analyzed in Chapter 4, but split by whether the answer was correct or incorrect. A fairly substantial gap in belief stability opens up between the two sets of respondents, especially at high levels of certainty. While respondents who claim to be certain of the correct answer assign a probability of 0.973 to it when asked the same question again in a follow-up survey, respondents who claim to be certain of the incorrect answer assign a probability of just 0.769 to it. Both figures indicate a considerably stronger belief than among respondents who indicate a lower level of certainty. However, the certainty scale does a better job of capturing variation in ignorance and knowledge on correct answers than of capturing variation between ignorance and "incorrect knowledge" on incorrect answers.

The first row's other six panels provide the first opportunity to evaluate the three hypotheses laid out above: incorrect knowledge, miseducated guesses, and non-belief. At high levels of certainty, incorrect answers never come close to the last chapter's benchmarks. Pooling across all six questions, the average respondent who claims to be completely certain of an incorrect answer assigns a probability of 0.570 to it, a substantially *lower* figure than on political awareness questions (difference = 0.199, s.e. = 0.080, p = 0.02). Moreover, the most stable beliefs in incorrect answers come on the two least prominent controversies, Trump's claim of unlimited Article II power (0.759) and Obama's reversal of his DAPA position (0.799).

Though the stability of belief in incorrect answers consistently falls short of the political awareness benchmark at high levels of certainty, the controversy questions include two cases in

Figure 5.2: Belief stability, politicized controversies.

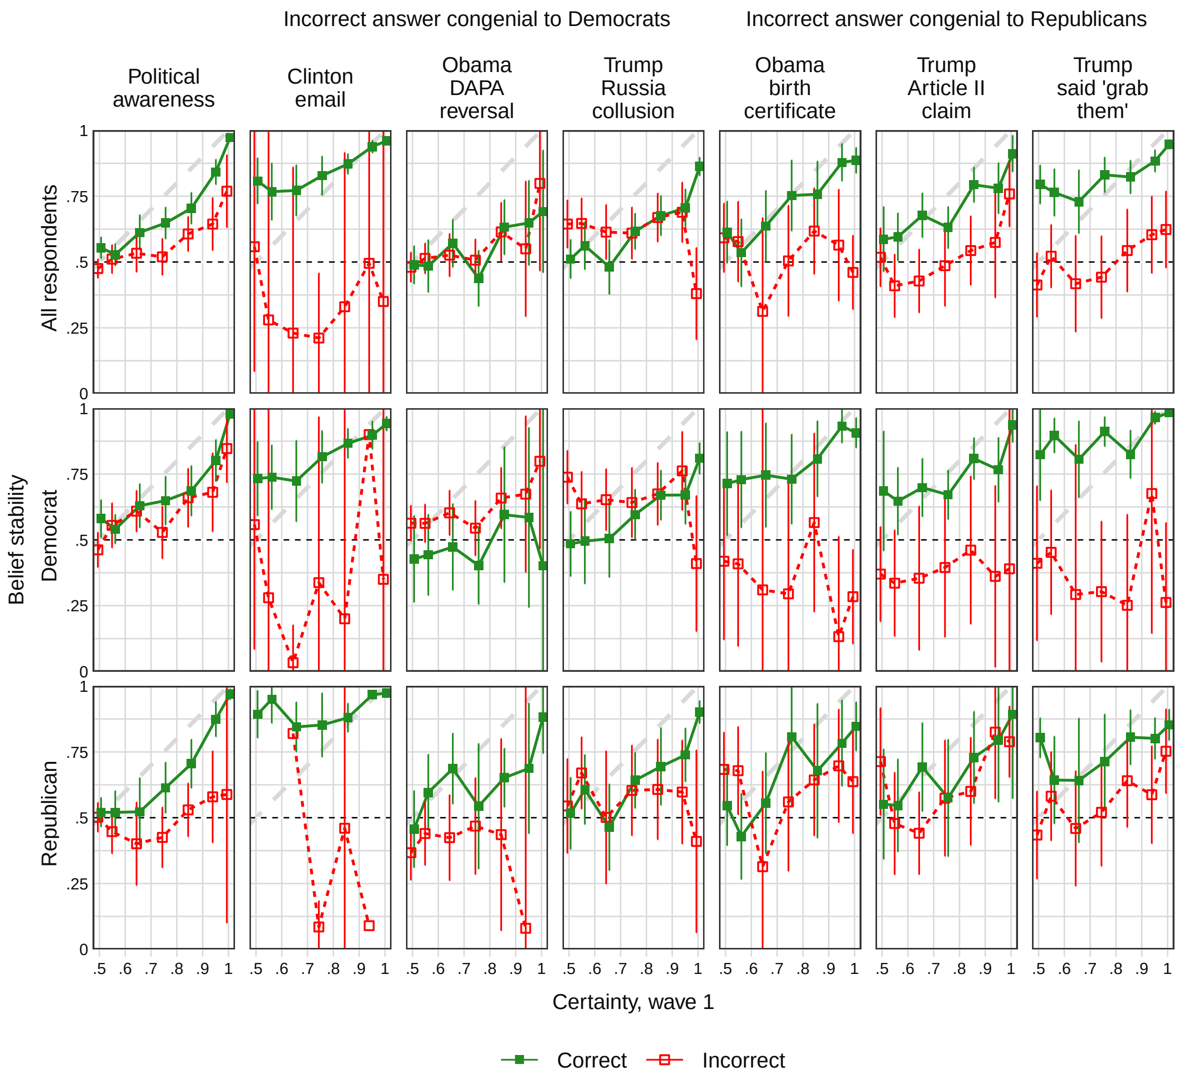

Note: This figure plots belief stability for each of the questions analyzed in this section. The green, solid squares connected by solid lines represent respondents who answered correctly in wave 1, while the red, hollow squares represent those who answered incorrectly. In all other ways, this figure is identical to the bottom row of Figure 4.2. Here and in all subsequent figures, these measures are binned using the following groups: 0.5; [0.51, 0.59]; [0.6, 0.69]; [0.7, 0.79]; [0.8, 0.89]; [0.9, 0.99]; 1. Appendix C presents all of these estimates in tabular form.

which correct and incorrect answers are about as stable as one another: Obama's DAPA reversal, and allegations that Trump personally colluded with Russia in 2016. With respect to Obama's DAPA reversal, correct answers also have fairly low belief stability across the board; among respondents who answered correctly, belief stability tops out at 0.692 among those who claim to be certain. The Trump-Russia item sees more of a spike among the completely certain, to 0.863. These two questions highlight that just as incorrect answers do not generally capture "incorrect knowledge," correct answers do not always capture correct knowledge either. To the degree that correct and incorrect answers converge, it is generally because the correct answers should also be viewed as guesses, not because the incorrect answers are incorrect knowledge.

On the three questions not yet summarized — Clinton email, Obama birth certificate, and Trump said 'grab them' — the difference between correct and incorrect answers is most stark. On these questions, respondents who claim to be certain of the correct answer seem to actually know it, while claims to be certain of the incorrect answer are highly unstable. Particularly striking is the finding that respondents who say they are certain that Obama never released his birth certificate are no more stable than chance.

These results show that for these six items, the incorrect knowledge hypothesis clearly fails. Support for the miseducated guess hypothesis is stronger: on at least some questions, belief stability seems to be a bit higher among respondents who claim to be certain of the incorrect answer. However, to the extent that the miseducated guess hypothesis is supported, it is supported in the same sense for political awareness questions.

The bottom two rows of Figure 5.2 display the same results split by political partisanship. Visually, they suggest that there may be something to the difference in interpretation hypothesis. On the three questions for which the incorrect answer is congenial to Democrats, Democrats appear to be more stable in their incorrect answers than Republicans, especially on the Obama DAPA question. On the three questions for which the incorrect answer is congenial to Republicans, the pattern is starker: Republicans who claim to be certain of the correct answer have belief stability figures of 0.637, 0.789, and 0.753, while for Democrats, these figures are 0.284, 0.390, and 0.262.

For the partisan difference in interpretations hypothesis, the graphical analysis furnishes two preliminary takeaways. First, partisan differences in stability are not strong enough to confirm the incorrect knowledge hypothesis for one party but not the other. Even the relatively high stability among Republicans on congenial, incorrect answers falls far short of the political awareness questions. Second, partisan differences in stability might be sufficient to predict a difference between miseducated guessing and complete ignorance. For more precise leverage on this question, I now turn to a more precise test.

Regression Test

To summarize the differences between correct and incorrect answers, I turn to OLS regression. For each of the same three groups of respondents plotted in Figure 5.2, I use OLS to estimate the parameters in

$$p_{i2j}^{\text{initial}} = \alpha + \beta_1 c_{i1j} + \beta_2 a_{i1j} + \beta_3 (c_{i1j} \times a_{i1j}) + \epsilon_{ij} \tag{5.1}$$

where i indexes respondents, j indexes questions, and $t \in \{1, 2\}$ indexes waves 1 and 2. As in the regressions specified in Chapter 4, the variables $c_{i1j} \in [0.5, 1]$ and $p_{i2j}^{\text{initial}} \in [0, 1]$ respectively measure the respondent's probabilistic certainty level and their probabilistic belief in their initial response. These tests are modified to examine the difference between correct and incorrect answers by the inclusion of $a_{i1j} \in \{0, 1\}$. Just as in Chapter 4, a_{i1j} measures whether the respondent's *a*nswer at $t = 1$ was correct (1) or incorrect (0). In this regression, β_1 is the main parameter of interest. It measures the slope of the relationship between certainty and response stability for incorrect answers. The other key parameter is β_3. It measures the difference in slopes between correct and incorrect answers. Consequently, the slope for correct answers is equal to $\beta_1 + \beta_3$.

To examine the hypothesis of a partisan difference in interpretation, this regression specification is adapted to test for partisan differences. For a comparison between Democrats and Republicans, I use the data from only these two groups of respondents. Though political independents can be included without changing the parameter estimates, doing so would serve only to distract from the comparison of interest. I use OLS to estimate the parameters in

$$\begin{aligned} p_{i2j}^{\text{initial}} = {} & \alpha_D + \beta_1 c_{i1j} + \beta_2 a_{i1j} + \beta_3 (c_{i1j} \times a_{i1j}) \\ & + \alpha_R R_i + \beta_4 (R_i \times c_{i1j}) + \beta_5 (R_i \times a_{i1j}) + \beta_6 (R_i \times c_{i1j} \times a_{i1j}) + \epsilon_{ij} \end{aligned} \tag{5.2}$$

where $R_i \in \{0, 1\}$ is an indicator for being a Republican (1) or a Democrat (0), including leaners. The parameters in the top row have the same interpretation as in equation (5.1) but now only apply to Democrats. The new parameters in the bottom row measure the difference between the parameter estimates for Democrats and Republicans. In this regression β_4 is the main parameter of interest. It measures the difference in the certainty-stability relationship between Democrats and Republicans on incorrect answers. For correct answers, the difference between Democrats and Republicans is

Table 5.1: Regression test: belief stability, political awareness questions.

Term	All respondents	Democrats	Republicans	Comparison
α Constant	0.278^{**}	0.234^{**}	0.370^{**}	0.234^{**}
	(0.050)	(0.077)	(0.076)	(0.077)
β_1 Certainty	0.385^{**}	0.497^{**}	0.161	0.497^{**}
	(0.074)	(0.113)	(0.114)	(0.113)
β_2 Correct	-0.288^{**}	-0.222^{*}	-0.444^{**}	-0.222^{*}
	(0.059)	(0.090)	(0.093)	(0.090)
β_3 Certainty $\times$ Correct	0.563^{**}	0.429^{**}	0.852^{**}	0.429^{**}
	(0.082)	(0.124)	(0.129)	(0.124)
α_R Republican				0.136
				(0.109)
β_4 Certainty $\times$ Rep.				-0.336^{*}
				(0.161)
β_5 Republican $\times$ Correct				-0.222
				(0.129)
β_6 Certainty $\times$ Corr. $\times$ Rep.				0.422^{*}
				(0.178)
Adj. R^2	0.413	0.372	0.458	0.414
Num. obs.	1572	842	530	1372

Note: This table presents OLS estimates of the parameters in (5.1) and (5.2). Standard errors clustered by respondent. $^{*}p < 0.05$, $^{**}p < 0.01$.

measured by $\beta_4 + \beta_6$.

Table 5.1 displays the results for the political awareness benchmark. The first three columns display the estimates of (5.1), while the last column displays the estimates of (5.2). These estimates are equivalent to those displayed in Table 4.4, but split by party and whether the answer is correct or incorrect.

Confirming the visual impressions from Figure 5.2, the estimate of β_1 suggests that for incorrect answers, certainty predicts some meaningful variation in response stability. However, the estimate of β_3 is larger in magnitude, indicating that claims to be certain are a bit more than twice as predictive of greater stability for correct answers than for incorrect answers. Whereas certainty scales capture the full range from knowledge to ignorance on correct answers, they only partly capture this range for incorrect answers.

Table 5.2a displays the results pooling across the three questions on which the incorrect answer is congenial to Democrats. None of the estimates of β_1 are statistically significant, and all of them are small in substantive terms, never exceeding 0.08. The estimate of β_4 shows no evidence partisan difference. This suggests that for these items, incorrect answers are indicators of complete ignorance, regardless of what the respondent claims about their certainty level.

By contrast, the estimates of β_3 are large and substantively significant, all in about the 0.6

Table 5.2: Regression test: belief stability, partisan questions.

(a) Incorrect answer congenial to Democrats

Term	All respondents	Democrats	Republicans	Comparison
α Constant	0.510^{**}	0.567^{**}	0.465^{**}	0.567^{**}
	(0.060)	(0.076)	(0.102)	(0.076)
β_1 Certainty	0.077	0.064	0.031	0.064
	(0.088)	(0.112)	(0.147)	(0.112)
β_2 Correct	-0.317^{**}	-0.361^{**}	-0.228	-0.361^{**}
	(0.070)	(0.094)	(0.115)	(0.094)
β_3 Certainty $\times$ Correct	0.595^{**}	0.553^{**}	0.633^{**}	0.553^{**}
	(0.096)	(0.128)	(0.157)	(0.128)
α_R Republican				-0.101
				(0.127)
β_4 Certainty $\times$ Rep.				-0.033
				(0.184)
β_5 Republican $\times$ Correct				0.133
				(0.149)
β_6 Certainty $\times$ Corr. $\times$ Rep.				0.081
				(0.203)
Adj. R^2	0.184	0.100	0.293	0.193
Num. obs.	1796	940	754	1694

(b) Incorrect answer congenial to Republicans

Term	All respondents	Democrats	Republicans	Comparison
α Constant	0.352^{**}	0.479^{**}	0.356^{**}	0.479^{**}
	(0.065)	(0.116)	(0.079)	(0.116)
β_1 Certainty	0.233^{**}	-0.145	0.334^{**}	-0.145
	(0.085)	(0.150)	(0.100)	(0.150)
β_2 Correct	-0.045	-0.046	-0.057	-0.046
	(0.075)	(0.123)	(0.100)	(0.123)
β_3 Certainty $\times$ Correct	0.373^{**}	0.656^{**}	0.205	0.656^{**}
	(0.094)	(0.155)	(0.122)	(0.155)
α_R Republican				-0.123
				(0.140)
β_4 Certainty $\times$ Rep.				0.479^{**}
				(0.180)
β_5 Republican $\times$ Correct				-0.011
				(0.158)
β_6 Certainty $\times$ Corr. $\times$ Rep.				-0.451^{*}
				(0.197)
Adj. R^2	0.240	0.428	0.099	0.289
Num. obs.	1797	942	754	1696

Note: This table presents OLS estimates of the parameters in (5.1) and (5.2). Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

range. This implies that claims to be certain of the correct answers to these questions captures meaningful variation in the respondents' probabilistic beliefs. The relationship between certainty and stability is weaker than observed for the political awareness questions, but clearly indicates that respondents who claim to be more certain are consistently drawing on the same, informative-on-average heuristics.

Turning to the questions on which the incorrect answer is congenial to Republicans, Table 5.2b displays the estimates. For these questions, a statistically significant relationship between certainty and stability can be detected: the estimate of β_1 for Republicans suggests that a one-unit increase in certainty translates to one-third of a unit increase in belief stability. This is about equal to the relationship between certainty and stability for incorrect answers to all respondents on political awareness questions (β_1 in Table 5.1). For Democrats, a statistically significant relationship cannot be detected. The estimate of the difference between Democrats and Republicans, β_4, is large and statistically significant. For this set of questions, the partisan difference in interpretations hypothesis appears to be confirmed: incorrect answers represent complete ignorance for Democrats, but a miseducated guess for Republicans.

Takeaways

Together, these results suggest that incorrect answers represent, at best, miseducated guesses. Even as certainty is strongly predictive of belief stability among respondents who answer correctly, the relationship between certainty and stability for incorrect answers is weak to nonexistent. Though there is some evidence of a partisan difference in interpretation, this difference is between total ignorance and miseducated guesses. No estimate of the stability of belief in incorrect answers comes close to the benchmarks set by general political awareness questions, on which certainty scales appeared to successfully capture the full range of variation between knowledge and ignorance.

There are two main reasons to doubt that these results should be taken at face value. First, stability may be affected by expressive responding. If Democrats and Republicans systematically misrepresent their beliefs as more congenial to their party than they really are, it is possible that stability on incorrect answers, as well as partisan differences in certainty, reflect a tendency to respond expressively rather than revealing one's sincere best guess.

Second, the classic worry with respect to response stability is that instability can occur due to genuine change in beliefs (Converse 1964). This is particularly worrisome for an asymmetry between correct and incorrect answers. If respondents learn the correct answer between the two surveys, incorrect responses may be less stable not because they are less reflective of what the respondent actually thinks about the world, but because the respondent's beliefs changed between the two surveys.

Reassuringly, the share of respondents answering correctly did not increase between the two surveys, which provides some evidence that learning did not cause the difference in stability between the correct and incorrect answers. For a design-based solution to both concerns, the next section presents replicates this analysis using the revealed belief measure.

Revealed Disbelief in Incorrect Answers

To address the two concerns just raised about the trustworthiness of response stability as a metric, this section presents results using the revealed belief measure included in the March 2020 MTurk survey. As described in Chapter 4 and Appendix E.2, this measure reveals respondents' probabilistic beliefs through a series of costly discrete choices between payment for a correct answer and entering a drawing for a financial reward. Because these responses are financially incentivized, respondents have a greater incentive to reveal their beliefs accurately, which addresses concerns about expressive responding (Bullock et al. 2015; Prior et al. 2015). Because the two distinct measures of belief are taken in the same survey, there is no concern that learning between the two surveys is the source of response stability.[7]

An added advantage of the revealed belief measure is the relative credibility of taking a second measure of belief on answers to questions about the state of the economy. Response stability across time is harder to expect because there is not just a danger, but an expectation that respondents' beliefs should evolve with the state of the economy (Parker-Stephen 2013; Stanig 2013). Consequently, this section adds two economic facts that were convenient for Republicans at the time of the March 2020 MTurk survey, and two that were convenient for Democrats. These are:

- **Budget deficit.** Respondents read a short definition of the federal budget deficit. Respondents were then asked, "Compared with the 2017 fiscal year, was 2019's budget deficit higher or lower?" The rising deficit was interpreted as congenial to Democrats and inconvenient for Republicans, who controlled the presidency and had predicted that their signature tax cut legislation, the Tax Cuts and Jobs Act (TCJA), would reduce the deficit.
- **GDP growth.** Respondents read a short definition that linked the change in gross domestic product (GDP) to economic growth. Respondents were then asked, "Over the past year, what was the rate of economic growth in the United States?," with the options to say "Below 4%" or "4% or more." Growth below 4 percent was interpreted as congenial to Democrats and inconvenient for Republicans, who controlled the presidency and had prominently predicted this rate of growth after the passage of the TCJA.

[7] As discussed in Appendix E.3, the survey analyzed in this section also builds in a check for looking up the answer during the survey. It estimates that for the partisan questions, undetected cases of information search represent less than 1 percent of the data.

- **Inflation.** Respondents read a short definition of inflation and were told its historical average since 1945. Respondents were then asked, "Over the past year, has inflation been higher or lower than the historical average?" Below-average inflation was interpreted as convenient for Republicans.
- **Unemployment.** Respondents read a short definition of the U-3 unemployment rate as it is defined by the Bureau of Labor Statistics. Respondents were then asked, "Over the past year, did the unemployment rate increase or decrease?" Decreasing unemployment was interpreted as congenial for Republicans and inconvenient for Democrats, who were marketing the low rate as President Trump's signature achievement at the time of the survey.

Responses to these questions tend to be stated with less confidence than responses to the political awareness or controversy questions. The average respondent who answered correctly chose a confidence level of 0.737 on the 0.5 to 1 scale, while the average respondent who answered incorrectly chose a confidence level of 0.683. As before, respondents who answered correctly were about 30 percent more certain.[8]

Analysis

To accommodate the larger number of questions considered here, the graphical analysis splits the questions into two figures, according to whether incorrect answer was congenial to Democrats (Figure 5.3) or to Republicans (Figure 5.4). The interpretation of this figure matches that of Figure 5.2, with the exception that revealed belief has been substituted for belief stability. The two questions that were included in the belief stability analysis appear under the "controversies" heading, while the two new economic questions appear under the "economic" heading.

The leftmost panels in both figures use political awareness questions as a benchmark for the others. These panels are based on the same data as the rightmost panel of Figure 4.3, but split by whether the answer was correct and, in the bottom two rows, by the respondent's partisanship. Similar to the belief stability results, the figure shows that for respondents who answer correctly, the certainty measure does a fairly good job of covering the full spectrum from knowledge to near-complete ignorance: the average respondent who initially claimed to be certain of the correct answer reveal a belief of 0.942 in that answer, while the average respondent who claimed to be totally uncertain revealed a belief of 0.591. By contrast, incorrect answers cover the spectrum from ignorance to a miseducated guess: the average respondent who initially claimed to be certain of the incorrect answer revealed an average belief of 0.752 in it, compared with 0.519 for respondents who claimed to be totally uncertain.

In terms of the incorrect knowledge, miseducated guess, and total ignorance hypotheses, the results for the controversy questions are quite similar to the belief stability results, with one crucial

[8] (0.737 - 0.683) / (0.683 - 0.5) = 0.296

Figure 5.3: Revealed belief by certainty level, incorrect answer congenial to Democrats.

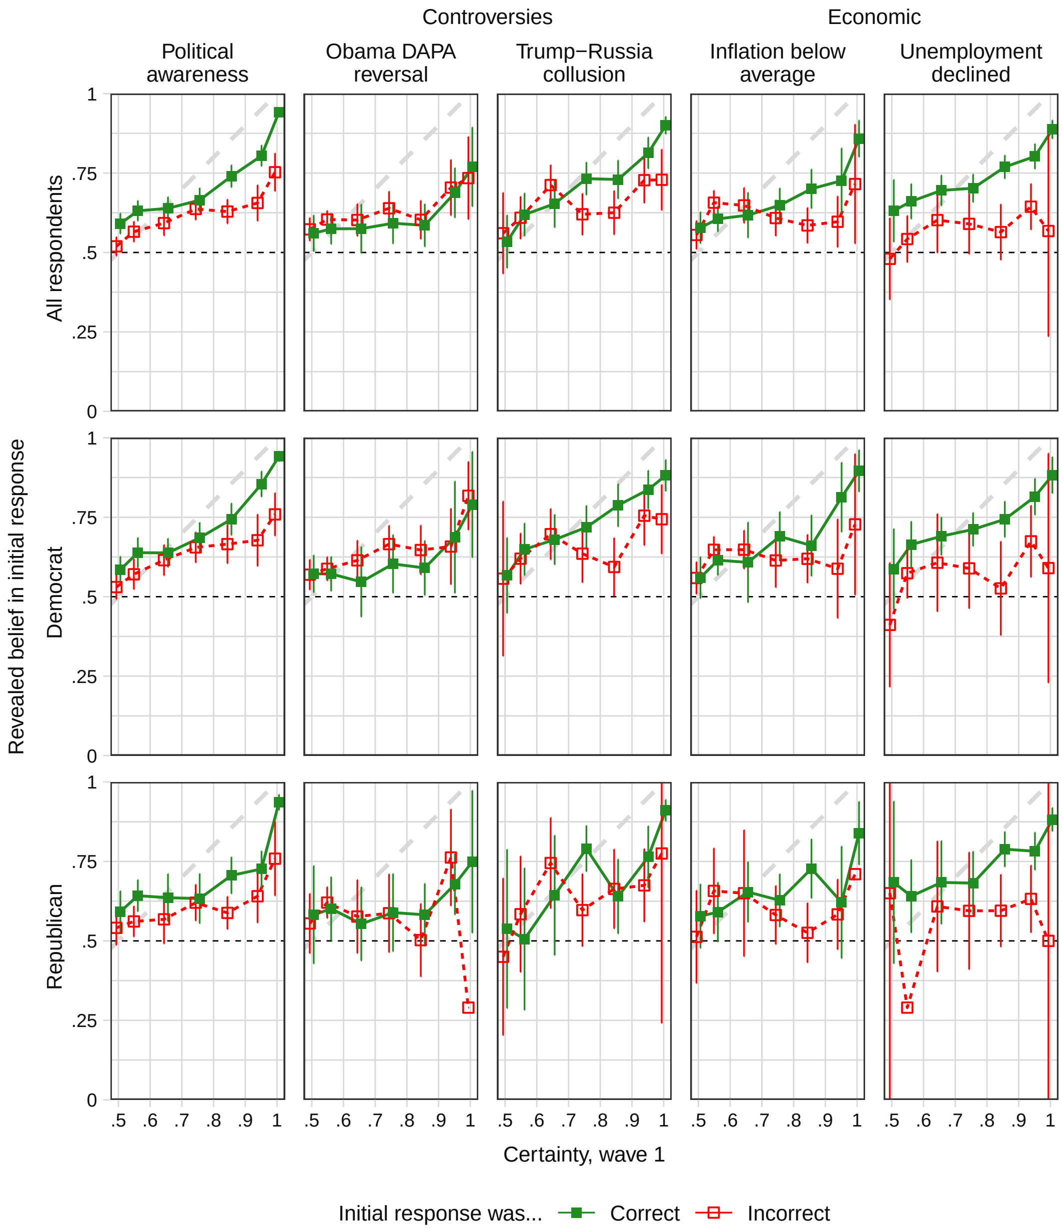

Note: This follows the format of Figure 5.2, with revealed belief on the y-axis rather than belief stability. Appendix C presents all of these estimates in tabular form.

Figure 5.4: Revealed belief by certainty level, incorrect answer congenial to Republicans.

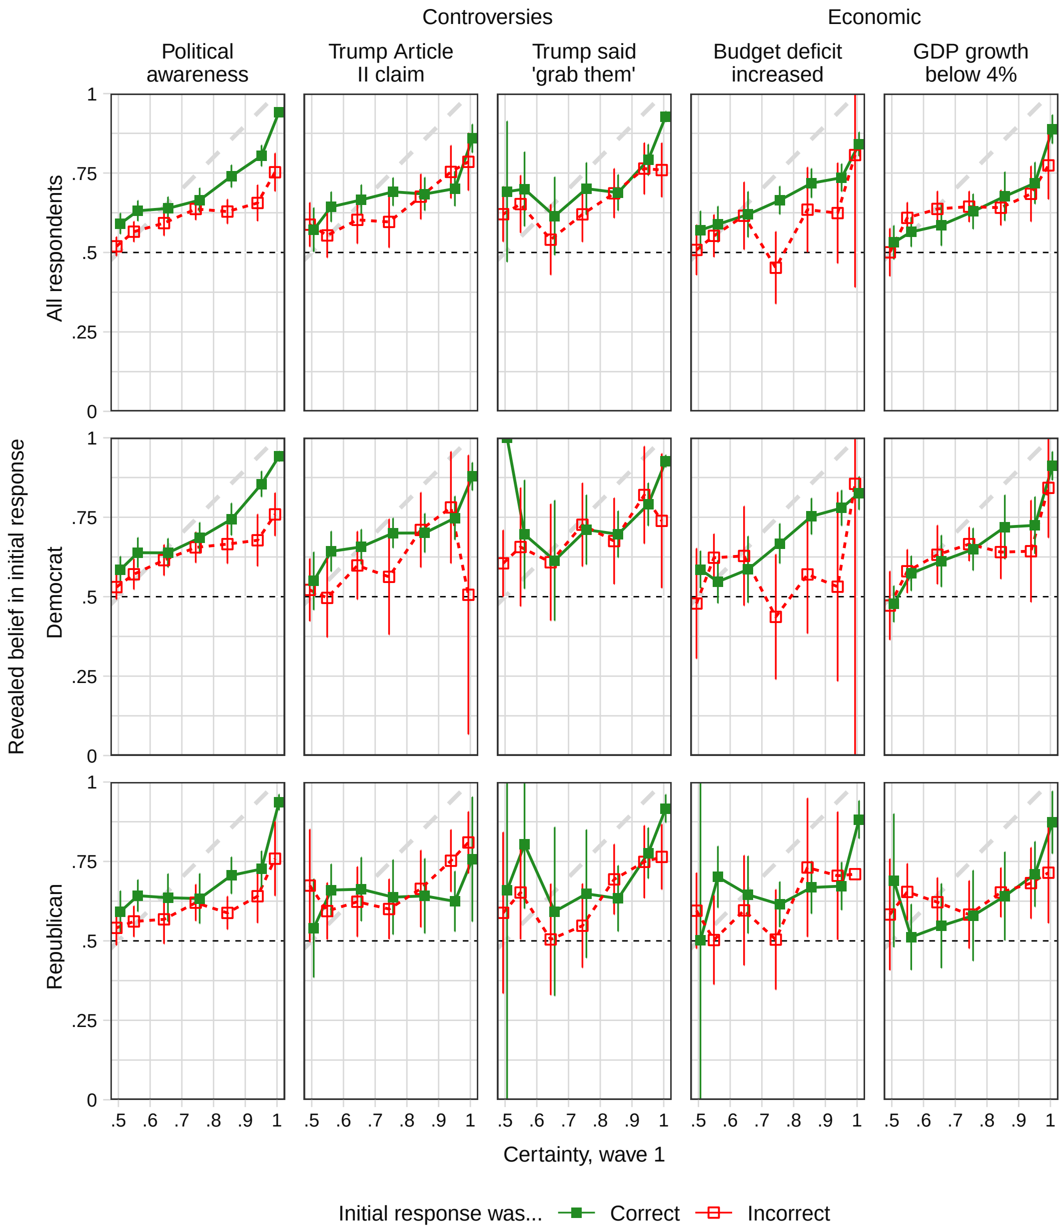

Note: This figure follows the format Figure 5.2, with revealed belief on the y-axis rather than belief stability. Appendix C presents all of these estimates in tabular form.

exception: there appears to be less evidence of a partisan asymmetry. Among both Democrats and Republicans, correct answers to the Trump 'grab them' and Trump-Russia questions cover close to the full spectrum from total ignorance to knowledge. Respondents who initially answered incorrectly display about an equal tendency to stick with their initial inference, except among those who initially indicated absolute certainty. Meanwhile, the Obama DAPA and Trump Article II questions see a fairly close tracking between correct and incorrect answers across the board. Especially on the Obama DAPA question, even the typical respondent who claims to be certain of the correct answer appears to be making an educated guess.

The economic questions display a similar pattern. Among all respondents, the average respondent who claimed to be certain of the incorrect answer revealed a belief of about 0.6 to 0.8 in their initial response, depending on the question. For the party with a partisan incentive to endorse the incorrect answer — Democrats on the two economic facts that are convenient for the incumbent Republican president in Figure 5.3, and Republicans on the two facts that are inconvenient in Figure 5.4 — the point estimate in every case is lower than the 0.752 observed on the political awareness questions. By comparison, claims to be certain of the correct answer fall in the 0.85 to 0.9 range. This falls short of the political awareness benchmark, but not by the same magnitude as incorrect answers.

For summary tests, I turn to the same OLS regression strategy developed in the previous section. Tables 5.3 and 5.4 display results from the same two sets of regressions, but with revealed belief substituted for belief stability. On the political awareness benchmark, we see the same pattern: claims to be certain of incorrect answers capture some meaningful variation in probabilistic belief, while correct answers capture about twice as much (Table 5.3). As was the case in in Chapter 4, the slope for revealed belief is not quite as steep as it is for response stability (Table 4.6).

For the questions whose incorrect answer was congenial to Democrats, the revealed belief measure provides some evidence that Democrats' claims to be certain capture meaningful variation in probabilistic beliefs (Table 5.4a). The estimate of β_1 implies that a one-unit increase in certainty implies about a one-quarter-unit increase in revealed belief in one's incorrect answer. By contrast, there is no statistical evidence that Republicans' claims to be certain about incorrect answers capture variation in probabilistic beliefs on these questions. The estimate of β_1 is small and statistically insignificant. For the partisan difference in interpretation hypothesis, this comparison is somewhat inconclusive. The estimate of β_4 provides suggestive evidence that the certainty scale is better at capturing variation in Democrats' belief in their answer than it is for Republicans, but the difference does not attain conventional levels of statistical significance.

Table 5.3: Regression test: revealed belief, political awareness questions.

Term	All respondents	Democrats	Republicans	Comparison
α Constant	0.364^{**}	0.351^{**}	0.402^{**}	0.351^{**}
	(0.032)	(0.043)	(0.054)	(0.043)
β_1 Certainty	0.337^{**}	0.381^{**}	0.264^{**}	0.381^{**}
	(0.045)	(0.060)	(0.075)	(0.060)
β_2 Correct	-0.184^{**}	-0.180^{**}	-0.196^{**}	-0.180^{**}
	(0.040)	(0.053)	(0.066)	(0.053)
β_3 Certainty $\times$ Correct	0.388^{**}	0.366^{**}	0.401^{**}	0.366^{**}
	(0.052)	(0.069)	(0.086)	(0.069)
α_R Republican				0.051
				(0.069)
β_4 Certainty $\times$ Rep.				-0.117
				(0.096)
β_5 Republican $\times$ Correct				-0.016
				(0.085)
β_6 Certainty $\times$ Corr. $\times$ Rep.				0.035
				(0.110)
Adj. R^2	0.283	0.310	0.205	0.271
Num. obs.	3417	1850	1205	3055

Note: This table presents OLS estimates of the parameters in (5.1) and (5.2), with revealed belief substituted for response stability. Standard errors clustered by respondent. $^{*}p < 0.05$, $^{**}p < 0.01$.

The most striking difference between the belief stability and revealed belief results comes in Table 5.4b, which presents the estimates for questions whose incorrect answer was congenial to Republicans. In contrast to the belief stability results in Table 5.2b, there is no evidence for the partisan difference in interpretation hypothesis on these questions. The estimates of β_1 are very similar for both parties, and the near-zero estimate of β_4 provides no evidence of a partisan difference.

Takeaways

Relative to the takeaways from the previous section, the results just presented should revise the reader's view of the results in two ways. First, the results in this section provide a bit more evidence in favor of the miseducated guess hypothesis. Regardless of party or question type, claims to be certain and uncertain about incorrect answers usually capture meaningful variation in revealed belief. However, claims to be certain of incorrect answers never come close to attaining the benchmarks set by the political awareness questions.

Second, they provide some reason to doubt the partisan difference in interpretations hypothesis. In the belief stability results, the key evidence for this hypothesis came on questions for which the incorrect answer is congenial to Republicans: conditional on certainty, Republicans were also more

Table 5.4: Regression test: revealed belief, partisan questions.

(a) Incorrect answer congenial to Democrats

Term	All respondents	Democrats	Republicans	Comparison
α Constant	0.497^{**}	0.465^{**}	0.527^{**}	0.465^{**}
	(0.028)	(0.034)	(0.059)	(0.034)
β_1 Certainty	0.168^{**}	0.223^{**}	0.096	0.223^{**}
	(0.041)	(0.051)	(0.082)	(0.051)
β_2 Correct	-0.232^{**}	-0.187^{**}	-0.285^{**}	-0.187^{**}
	(0.035)	(0.043)	(0.076)	(0.043)
β_3 Certainty $\times$ Correct	0.398^{**}	0.336^{**}	0.479^{**}	0.336^{**}
	(0.048)	(0.060)	(0.099)	(0.060)
α_R Republican				0.062
				(0.068)
β_4 Certainty $\times$ Rep.				-0.127
				(0.097)
β_5 Republican $\times$ Correct				-0.098
				(0.088)
β_6 Certainty $\times$ Corr. $\times$ Rep.				0.143
				(0.116)
Adj. R^2	0.124	0.123	0.118	0.120
Num. obs.	3496	1882	1241	3123

(b) Incorrect answer congenial to Republicans

Term	All respondents	Democrats	Republicans	Comparison
α Constant	0.352^{**}	0.352^{**}	0.360^{**}	0.352^{**}
	(0.035)	(0.055)	(0.060)	(0.055)
β_1 Certainty	0.375^{**}	0.365^{**}	0.367^{**}	0.365^{**}
	(0.047)	(0.076)	(0.075)	(0.076)
β_2 Correct	-0.149^{**}	-0.189^{**}	-0.068	-0.189^{**}
	(0.042)	(0.063)	(0.085)	(0.063)
β_3 Certainty $\times$ Correct	0.258^{**}	0.329^{**}	0.116	0.329^{**}
	(0.055)	(0.082)	(0.104)	(0.082)
α_R Republican				0.008
				(0.081)
β_4 Certainty $\times$ Rep.				0.002
				(0.106)
β_5 Republican $\times$ Correct				0.121
				(0.106)
β_6 Certainty $\times$ Corr. $\times$ Rep.				-0.213
				(0.133)
Adj. R^2	0.151	0.209	0.061	0.150
Num. obs.	3478	1888	1223	3111

Note: This table presents OLS estimates of the parameters in (5.1) and (5.2), with revealed belief substituted for response stability. Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

stable in their responses to these questions. The lack of such a difference using the revealed belief measure suggests that expressive responding, or differences in information exposure between the two surveys, may be a better explanation for the appearance that claims to be certain of incorrect answers mean different things for different parties. By extension, this implies that absent incentives to accurately reveal one's certainty level, expressive tendencies may artificially inflate the stability of responses that are congenial to the respondent's partisanship or other attitudes.

Evidence from Subjective Scales

The results presented so far in this chapter raise substantial concern about the "incorrect knowledge" interpretation of incorrect answers to factual questions. A reasonable reader could wonder if confusion about the meaning of numerical probabilities was responsible for the results. Most political science research that measures certainty or confidence tends to make use of subjective scales, by which I mean ordered verbal scale points that do not require respondents to think about numbers at all. Could some difference between numerical and subjective scales be responsible for the results in the previous section? Though it is not immediately clear what such flaw would uniquely affect incorrect answers, a core ethos of this volume is that when possible, questions about interpretation should be resolved through measurement.

Accordingly, this section replicates the findings from the previous sections using three sets of data that use four separate subjective certainty scales. The first dataset comes from the first investigation of respondents' certainty about their answers of which I am aware: the certainty scales included in the ANES during the 1990s and early 2000s (Alvarez and Franklin 1994). The second is a horse race between horse race between two scales directly copied from previously published work (Graham 2020; Pasek et al. 2015). The third uses data from the 2018 CCES. Each of these datasets features a certainty scale that I had no hand in developing.

A downside of subjective scales is that they lack a natural probabilistic interpretation. This means that they are not as easily converted into a measure of belief stability, which as discussed in Chapter 4, has easily interpreted expectations: if there were no measurement error, respondents would express the same degree of belief both times. Consequently, the analysis in this section focuses on response stability, which was defined in Chapter 4 as the percentage of respondents who choose the same answer when asked the question again in a follow-up survey. The analysis will be informative as to the difference between correct and incorrect answers, but the benchmarks that help adjudicate between the three hypotheses — incorrect knowledge, miseducated guesses, or total ignorance — will not be as clear. Fortunately, the results are stark enough as to leave little question as to the relative dependability of claims to be certain of correct or incorrect answers.

Bill Clinton's Ideology

It is only fitting that the first evidence from subjective scales come from the first large-scale investigation of respondents' certainty about their answers to be performed by political scientists. As described in more detail in Chapter 3, more than 70 questions that were asked in 1993-2000 ANES surveys included a three-point certainty scales with the scale points "not very certain," "pretty certain," and "very certain." Follow the same scoring rule used in Chapter 3, I code these levels as 0, 0.5, and 1 in regression analysis.

Though most of the ANES questions with certainty scales attached were about subjective matters like the respondent's own policy positions and character traits of politicians, one of these questions in the 1994 post-election survey was about something that is arguably factual: former president Bill Clinton's ideology. After placing themselves on a seven-point political ideology scale, respondents were asked, "Where would you place Bill Clinton on this ruler?" Respondents who classified Clinton as liberal were considered to have answered correctly, while respondents who classified him otherwise were considered to have answered incorrectly.[9] This question was asked again in both the pre- and post-election waves of the 1996 ANES.

By this definition of a correct and incorrect answer, correct answers were provided by 68 percent of respondents to the 1994 post-election survey and 64 percent of respondents to the 1996 pre- and post-election surveys. Among respondents who answered correctly, 25 percent indicated that they were not at all certain, 51 percent pretty certain, and 24 percent very certain (average = 0.49). Among respondents who answered incorrectly, this split was 47 percent, 40 percent, 13 percent (average = 0.33).

Figure 5.5 displays response stability conditional on whether the respondent answered correctly or incorrectly. Among respondents who answered correctly, response stability increases markedly in certainty, from 68 percent among the not very certain to 91 percent among the very certain. Among respondents who answered incorrectly, the same increase is from 63 percent to 67 percent. The high level of stability among respondents who are certain and correct indicates that these respondents likely had a well-formed impression of Clinton as a liberal.

[9]Though Clinton was clearly on the leftward half of the ideological spectrum as it existed in the 1990s, classifying "moderate" as an incorrect answer admittedly does some injustice to Clinton's position on a few key issues, like welfare reform. To the extent that I have erred in scoring respondents who classified Clinton as a moderate as having answered incorrectly, it should work against my results.

Figure 5.5: Response stability by certainty level, 1994-1996 ANES.

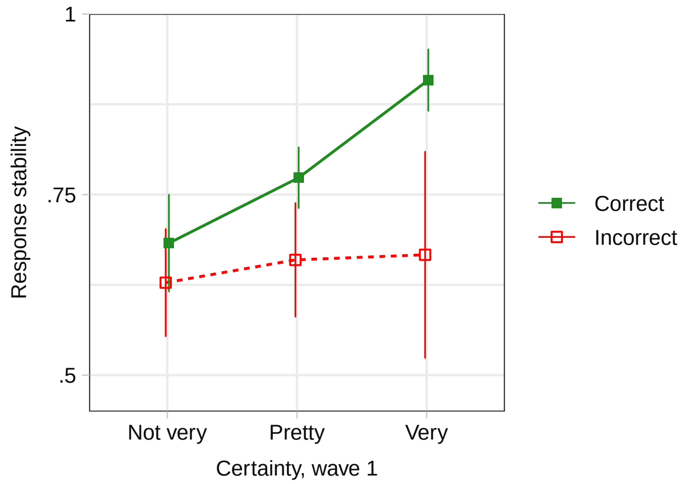

Note: This figure plots response stability (y-axis) by certainty level (x-axis) for the 1994-96 ANES question about Bill Clinton's ideology. It follows the format of the figures above.

To summarize the difference between correct and incorrect answers, I used OLS to estimate the parameters in (5.1). The regression test finds the same difference between correct and incorrect answers that is observed above. In Table 5.5, the estimate of β_1 shows that among respondents who answered incorrectly, certainty predicts only a slight increase in response stability, which is not statistically distinguishable from zero. By contrast, the estimate of β_3 shows that correct answers are about 18 percentage points more stable.

At first blush, the results of these tests could appear consistent with the total ignorance hypothesis: among respondents who answer incorrectly, certainty does not appear to predict much responsibility at all. However, if respondents chose between the seven response options at random, stability would be $3/7 = 0.429$, which is well below the stability observed at any certainty level. Moreover, because DK responses were allowed, many of the least certain — and likely least stable — responses were filtered out. Consequently, although these tests help confirm that the difference between correct and incorrect answers observed above is not an artifact of the particular samples or measurement choices, the lack of a clear floor on performance makes it hard to say whether the results are more supportive of the miseducated guess hypothesis or the total ignorance hypothesis.

Foreign Aid Spending

The second test using subjective scales comes from the panel dataset analyzed by Graham and Svolik (2020). As part of a battery of political knowledge questions taken directly from the 2016 ANES, respondents were asked, "On which of the following does the U.S. federal government

Table 5.5: Regression test: response stability, 1994-96 ANES

Term	Estimate
α Constant	0.631**
	(0.036)
β_1 Certainty	0.046
	(0.074)
β_2 Correct	0.041
	(0.046)
β_3 Certainty × correct	0.180*
	(0.084)
Adj. R^2	0.041
Num. obs.	1086

Note: This table presents OLS estimates of the parameters in (5.1). Robust standard errors in parentheses. $^{*}p < 0.05, ^{**}p < 0.01$.

currently spend the least?" and presented with four response options in random order: foreign aid, Medicare, national defense, or Social Security. The first option, foreign aid, is correct. After one to three weeks, respondents answered the same question in a follow-up survey, with no certainty scale.

When they first answered the question in the wave 1 survey, respondents were randomly assigned to use one of two certainty scales: the certainty scale used by Graham (2020) and the certainty scale used by Pasek, Sood and Krosnick (2015). The Graham (2020) scale asked respondents, "How certain are you that your answer is correct?" and allowed respondents to choose one of five options ranging from "don't know" to "absolutely certain." The Pasek et al. (2015) scale asked respondents, "How sure are you about that?" and allowed respondents to choose one of five options ranging from "Not sure at all" to "Extremely sure." I originally conducted this test to provide some exploratory evidence as to the differences in measurement properties between the two scales. For the present purposes, the causal nature of the estimated difference between the two scales is secondary; what matters is whether the basic pattern of results from above replicates with a subjective scale.

In the first wave, 28.4 percent of respondents answered correctly, including 27.3 percent of those assigned to use the Graham scale and 29.5 percent of those assigned to use Pasek and colleagues' scale. Average certainty was 3.03 and 2.83 among respondents who answered correctly, and 2.92 and 2.74 among respondents who answered incorrectly.[10] In the second wave, 29.1 percent of respondents answered correctly, suggesting that learning was probably not a major cause of response instability.

Figure 5.6 displays response stability conditional on certainty. Regardless of which scale is used, correct answers are substantially more stable than incorrect answers at high levels of certainty.

[10] These differences in average certainty are statistically significant in t-tests, with t-statistics of 3.4 for all responses, 2.9 for correct answers, and 1.9 for incorrect answers.

Figure 5.6: Response stability by certainty level, foreign aid.

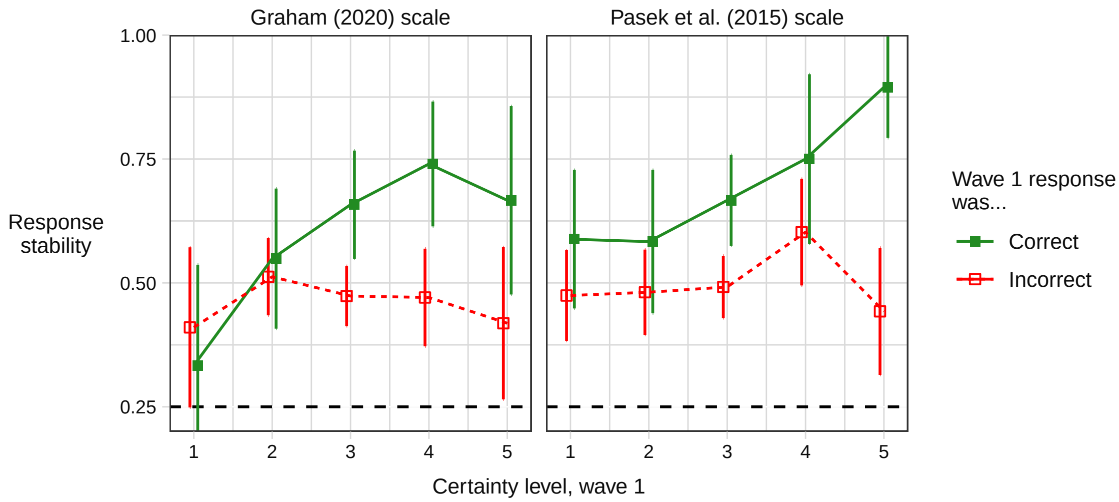

Note: This figure plots response stability (y-axis) by certainty level (x-axis) for the foreign aid question. It otherwise follows the format of the figures above.

Among respondents who answered correctly, the Graham (2020) scale seems better at identifying respondents who are truly uncertain, while the Pasek et al. (2015) scale seems better at identifying respondents who actually know that foreign aid is a smaller budget item than the other three.

Because respondents to these questions were not offered a DK response option, there is a clear floor for response stability: if respondents were choosing completely at random, they would choose the same response option 25 percent of the time. Incorrect answers consistently sit above this floor, regardless of the respondent's certainty level. This suggests that incorrect answers reflect some meaningful probabilistic belief, but that these certainty scales do not do a very good job of picking up variation in it.

To summarize the difference between these scales, I adapt the regression-based strategy for testing for partisan differences. Table 5.6 displays the results. The estimates of β_1 indicate that there is no statistical evidence that claims to be certain of incorrect answers predict greater response stability using either measure. The estimates of β_4 indicate that on correct answers, certainty is reasonably predictive of response stability. The lack of a difference in slopes is probably a little deceptive, given the apparent performance difference among respondents who indicate low and high levels of certainty. As discussed in Chapter 4, over-reliance on regression can hide important features of data.

Table 5.6: Regression test: foreign aid

Term	All responses	Randomly assigned scale: Graham	Pasek et al.	Comparison
Constant	0.484**	0.496**	0.478**	0.496**
	(0.028)	(0.044)	(0.035)	(0.044)
Correct	0.019	−0.049	0.061	−0.049
	(0.049)	(0.079)	(0.063)	(0.079)
Certainty (rescaled [0,1])	0.005	−0.043	0.042	−0.043
	(0.052)	(0.082)	(0.067)	(0.082)
Correct × Certainty	0.304**	0.383**	0.261*	0.383**
	(0.085)	(0.138)	(0.106)	(0.138)
Pasek et al. (2015) scale				−0.018
				(0.057)
Correct × Pasek et al.				0.110
				(0.101)
Certainty × Pasek et al.				0.085
				(0.106)
Correct × Certainty × Pasek et al.				−0.121
				(0.174)
Adj. R^2	0.031	0.024	0.036	0.031
Num. obs.	1749	838	911	1749

Note: This table presents OLS estimates of the parameters in (5.1) and (5.2). Standard errors clustered by respondent. $^{*}p < 0.05$, $^{**}p < 0.01$.

Government Policy

The third set of evidence comes from another set of data that I did not collect. In the 2018 Cooperative Congressional Election Survey (CCES), the Syracuse team module includes questions designed to tap misconceptions about status quo government policy in both the pre- and post-election surveys.[11] The questions asked respondents which of two statements was more likely to be true. For each pair of statements, brackets enclose the differences between the two and the correct version is printed in boldface:

- Currently, there is [**a** / no] federal limit on how long a person can receive welfare (TANF) benefits.[12]
- China holds [more / **less**] than half of US debt.[13]
- Interest on the federal debt is [more / **less**] than half of federal spending.[14]
- Undocumented immigrants [are / **are not**] eligible to receive food stamps.[15]

[11] One item from the pre-election survey, "[More / Less] than 30% of U.S. citizens pay no federal income tax.", was not included in the post-election survey.

[12] "Policy Basics: Temporary Assistance for Needy Families," Center on Budget and Policy Priorities, February 6, 2020.

[13] "No, China does not hold more than 50 percent of U.S. debt," *The Washington Post*, December 29, 2014.

[14] "5 facts about the national debt," Pew Research Center, July 24, 2019.

[15] Unauthorized immigrants are not themselves eligible to receive Supplemental Nutrition Assistance Program

- In the U.S., corporations are taxed at a [**higher** / lower] rate than in most other developed countries.[16]
- Federal law [**does** / does not] require licensed gun dealers to conduct a background check before selling a gun to a customer.[17]
- Planned Parenthood receives federal funding and [is allowed / **is not allowed**] to use it to provide abortion services.[18]

After answering each question, respondents stated their certainty level on a three-point scale of "not confident," "somewhat confident," and "very confident." For analysis this scale is recoded to $\{0, 0.5, 1\}$, as in the ANES analysis above and in Chapter 3.

Among the three studies using subjective scales, this battery of questions poses the toughest test, and not only because it has the most questions. Figure 5.7 plots the distribution of certainty for each question in the study. Relative to the distributions plotted above in Figure 5.1, respondents to these questions were more similarly confident in their correct and incorrect answers. Pooling across the questions, just 51 percent of respondents answered the two-option questions correctly. On the 0 to 1 scale, respondents who answered correctly were about 15 percent more confident in their correct answers than in their incorrect answers — a substantially smaller gap than appeared on the controversy questions. For two of the questions, the average incorrect answer is stated with greater confidence than the average incorrect answer.

Figure 5.8 plots average response stability for correct and incorrect answers at each level of certainty. Respondents who claim to be certain of the correct answer are stable enough in their responses that they might, in a rough sense, be said to know the answer: 89 percent of respondent who said they were very confident of the correct answer in October chose the same response again in November. Among the correct and not confident, this figure was 64 percent. By contrasts, claims to be certain of incorrect answers are much less indicative of incorrect knowledge: response stability ranges from 51 percent among the not confident to 61 percent among the very confident. Even the allegedly most confident incorrect answers are less stable on average than the allegedly least confident incorrect answers.

The unusually large gap between completely certain and completely uncertain responses may owe in part to the coarseness of the scale. Although it is possible to have too many scale points,

(SNAP) benefits, but households headed by unauthorized immigrants can access food stamps if the beneficiary is a U.S. citizen child.

[16] This was the hardest claim to assess. Prior to the December 2017 passage of the Tax Cuts and Jobs Act, the U.S. had the highest statutory corporate tax rate and fourth-highest effective corporate tax rate among G20 countries (Source: Congressional Budget Office). The Act lowered the rate substantially, but the question was asked before the end of the first tax year under the law. Fortunately, the unusually close correspondence in response stability between correct and incorrect answers on this question means that this scoring decision has little effect on the results.

[17] "Universal Background Checks," Giffords Law Center, Accessed September 25, 2020.

[18] Federal funding for abortions was banned by the Hyde Amendment in 1976.

Figure 5.7: Distribution of wave 1 certainty, Study 2.

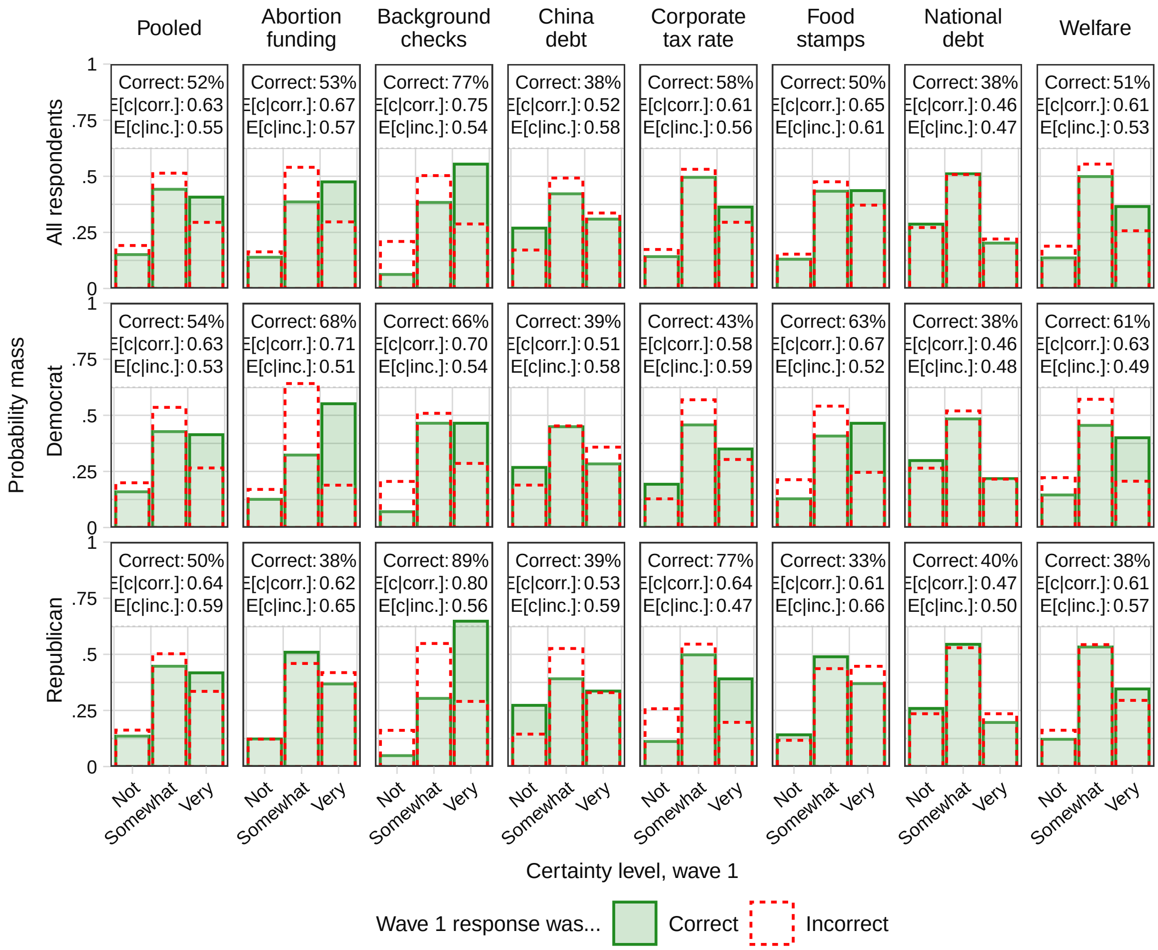

Note: This figure plots the distribution of certainty for the government policy questions. The statistics printed in each panel have the same meaning as those in Figure 5.1.

Figure 5.8: Response stability by certainty level, Study 2.

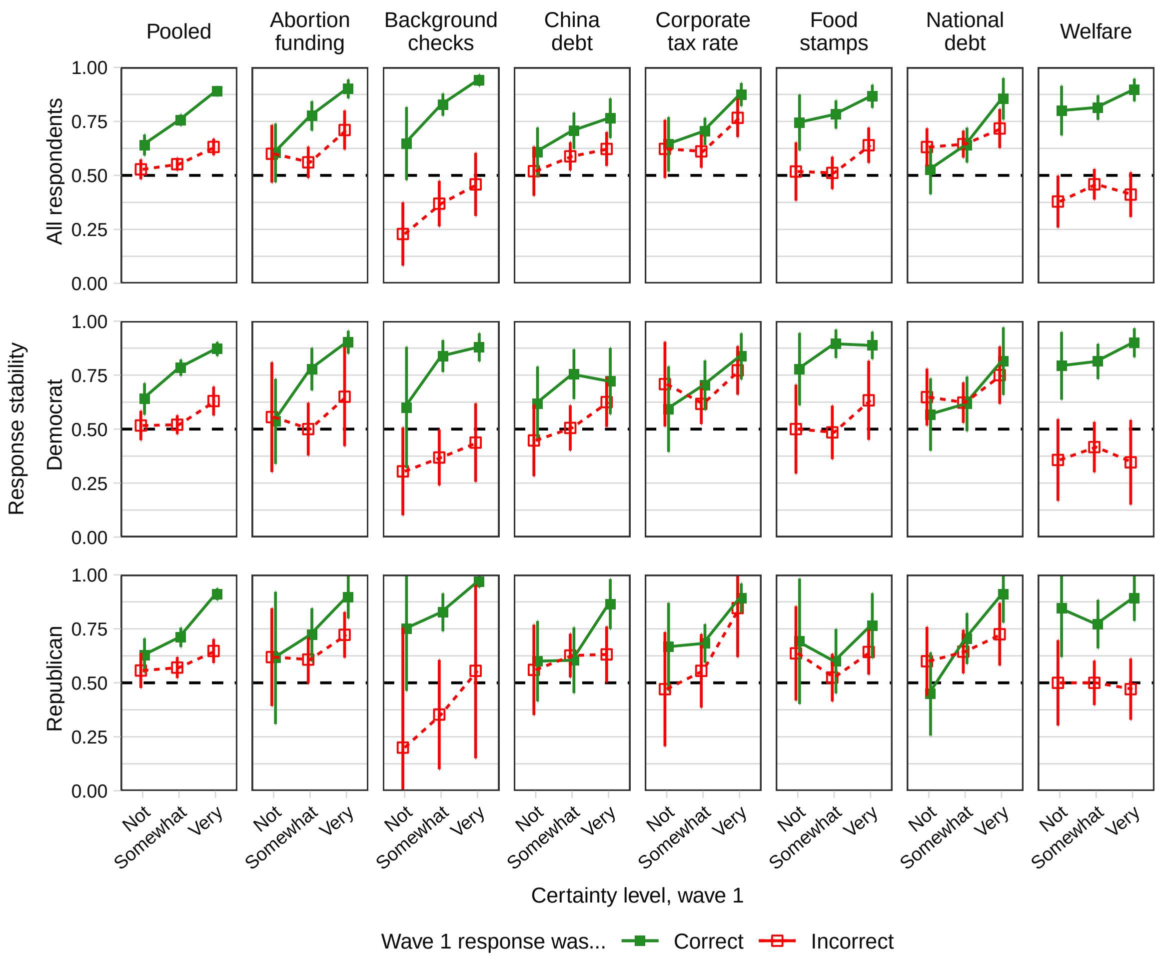

Note: This figure plots response stability (y-axis) by certainty level (x-axis) for the government policy questions. It otherwise follows the format of the figures above.

three levels does not allow much space to separate the very most certain respondents from very least. Figures 5.2, 5.3, and 5.4, which use seven bins, suggest that finer scales capture a wide range of variation in stability. If respondents who are correct and incorrect differ systematically within the three bins, differences in stability could crop up as a result.

The regression test reflects these differences as well. Table 5.7 presents the results of the same regression-based test that has been used throughout. The coefficient estimate for certainty suggests that Democrats and Republicans who claim to be more certain do, on average, assign greater credence to their answers. On average, a full scale increase in certainty predicts about an 11 percentage point increase in response stability. By contrast, correct answers are about 12 percentage points more stable on average at baseline, and a full scale increase predicts a 26 percentage point

Table 5.7: Regression test: response stability, policy status quo

Term	All respondents	Democrats	Republicans	Comparison
Constant	0.510^{**}	0.484^{**}	0.534^{**}	0.484^{**}
	(0.020)	(0.031)	(0.034)	(0.031)
Correct	0.124^{**}	0.176^{**}	0.047	0.176^{**}
	(0.028)	(0.043)	(0.047)	(0.043)
Certainty	0.110^{**}	0.121^{*}	0.102^{*}	0.121^{*}
	(0.030)	(0.050)	(0.048)	(0.050)
Certainty × correct	0.143^{**}	0.098	0.213^{**}	0.098
	(0.039)	(0.063)	(0.061)	(0.063)
Republican				0.050
				(0.046)
Correct × Republican				-0.129^{*}
				(0.064)
Certainty × Republican				-0.019
				(0.069)
Certainty × corr. × Rep.				0.115
				(0.087)
Adj. R^2	0.079	0.089	0.071	0.081
Num. obs.	10078	4594	3924	8518

Note: This figure presents OLS estimates of the parameters in (5.1) and (5.2). Standard errors clustered by respondent. $^{*}p < 0.05, ^{**}p < 0.01$.

increase in response stability.

Implications

The results above demonstrate that measurement error frequently creates the illusion that misperceptions are more deeply-held than they really are. Even the "24 carat gold standard" for measuring misperceptions — focusing on respondents who claim to be absolutely certain of the incorrect answer — fails to identify respondents who think they know, but are wrong. Given that most survey-based analysis of political misperceptions does not go nearly this far in trying to isolate respondents who are genuinely certain of their answer, strong claims about the prevalence of misperceptions based on survey data are very likely to be exaggerations.

At the same time, the results demonstrate that incorrect answers to factual survey questions are not completely meaningless. Respondents who claim to be more certain of incorrect answers are more stable in their responses and reveal a higher degree of belief in them when financial incentives are involved. Thinking of incorrect answers as miseducated guesses captures this reality. Respondents who claim to believe their answers seem to be drawing on some consistent basis for their inference. These are meaningful beliefs in the probabilistic sense, but not in the high-threshold sense that pervades authoritative definitions of misperceptions.

Paired with the findings on the properties of DK response options in Chapter 3, this suggests that DK response options help identify a set of incorrect responses that can be fairly thought of as miseducated guesses. Respondents who say that they are completely uncertain tend to be the least stable in their responses. Offering DK response option and scoring these responses as an 0.5 is quite justifiable. The unjustifiable practice is to score those who do not say DK as a 1, as if they completely believe their answer.

The chief limitation of the evidence presented here is that it did not consider every question — or even most of the questions, if I may lean on a vaguely-defined denominator — on which scholars and other observers of politics have claimed to have measured misperceptions. It is possible that for some or even many survey items, claims to be certain of incorrect answers are more meaningful than they were on the items analyzed here. The true strength of this chapter's framework is to show that such claims are provable. To the extent that scholars think that their favorite question or measurement technology captures misperceptions more effectively than the questions and measurement technologies deployed here, they can use these empirical strategies to test that hypothesis.

The major implication for future research, then, is that researchers who want to make stronger claims than the miseducated guess interpretation can and should assume the burden of proof for their interpretation. Independent of our theories about who should believe in false claims, measurement can provide a basis for verifying the extent to which people actually believe them. Though the results in this chapter suggest that identifying respondents who are certain of falsehoods is a taller task than many have assumed, it may not be impossible. My hope is that the methods developed here ultimately provide a toolkit for solving these measurement problems or at least discriminating between instances where they do and do not appear, as opposed to simply pointing them out.

Chapter 6

Partisan Belief Differences

Abstract. Prevailing practices in survey research treat all respondents as if they believe their answers with complete certainty. This chapter shows that the resulting bias — certainty bias — distorts the apparent size and nature of partisan differences in factual beliefs. On average, certainty bias increases measured belief differences between Democrats and Republicans by 40 percent. This varies substantially across questions, sometimes more than doubling the measured belief difference and other times cutting it in half. This variation systematically distorts portraits of partisan disagreement: certainty bias is substantially more pronounced on questions about economic performance than on questions about abuses of presidential power, partisan rumors, and other politicized controversies. After accounting for certainty bias, the typical partisan belief difference is better-characterized as a knowledge gap than as evidence of misperceptions. These findings demonstrate that accounting for respondents' uncertainty about their beliefs does more than just shrink measured belief differences between Democrats and Republicans. Instead, it sharpens the observer's sense of what partisan differences mean and where these differences are most pronounced.

Survey measures suggest that Democrats and Republicans disagree over matters of fact. Taken at face value, responses to questions about objective facts suggest perceptual differences in a wide range of dimensions, including the economy (Bartels 2002; Conover et al. 1986; Evans and Andersen 2006), war (Kull et al. 2003; Prasad et al. 2009; Jacobson 2010; but see Gaines et al. 2007), and politicized rumors and controversies (Krosnick et al. 2014; Berinsky 2018). As discussed in the introduction to this volume, these differences raise fundamental questions for democratic citizens' ability to hold politicians accountable for their performance and actions. To the extent that partisanship encourages people to perceive reality in a manner that is favorable to their party and unfavorable to the other party, partisan voters are may be less likely to hold their side accountable for bad performance or reward the other side for good performance.

The large majority of evidence of partisan differences in perceptions of reality is based on the sort of survey questions whose measurement properties were examined in the previous chapter. The canonical case is the state of the economy. For example, researchers ask respondents to guess the unemployment rate or the change in it over some time period. The defining advantage of such questions is that unlike subjective evaluations (e.g., is the economy getting better or worse?), respondents who answer the question as intended should only base their response on their perceptions of reality, not on attitudes that correlate with partisanship (e.g., concerning how the economy should be managed; Bartels 2002).

The essential drawback of factual questions is that not everyone who answers survey questions has a pre-existing belief about the answer. As discussed in Chapter 2, existing research highlights two reasons why survey measures of belief might not be an accurate reflect on of respondents' beliefs. First, expressive responding, or respondents' tendency to report beliefs that are more party-congenial than their underlying perceptions. Second, uncertainty, or respondents' tendency to guess even when they do not know the correct answer.

Accounts of expressive responding focus on respondents' motivations as they form and express their beliefs. A fully accuracy motivated respondent would read or listen to the survey question, take an unbiased sample of their preexisting perceptions, mold those perceptions into a belief, and report that belief honestly. A lack of accuracy motivation could distort this process in one of two ways: respondents may either form their belief based on a biased sample of their underlying perceptions (Kunda 1990; Jerit and Zhao 2020), or they may "cheerlead" by insincerely reporting a belief that is more party-congenial than the belief they actually formed (Bullock et al. 2015). To simultaneously address both possible mechanisms, researchers use accuracy incentives, which are thought to encourage both a more even-handed inferential process and greater honesty in reporting the resulting belief. Through some combination of even-handed sampling and honest reporting, this tightens the link between expressed beliefs and underlying perceptions.

Uncertainty is a concern because surveys do not usually allow respondents to express the uncertainty in their beliefs. Standard surveys record only an all-or-nothing judgment: respondents must declare their belief in one option or the other.[1] This practice is not neutral vis-a-vis respondents' uncertainty. To see this, suppose a respondent reads a question with two response options, 0 and 1, and forms a probabilistic belief p_i, corresponding to the probability assigned to option 1. If $p_i = 1$, the respondent is certain about the correct answer, meeting the definition of knowledge from Chapter 4. If $p_i = 0$, the respondent is certain about the incorrect answer. If p_i falls somewhere in between, the respondent is at least somewhat uncertain: either some of their perceptions conflict,

[1]Below, the empirical framework covers "don't know" responses.

or are not dispositive enough, to support a completely certain inference one way or the other. The researcher, never seeing p_i, requests and records a 0 or 1 — exactly as if the respondent were completely certain of their answer. This tends to exaggerate the partisan difference in beliefs. In an extreme case, if p_i were 0.4 for all Democrats and 0.6 for all Republicans, common practices could warp a true belief difference of 0.2 into a measured difference of 1.0.

Despite a litany of evidence that people guess the answers to survey questions even when they are uncertain, little is known about how individual-level uncertainty affects estimates of partisan differences in factual beliefs. Bullock and Lenz (2019) write:

> [T]he extent to which apparent partisan differences in factual beliefs depend on a lack of confidence in the correct answers ... is deeply underappreciated in research on partisan differences. As we have argued, partisan differences in survey responses will sometimes be due neither to cheerleading nor to sincere differences in beliefs, but to the use of pro-party responses among those who have little notion about the correct answers. (337-38)

To address this problem, this chapter applies the probabilistic belief framework to the question of partisan differences in factual beliefs. It finds that treating respondents as if they are completely certain of their answers substantially inflates estimates of partisan perceptual differences. The difference between the average difference in probabilistic beliefs and the measured difference when respondents are assumed to be certain of their answers is termed *certainty bias*.

Study 1, which examines 25 questions from four surveys conducted in 2018 and 2019, finds that certainty bias inflated partisan belief differences by nearly 40 percent on average, from 0.106 to 0.147 on a 0 to 1 probability scale. Differences on factual questions about the state of the economy were inflated by 55 percent (from 0.086 to 0.135), while differences over politicized controversies were inflated by a bit more than 25 percent (from 0.128 to 0.158). Using the same measure of revealed belief developed earlier in this volume, a second study uses the revealed belief measure to show that certainty bias is present even when expressive responding is accounted for.

Allowing respondents to express uncertainty about their answers provides a unique window into the degree to which partisan belief differences are driven by two possible departures from perfect knowledge: ignorance, defined above as a total lack of certainty about one's beliefs (Chapter 4), and misperceptions, which are conventionally defined as certainty about the incorrect answer (Chapter 5). The findings suggest that partisan differences are better-attributed to knowledge gaps than to misperceptions. Pooling across the two studies, the average correct answer is stated with about 30 percent greater confidence than the average incorrect answer. For example, among both Democrats and Republicans, respondents who correctly say that former President Barack Obama released his birth certificate are more certain about their answer than respondents who incorrectly say that he did not.

These results are generalizations, not universal rules. At the question level, there is substantial variation in certainty bias. At times, certainty bias is negative, meaning that estimates of partisan belief differences can *increase* once uncertainty is accounted for. Below, analysis of two factors that contribute to the magnitude of certainty bias — generalized uncertainty, and differences in certainty between Democrats and Republicans — sheds light on why it is that estimates of partisan differences over politicized controversies are less-affected by certainty bias than are differences over the state of the economy.

Theoretical Framework

Consider two hypothetical voters, a Democrat and a Republican. A survey seeks to measure their beliefs about the change in the unemployment rate over the past year. Before reading the question, neither knows the exact value of the unemployment rate or its change over time. Instead, they construct a belief about it on the spot. The Democrat's best guess that unemployment increased by 2 percent, while the Republican's best guess is a decline of the same magnitude. The partisan difference in best guesses is 4 percent.

To fully characterize the difference between these two individuals' beliefs, one also needs information about their uncertainty. Figure 6.1 plots two pairs of personal probability distributions — that is, the individuals' beliefs about unemployment — both of which are consistent with a 4 percent difference in best guesses. In the left panel, both voters have fairly certain beliefs, with no probability assigned to any of the same possibilities. In the right panel, both voters are quite uncertain, resulting in substantially more overlap: the Democrat and Republican assign almost all of their personal probability mass to the same possibilities. This greater similarity is not captured by measurement technologies that only ask for a best guess. A difference in best guesses clearly indicates some difference in beliefs, but is strikingly uninformative as to how different the Democrat's belief is from the Republican's.

Measuring beliefs by eliciting best guesses about numerical quantities is a common empirical strategy (e.g., Ahler and Sood 2018; Alesina et al. 2020; Ansolabehere et al. 2013; Bullock et al. 2015). Its complete insensitivity to individual-level uncertainty highlights a fundamental obstacle to incorporating such uncertainty into estimates of partisan belief differences: researchers tend to favor single-parameter measures of beliefs and estimators of belief differences, but uncertainty enters

Figure 6.1: Hypothetical individual-level beliefs about the change in unemployment.

most probability distributions as a second parameter. An empirical strategy's ability to incorporate individual-level uncertainty into estimates of partisan perceptual differences depends on the mapping between question formats, probability distributions, and estimators.

To streamline this mapping, this chapter focuses on "binary" questions, meaning questions that have two response options. Because binary questions map onto a binomial distribution, each respondent's entire personal probability distribution can be summarized by a single parameter p, the probability the respondent assigns to the correct answer. In turn, a difference in means estimate that reflects individual-level uncertainty can be computed. In Figure 6.1, each respondent's personal probability that unemployment decreased is represented by the amount of each distribution that falls to the left of zero; that it increased, to the right of zero.

The remainder of this section considers three strategies for estimating the difference between the average Democrat's belief and the average Republican's belief. Define the partisan belief difference as

$$\Delta = \frac{1}{N_D} \sum_{i \in \mathbb{D}} p_i - \frac{1}{N_R} \sum_{i \in \mathbb{R}} p_i, \tag{6.1}$$

where p_i is the probability that each respondent assigns to the correct answer, $\mathbb{D}$ and $\mathbb{R}$ are the sets of respondents who prefer the Democratic and Republican parties, and N_D and N_R are the total number of respondents in these sets. This can be estimated using the measures of probabilistic beliefs that were developed and tested in previous chapters.

Most surveys used to measure misperceptions and partisan belief differences ask respondents only for a best guess (Luskin et al. 2018). Such questions require the respondent to format their considerations into a 0 or a 1, corresponding to the answer they find most probable. To express this task as researchers ask respondents to carry it out, suppose that respondents use the function $\mathbb{1}(p_i > 0.5)$, where $\mathbb{1}$, the indicator function, returns a 1 if the condition is met and 0 otherwise. The *best guess estimator* of the partisan difference in beliefs is

$$\Delta_{\text{BG}} = \frac{1}{N_D} \sum_{i \in \mathbb{D}} \mathbb{1}(p_i > 0.5) - \frac{1}{N_R} \sum_{i \in \mathbb{R}} \mathbb{1}(p_i > 0.5). \tag{6.2}$$

Notice that the best guess estimator is not neutral vis-a-vis the respondent's certainty level. Instead, by scoring each respondent as a 0 or a 1, the best guess estimator treats all respondents as if they were completely certain. In recognition of this property, define *certainty bias* as the difference between the estimand (6.1) and the measured belief difference (6.2). More specifically, certainty bias is equal to $\Delta_{BG} - \Delta$.

Many researchers find fault with strategies that treat all survey responses, even those offered by respondents who had never thought about the question before, as if the respondent was completely certain. As described in Chapter 2, concern about this problem has consistently motivated political scientists to turn to threshold-based strategies. The usual solution is to allow respondents to opt out of providing a best guess if they are insufficiently certain, e.g., by saying DK. In showing that providing a DK response option can be thought of as a threshold-based strategy, Chapter 3 highlighted an important pitfall: allowing DK responses sets a low threshold, leaving substantial uncertainty among those who answer the question. To raise the bar, many recommend measuring certainty and reserving the term "believe" for respondents who choose the top two scale points on a four- or five-point certainty scale (Kuklinski et al. 2000; Pasek et al. 2015; Graham 2020). A similarly motivated device in attitude surveys is the "opinion filter," which raises the bar by encouraging respondents to say DK if they are not sure (Schuman and Presser 1981).

To formally express threshold strategies, recall that the previous chapters defined c_i as the probability the respondent assigns to their best guess. Certainty is a "folded" probability scale ranging from 0.5 to 1.[2] Threshold-based measurement technologies ask respondents to answer the question if their certainty level exceeds some latent threshold, τ, and to say "DK" otherwise.

Threshold-based estimators are deployed in two main ways. First, researchers may drop respondents who say DK and use the best guess estimator (i.e., equation 6.2). Peterson and Iyengar (2020) apply this approach to an estimate of partisan differences among respondents who were provided information about the questions before answering them. Alternatively, researchers may score

[2] Formally, $c_i = p_i$ if the best guess is correct and $1 - p_i$ if the best guess is incorrect. More compactly, $c_i = \max(p_i, 1 - p_i)$.

respondents who are below the threshold as an 0.5. When below-threshold responses are scored as an 0.5, as is common when explicit DK options are offered, the measured belief difference is equal to

$$\Delta_{\text{DK}} = \frac{1}{N_D}\sum_{i\in\mathbb{D}}\Big(\mathbb{1}(p_i > 0.5, c_i > \tau) + 0.5\times\mathbb{1}(c_i < \tau)\Big) - \frac{1}{N_R}\sum_{i\in\mathbb{R}}\Big(\mathbb{1}(p_i > 0.5, c_i > \tau) + 0.5\times\mathbb{1}(c_i < \tau)\Big). \tag{6.3}$$

Neither the "drop" strategy nor the "treat-as-completely-uncertain" strategy should be expected to eliminate certainty bias, except perhaps by chance. For respondents whose certainty level exceeds the threshold, the bias is equal to $\frac{1}{N}\sum(\mathbb{1}(p_i > 0.5) - p_i)$, no different than the best guess estimator. For these respondents, the bias is either undefined or $\frac{1}{N}\sum(0.5 - p_i)$.

We saw in Chapter 3 that for DK response options, τ is fairly low. This suggests that the treat-as-uncertain method may not do enough eliminate certainty bias. One attempt to compensate for this by choosing a higher-threshold strategy, e.g., by encouraging rather than just offering DK responses. The best-case scenario is that such a strategy could offer a reliable approximation; the worst case, that threshold-based strategies measure belief differences about as reliably as a stopped clock measures time. To illustrate this potential for arbitrariness, Appendix D presents a simple simulation study using the data analyzed below. It calculates the estimated partisan belief difference using three strategies: a probabilistic estimate based on the survey measure of p_i for each respondent, a threshold-based estimate that drops all respondents below the threshold (i.e., if $c_i < \tau$), and a threshold-based estimate that scores all respondents below the threshold as an 0.5, as if they were completely ignorant. Each strategy is applied using every possible threshold in the set $\{0.50, 0.51, 0.52, ..., 0.99, 1\}$. The results cast threshold-based estimators as a haphazard, shoot-from-the-hip strategy that is vulnerable to substantial under- and over-correction. Accordingly, the empirical portion of this chapter focuses on certainty bias as it is defined above: the difference between the probabilistic belief difference and the difference in best guesses.

Approach

This chapter's empirical analysis begins by dividing six component surveys into two studies. Study 1 uses data from four surveys conducted online in 2018 and 2019 (total $N = 8,428$). Study 2 makes use of two additional surveys conducted in 2020 (total $N = 1,471$). Respondents were recruited through Lucid, which quota samples online survey respondents to match Census demo-

Table 6.1: List of question topics.

			Pr(best guess is correct)		
Category	Correct answer (partisan valence)	Surveys	D	R	Diff
Economic	Budget deficit increased (D)	3, 5, 6	0.83	0.76	-0.08
	GDP growth below 4% (D)	5, 6	0.58	0.36	-0.21
	Health insurance declined (D)	2, 3, 4, 5	0.54	0.41	-0.13
	Record trade deficit w/ China (D)	2	0.87	0.84	-0.03
	Trade deficit increased (D)	3	0.69	0.51	-0.18
	Inflation below average (R)	3, 4, 5, 6	0.36	0.54	0.18
	Median real wage increased (R)	2, 3, 4, 5	0.64	0.84	0.21
	Stock market increased (R)	3	0.72	0.88	0.17
	Unemployment declined (R)	2, 3, 4, 5, 6	0.71	0.82	0.11
Controversies	Cohen paid off Stormy Daniels (D)	3	0.92	0.83	-0.10
	Obama released birth certificate (D)	1, 2	0.78	0.56	-0.22
	Rs allowed to call impeach. witnesses (D)	4	0.61	0.28	-0.33
	Trump Article II statement (D)	6	0.74	0.48	-0.26
	Trump said 'grab them' (D)	2, 3, 6	0.89	0.71	-0.18
	Clinton emailed classified info (R)	2, 3	0.90	0.97	0.06
	Mueller did not say Trump colluded (R)	2, 6	0.61	0.73	0.13
	Obama said DAPA would ignore law (R)	6	0.37	0.52	0.15
	Sanders supporter shot Scalise (R)	3	0.51	0.67	0.17

graphics, and Amazon Mechanical Turk (MTurk), which provides diverse national samples that skew young and left-leaning. The subset of these surveys that included political awareness questions, a panel component, or the revealed belief measure were analyzed above in Chapters 4 and 5. Table 6.1 lists the topic of each question, the answer choices, and the percentage of respondents selecting the correct answer by political party. Appendix E provides complete information about each survey.

Both studies featured a series of binary (two-choice) questions. Immediately after each respondent chose their answer to each question, a 50-100 scale appeared. This measures c_i, defined above the probability that each respondent assigns to their best guess. For the same set of facts, study 2 also included a revealed preference measure of beliefs. Chapter 4 describes both of these measures in greater depth, and Appendix E provides further details.

Together, these measures provide a basis to compare an estimate of the probabilistic belief difference (equation 6.1), the partisan difference in best guesses (equation 6.2), and certainty bias, which is the difference between them. Study 1 calculates these quantities using only stated measures of belief, while Study 2 incorporates the revealed preference measure of belief.

How Large is Certainty Bias?

For a first look at the effect of certainty bias on estimates of partisan differences in factual beliefs, Study 1 pools across four surveys conducted in 2018 and 2019. Of these 25 question-survey pairs, 15 concerned statistics about economic conditions and 10 concerned politicized controversies. Across the questions included in Study 1, the average partisan difference in the percentage of respondents answering correctly was about 15 percent, near the high end of what is typically observed in studies of partisan perceptual differences (Jerit and Barabas 2012; Roush and Sood 2020).

To crystallize the basis for the aggregate analysis below, it is helpful to first examine partisan belief differences on each of the individual questions. Throughout the analysis, responses are coded on a 0 to 1 scale, where 0 represents absolute certainty about the incorrect answer and 1 represents absolute certainty about the correct answer. Separately for each question, study, and empirical strategy, the partisan difference is estimated by

$$\text{Belief}_i = \alpha + \beta \text{ Republican}_i + \epsilon_i, \tag{6.4}$$

where Belief_i is the measure of belief in the correct answer, α is the mean among Democrats, and β is the difference in means between Democrats and Republicans.

For each of the questions, Figure 6.2 plots two estimates of β. The top row of each panel uses the best guess estimator (equation 6.2), while the bottom row of each panel uses the probabilistic estimator (equation 6.1). Comparing the two rows gives an estimate of the certainty bias that is introduced by empirical strategies that only focus on the respondent's best guess. The figure is organized so that when possible, questions with the same partisan valence appear on the same side of the figure and questions on the same topic are stacked atop one another.

On the economic questions, the difference in best guesses ranges from a high of 0.25 points (wages increased, survey 4) to a low of 0.03 points (record trade deficit with China, survey 2). Switching to the probabilistic estimator, which is not affected by certainty bias, shrinks this range from 0.16 points to 0.01 point. Across questions, the size of certainty bias seems to vary systematically with the subject matter. On the questions about wage growth, inflation, and the percentage of Americans with health insurance, certainty bias approximately doubles the estimate of partisan belief differences. On the unemployment and budget deficit questions, certainty bias has little effect.

On the controversy questions, the difference in best guesses ranges from 0.33 points (Republicans allowed to call impeachment witnesses, survey 4) to a low of 0.04 points (Clinton emailed classified information, survey 2). Switching to the probabilistic estimator shrinks this range to a high 0.24 points and a low of 0.08 points. Variation again appears between questions. In three instances — the two Clinton email questions, and the question about then-Trump lawyer Michael Cohen paying off adult film actress Stormy Daniels — removing certainty bias at least slightly *in-*

Figure 6.2: Estimated partisan difference by question and estimator, study 1.

(a) Economic and fiscal questions

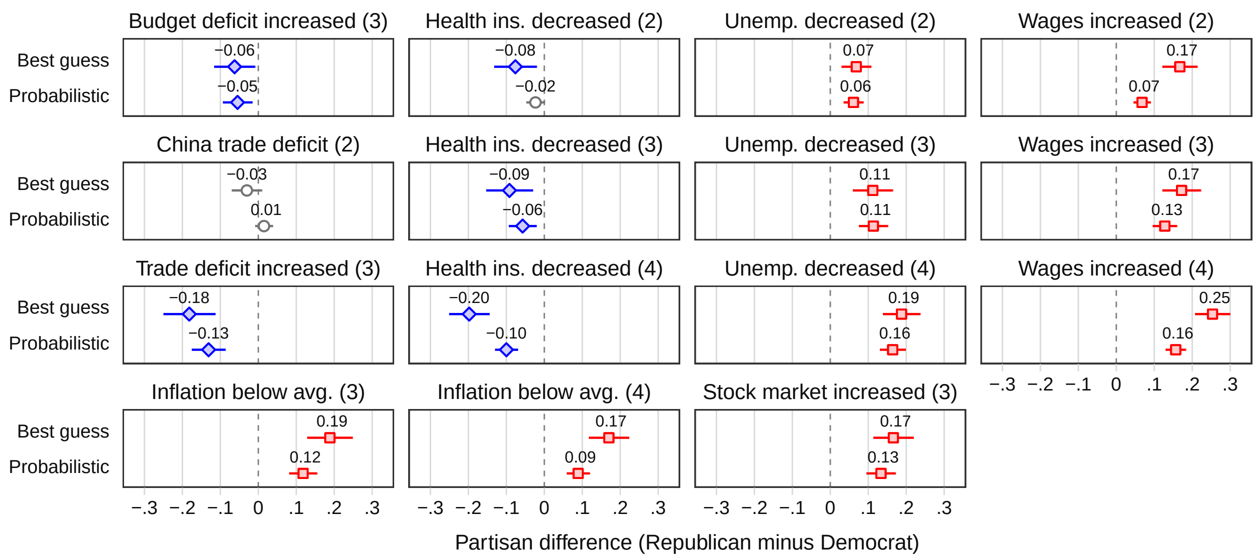

(b) Politicized controversies

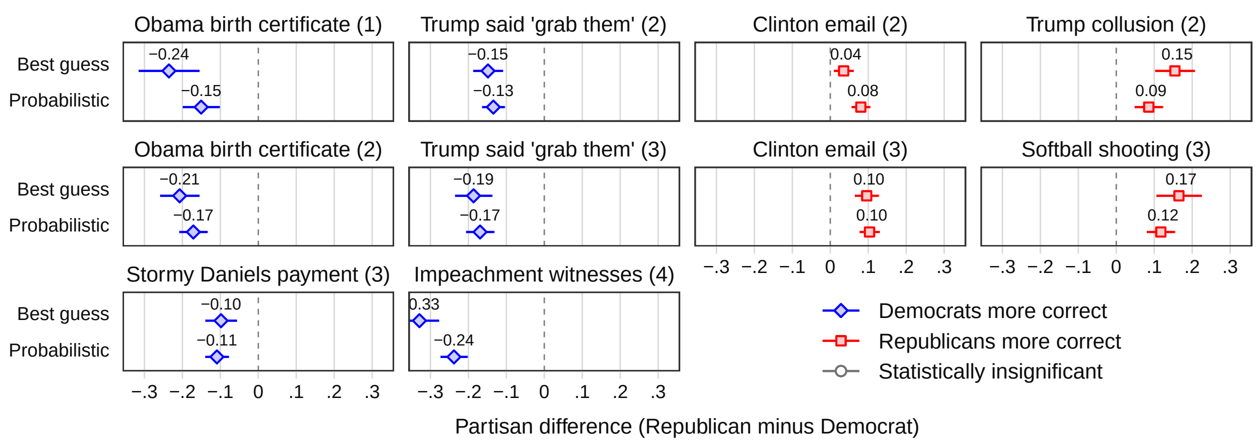

Note: Separately for each question, this figure presents estimates of β from equation (6.4) using the two estimators (best guess and probabilistic). Horizontal bars represent 95 percent confidence intervals.

creases the estimates of partisan differences. The reasons for this heterogeneity are examined in detail below.

To aggregate across questions, partisanship is recoded to align with the partisan valence of the correct answer to each question. Define the *congenial party* as the party whose political interests are supported by the facts and the *uncongenial party* as the party whose political interests are inconsistent with the facts. For example, at the time the surveys were fielded, unemployment had declined over the previous year and wages had risen; this was congenial to Republicans, who controlled the presidency, and by a zero-sum logic, uncongenial to Democrats. Effectively, the cross-question estimates fold each panel of Figure 6.2 at its midpoint, then take the average.

To estimate the overall average level of certainty bias, OLS was used to estimate the parameters in

$$\text{Correct}_{ijk} = \alpha + \beta_1 \text{Congenial}_{ij} + \beta_2(\text{Congenial}_{ij} \times \text{Best guess}_{ik}) + \beta_3 \text{Best guess}_{ik} + \epsilon_i, \quad (6.5)$$

where i indexes respondents, j indexes questions, k indexes measurement strategies, Congenial_{ij} indicates whether the correct answer was congenial (1) or uncongenial (0) to the respondent's partisanship, and Best guess_{ik} indicates whether the partisan difference is estimated using the probabilistic (0) or best guess (1) estimator. To account for the fact that each respondent contributes multiple observations, standard errors are clustered at the respondent level.

Table 6.2 displays the estimates, with clustered standard errors in parentheses. The estimate of β_1 gives the difference in probabilistic beliefs, and the estimate of β_2 gives the degree of certainty bias. Dividing the estimate of β_2 by the estimate of β_1 expresses certainty bias as a percentage of the difference in probabilistic beliefs. Likewise, adding the estimates of β_1 and β_2 yields the partisan difference that would be estimated in the typical survey analysis.

Across all 25 question-topic pairs, certainty bias inflates the estimate of partisan perceptual differences by 39 percent, from 0.106 to 0.147 points on the 0-1 scale of belief in the correct answer (i.e., p_i above). This includes 57 percent on the economic questions, from 0.86 to 0.135 ($\hat{\beta}_2 = 0.049$, s.e. $= 0.004$), and 23 percent on the controversy questions, from 0.128 to 0.158 ($\hat{\beta}_2 = 0.030$, s.e. $= 0.005$). Just as suggested by the question-by-question results, certainty bias tends to be larger on economic questions than on questions about politicized controversies, both in scale points (difference $= 0.019$, s.e. $= 0.007$) and as a percentage of the difference in probabilistic beliefs ($57 - 23 = 34$ percent).

Table 6.2: Regression estimates of certainty bias, study 1.

	Economic (1)	Controversies (2)	Pooled (3)
	Dependent variable: belief in correct answer.		
β_1 Congenial	0.086^{**}	0.128^{**}	0.106^{**}
	(0.004)	(0.006)	(0.004)
β_2 Congenial × best guess	0.049^{**}	0.030^{**}	0.041^{**}
	(0.004)	(0.005)	(0.003)
β_3 Best guess	0.034^{**}	0.027^{**}	0.031^{**}
	(0.003)	(0.004)	(0.003)
α Constant	0.556^{**}	0.644^{**}	0.591^{**}
	(0.003)	(0.005)	(0.003)
Adj. R^2	0.026	0.039	0.031
Num. obs.	34650	22464	57114

Note: $^{*}p < 0.05$, $^{**}p < 0.01$. Clustered standard errors in parentheses. Each cell entry is a coefficient estimate from equation (6.5). For a full list of survey questions, see Appendix E.

Is Certainty Bias Distinct from Expressive Responding?

Though these estimates suggest that certainty bias substantially inflates estimates of partisan differences in factual beliefs, it could be that the results are distorted by expressive responding. Moreover, given the theorized relationship between uncertainty and expressive responding (Bullock et al. 2015; Bullock and Lenz 2019), it could be that reducing or eliminating one of the biases also helps eliminate the other.

To examine certainty bias in a context that also accounts for expressive responding, study 2 paired a similar set of survey questions with the revealed belief measure, which infers respondents' best guesses and probabilistic beliefs through a series of costly choices. As discussed above, the financial incentives should encourage respondents to reveal their beliefs accurately rather than cheerleading or taking a biased sample of their considerations. By estimating certainty bias using both measures of belief, one can get a sense of whether expressive responding affects estimates of certainty bias.

The empirical strategy most similar to standard survey practices is the "best guess, stated" strategy, which measures beliefs using direct questions and makes use of the best guess estimator (equation 6.1). The conventional strategy suggests that partisan perceptual bias is often substantial, and varies widely across questions. Among the economic questions, the partisan difference in stated best guesses ranges from 0.26 and 0.27 on the wage growth and economic growth questions (survey 2) to 0.02 on the unemployment question (survey 1). Similarly wide variation obtains on the questions about political controversies, from a low of 0.07 to a high of 0.25.

To examine the degree to which expressive responding and certainty bias affect the estimates of partisan differences, Figure 6.3 displays the partisan difference using each of the four combinations of measurement strategies (stated and revealed) and estimators (best guess and probabilistic). The first row displays the difference in stated best guesses, described just above. The other three rows account for expressive responding, certainty bias, and their combination.

To examine how expressive responding affects the partisan difference in best guesses, one can compare the first row to the second row. This is the closest comparison to previous studies of expressive responding. Consistent with these studies, accuracy incentives tend to reduce partisan differences. The bottom two rows use the same two measurement technologies, but account for certainty bias by switching from the best guess estimator (equation 6.2) to the probabilistic estimator (equation 6.1). Comparing these two rows gives this chapter's preferred estimate of the bias introduced by expressive responding, which also accounts not just for respondents' best guesses, but the degree to which respondents claim to be certain or uncertain. Using either measure, the patterns are broadly similar.

Another set of comparisons across the rows shows the effect of certainty bias. Comparing the first row, "stated best guess," to the third row, "stated probability," shows the effect of certainty bias when no financial incentives are present. Comparing the second row, "revealed probability," to the fourth row, "stated probability," shows this effect when respondents reveal their best guess and probabilistic belief through costly choices. Both comparisons suggest that certainty bias inflates estimates of partisan perceptual differences, and tends to do so by a modestly larger amount than expressive responding.

The total effect of certainty bias and expressive responding can be seen by comparing the top row, "stated best guess," to the bottom row, "revealed probability." Together, uncertainty and accuracy incentives shrink partisan differences substantially, with the largest reductions coming on the economic questions with the initially largest partisan differences. After removing these biases, perceptual differences on the economic questions range between 0.03 percent and 0.11.

Removing the influence of certainty bias and expressive responding also suggests that there is a bigger difference between economic and controversy questions than was evident from the differences in best guesses. Whereas partisan differences in best guesses over the two facts that are inconvenient for President Trump — his "grab them" boast and his claim to unlimited power under Article II — were in the same vicinity as partisan differences in best guesses about the state of the economy, the revealed measure shows belief differences that are 50 percent larger than on any economic question

Figure 6.3: Estimated partisan difference by question and measurement strategy.

(a) Economic and fiscal questions

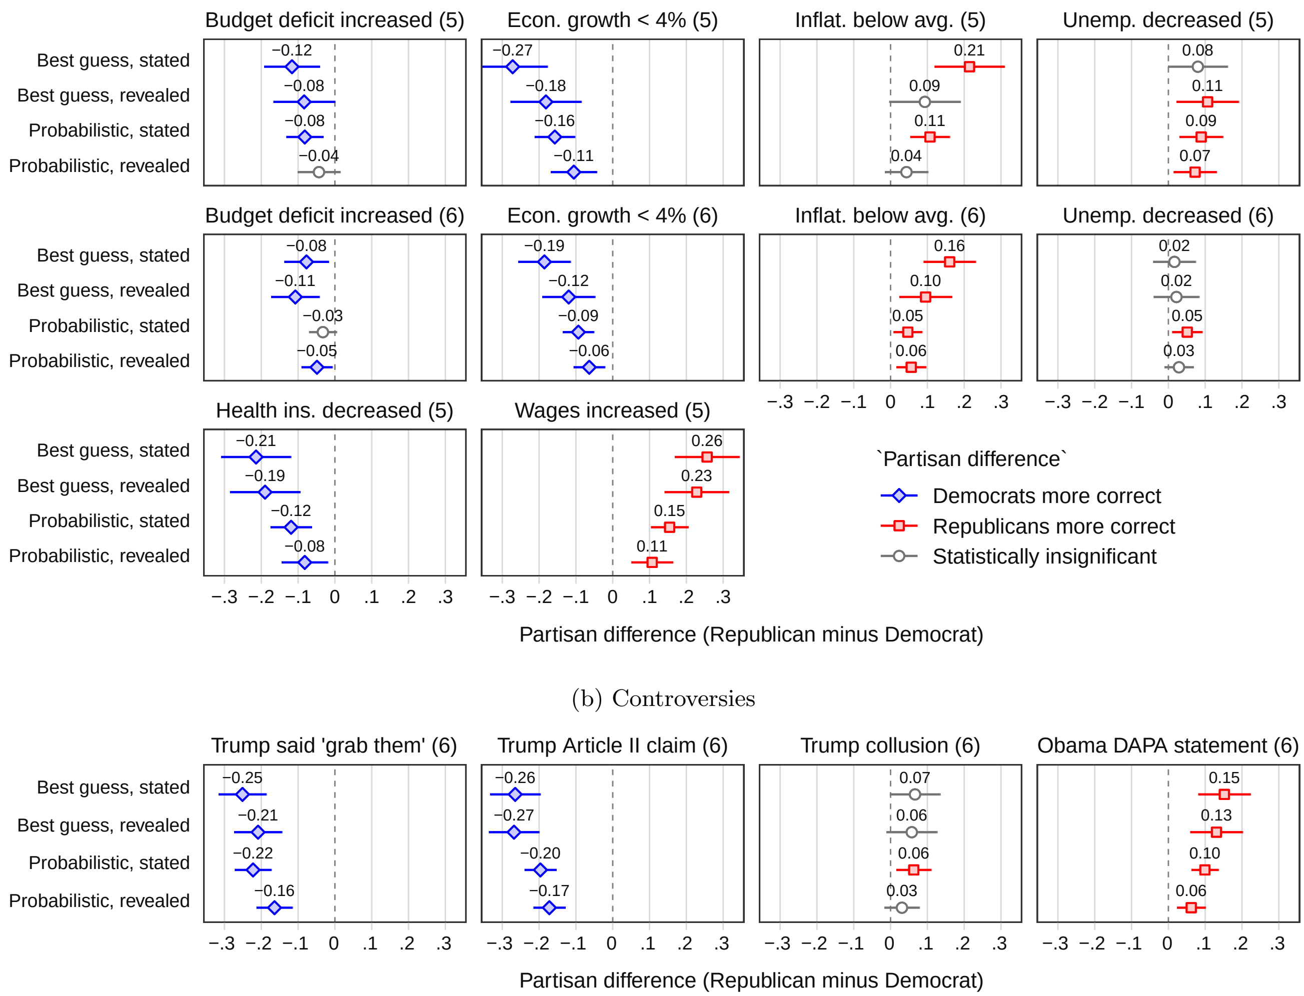

(b) Controversies

Note: Separately for each question and study, this figure presents estimates of β from equation (6.4) using two estimators (best guess and probabilistic) and measurements (incentive and no incentive). Horizontal bars represent 95 percent confidence intervals.

Table 6.3: Regression estimates of certainty bias and expressive bias.

	Dependent variable: belief in correct answer.					
	Survey 5	Survey 6			Pooled	
	Econ. (1)	Econ. (2)	Cont. (3)	All (4)	Econ. (5)	All (6)
$\sum_{k=1}^{4} \beta_k$ Difference in stated best guesses	0.190^{**}	0.102^{**}	0.208^{**}	0.156^{**}	0.141^{**}	0.165^{**}
	(0.020)	(0.018)	(0.018)	(0.013)	(0.013)	(0.011)
β_1 Revealed partisan difference	0.073^{**}	0.040^{**}	0.122^{**}	0.081^{**}	0.055^{**}	0.079^{**}
	(0.012)	(0.010)	(0.011)	(0.008)	(0.008)	(0.007)
β_2 Certainty bias	0.070^{**}	0.034^{**}	0.068^{**}	0.051^{**}	0.050^{**}	0.056^{**}
	(0.012)	(0.012)	(0.010)	(0.009)	(0.009)	(0.007)
β_3 Expressive responding	0.043^{**}	0.013	0.040^{**}	0.027^{**}	0.027^{**}	0.032^{**}
	(0.013)	(0.010)	(0.010)	(0.007)	(0.008)	(0.006)
β_4 Certainty × expressive responding	0.003	0.014	−0.021	−0.004	0.009	−0.002
	(0.014)	(0.012)	(0.011)	(0.009)	(0.009)	(0.007)
Adj. R^2	0.030	0.010	0.044	0.023	0.017	0.025
Num. obs.	9864	12352	12584	24936	22216	34800

Note: $^{*}p < 0.05$, $^{**}p < 0.01$. Clustered standard errors in parentheses. Each cell entry is a coefficient estimate from equation (6.6). Columns correspond to the study and question topics as listed in Table 6.1. The excluded parameter estimates, β_5 through β_7, are all statistically insignificant in the pooled estimate (column 6).

(0.16 and 0.17, versus a maximum of 0.11).

To estimate how the estimates of partisan differences are inflated by the combination of expressive responding and certainty bias, OLS was used to estimate

$$\begin{aligned}\text{Correct}_{ijk} = \alpha &+ \beta_1 \text{Congenial}_{ij} + \beta_2(\text{Congenial}_{ij} \times \text{Best guess}_{ik}) + \beta_3(\text{Congenial}_i \times \text{Stated}_{ik}) \\ &+ \beta_4(\text{Congenial}_{ij} \times \text{Best guess}_{ik} \times \text{Stated}_{ik}) + \beta_5 \text{Best guess}_{ik} + \beta_6 \text{Stated}_{ik} \\ &+ \beta_7(\text{Best guess}_{ik} \times \text{Stated}_{ik}) + \epsilon_i, \qquad (6.6)\end{aligned}$$

where Stated_{ik} indicates whether the estimate takes respondents' claims about their beliefs at face value (1) or infers them from costly choices (0). All other terms are defined as they were in equation (6.5).

Table 6.3 reports estimates of equation (6.6) for each study and question category, as well as estimates that combine data from both studies. Rather than copying the variable names from equation (6.6), the labels on β_1 through β_4 indicate their relationship to the quantities discussed in the empirical framework. The top row uses a separate regression to estimate the sum of β_1 through β_4. This specification simply switches the reference group in equation (6.6) from revealed probabilities to stated best guesses.

Comparing the estimates of β_1 suggest that survey estimates of partisan perceptual bias are substantially inflated by, but not entirely a result of, certainty bias and expressive bias. Pooling across both studies and question categories, the average difference in probabilistic beliefs is 0.079 on a 0 to 1 scale. This amounts to less than half the pooled difference in the stated best guesses, 0.165. This includes 0.073 on survey 5's economic questions, 0.040 on survey 6's economic questions, and 0.122 on survey 6's political controversy questions. On the economic questions, the combination of certainty bias and expressive bias boosted the partisan difference by more than 150 percent,[3] while on the controversy questions, the two biases increased the difference by 70 percent.[4]

Certainty bias is quantified by the estimates of β_2. This estimate is the cross-question average difference between the third and fourth rows of Figure 6.3. In each column of Table 6.3, the estimate is positive, indicating that certainty bias generally inflates the estimated partisan difference. For the economic questions, certainty bias increases the estimated partisan belief difference from 0.055 to 0.105, an increase of about 91 percent. For survey 5, this figure is 96 percent, and for survey 6, 85 percent. For the political controversy questions in survey 6, certainty bias inflates the estimated partisan belief difference from 0.122 to 0.190, an increase of 56 percent. Aggregating across all questions, the increase from 0.079 to 0.135 works out to 71 percent. These percentage differences are larger than observed in Study 1 partly because the point estimates are larger and partly because the baseline also accounts for expressive bias.

The effect of expressive responding is quantified by the estimates of β_3. This estimate is the cross-question average difference between the second and fourth rows of Figure 6.3. For the economic questions, expressive bias increases the estimate of partisan belief differences from 0.55 to 0.82, an increase of about 50 percent. Expressive bias on the questions about political controversies appears somewhat larger in absolute terms, but smaller in percentage terms: the increase from 0.122 to 0.162 is equal to a bit less than one-third. Pooling across all questions and studies, expressive bias inflated partisan differences from 0.079 to 0.111, a difference of 0.032 (41 percent).

Would the estimates of certainty bias be different if expressive bias had not been accounted for, as in Study 1? The estimates of β_4 check for the degree of overlap between the two biases. These estimates quantify the difference in certainty bias between the stated and revealed belief measures . Pooling across all questions, there is not much evidence that the estimates of certainty bias depended on the inclusion of accuracy incentives ($\hat{\beta}_4 = 0.002$, s.e. $= 0.007$).

Together, these results replicate the patterns observed in study 1, and further suggest that certainty bias is at least largely empirically distinct from expressive bias. Though the results here do not rule out the potential for modest overlap between the two biases, unincentivized measures of

[3] (0.141 - 0.055) / 0.055 = 1.56
[4] (0.208 - 0.122) / 0.122 = 0.70

certainty bias are likely to be fairly trustworthy.

Partisan Differences in Certainty

Across the 39 questions included in the two studies, the size of certainty bias ranged from 0.11 points (+186 percent; inflation, survey 5) to -0.05 points (-56 percent; Clinton email, survey 2). What explains this variation? Appendix D shows that certainty bias can be broken into two components: the effect of generalized uncertainty, and the effect of partisan differences in certainty.

As a first look at these factors, Figure 6.4 plots respondents' average certainty level conditional on two factors: whether the fact is positively valenced ("congenial") for the respondent's party, and whether the respondent answers correctly or incorrectly. The most striking pattern in Figure 6.4 is that correct answers tend to be stated with considerably more confidence than incorrect answers. Pooling across all questions, the average correct answer is stated with 30 percent greater confidence than the average incorrect answer.[5] Regardless of party, correct answers are closer to knowledge than incorrect answers are to "incorrect knowledge" (Hochschild and Einstein 2015, 11), in the sense in which these terms were respectively defined in Chapters 4 and 5.

Among respondents who answer correctly, those for whom the truth is politically convenient are 0.037 points more certain of it (s.e. = 0.002). On the 0.5 to 1 scale, this means that respondents who provide the correct answer about convenient truths are 13 percent closer to complete certainty than are respondents who provide the correct answer about inconvenient truths. Split by category, the certainty gap on correct answers is 0.044 on controversy questions (s.e. = 0.003) and 0.029 on economic questions (s.e. = 0.003). This means that correct answers to questions about convenient truths tend to be closer to knowledge, while correct answers to inconvenient truths reflect a larger dose of lucky guessing.

A surprising pattern, at least in light of contemporary scholarship's emphasis on partisan incentives to believe in misinformation, is the lack of an equivalent certainty gap among respondents who choose the incorrect answer. The uncongenial party, which has a partisan incentive to endorse the incorrect answer, is only about 0.001 points more certain of it on average, or 0.6 percent closer to complete certainty on the 0.5 to 1 scale. Split by category, there is a small but statistically significant gap of 0.015 points on the controversy questions (s.e. = 0.005) and a negative gap that is only borderline statistically significant on the economic questions (-0.006, s.e. = 0.003, $p = 0.08$). This is traceable to a tendency for Republicans to be more certain about incorrect answers

[5] The calculation is $(.8078 - .7375)/(.7375 - 0.5)$.

Figure 6.4: The partisan certainty gap.

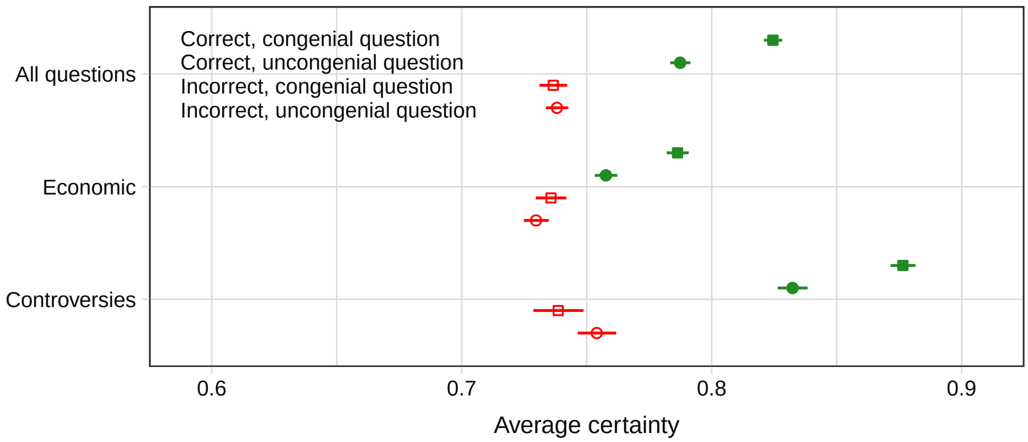

to economic questions than Democrats, regardless of party (see Figure 6.6 below).

To set more concrete expectations for how these factors should impact estimates of partisan differences, Appendix D.1 decomposes certainty bias into two components: the effect of generalized uncertainty, and the effect of the partisan certainty gap. This can be expressed as

$$\underbrace{\hat{\Delta}_{BG} - \hat{\Delta}}_{\text{Total certainty bias}} = \underbrace{\hat{\Delta}_{BG} - \hat{\Delta}_{EC}}_{\text{Effect of general uncert.}} + \underbrace{\hat{\Delta}_{EC} - \hat{\Delta}}_{\text{Effect of partisan cert. gap}}, \tag{6.7}$$

where Δ_{EC} is the partisan belief difference under a counterfactual in which the average Democrat and Republican $e$qually $c$ertain of their answers, and all other terms are defined above.

If the introduction of an equal-certainty counterfactual strikes the reader as odd, consider that in assuming all respondents are perfectly certain of their answer, the best guess estimator also treats all respondents as if they are equally certain. In fact, the Appendix D.1 shows that the best guess estimator is a special case of the equal certainty counterfactual. As it is operationalized here, the equal-certainty counterfactual allows two implications of the best guess estimator's defining assumption to be unpacked in stages: first relax the assumption that all respondents are completely certain of their answers, then relax the assumption that Democrats and Republicans are equally certain.

Based on the decomposition in equation (6.7), Appendix D.1 shows the following:

1. Generalized uncertainty always shrinks the partisan difference toward zero.
2. Partisan differences in certainty increase partisan belief differences, and counteract certainty bias, whenever two conditions are met: the party that answers correctly more often is (a) more certain on average about its correct answers and (b) less certain on average about its incorrect answers, relative to the other party.

3. Partisan differences in certainty decrease partisan belief differences, and reinforce certainty bias, whenever two conditions are met: the party that answers correctly more often is (a) less certain about its correct answers and (b) more certain about its incorrect answers, relative to the other party.

In the results above, the certainty gap on correct answers supports condition 2a: the congenial party, which answers correctly more often, also tends to be more certain of its correct answers. The less consistent certainty gap on incorrect answers suggests that condition 2b holds on average for controversy questions, but holds weakly if at all for economic conditions. This means that in general, one should expect generalized uncertainty and the partisan certainty gap to have countervailing effects, especially on questions about politicized controversies.

Consistent with these analytic results, the data show that generalized uncertainty and the partisan certainty gap tend to have countervailing effects on partisan belief differences. To show this, Figure 6.5 plots three separate estimates of the partisan belief difference. In addition to the two estimates that were compared in the previous sections (the difference in best guesses, $\hat{\Delta}_{BG}$, and the difference in beliefs, $\hat{\Delta}$), the equal-certainty counterfactual ($\hat{\Delta}_{EC}$) is added, permitting a visualization of the both components of certainty bias.

The most general implication of Figure 6.5 is that absent partisan certainty gaps, certainty bias would be larger. To see this, first compare the partisan difference in best guesses (hollow grey triangles) to the partisan belief difference that would realize under equal certainty (hollow grey circles). Under this assumption, removing certainty bias would reduce the average partisan difference across all questions by 0.067, from 0.152 to 0.085. The partisan certainty gap counteracts the effect of generalized uncertainty. The black, solid circle representing $\hat{\Delta}$ sits at 0.108, 0.023 above the equal-certainty estimate and 0.044 below the difference in best guesses. This implies that across the board, the certainty gap increases partisan differences in factual beliefs by about 20 percent, and reduces the magnitude of certainty bias by more than one-third.

The certainty gap's effect also reflects the expected variation by question type. On questions about politicized controversies, the certainty gap's influence on partisan differences is about twice as large as it is on economic questions (0.034 versus 0.014, difference = 0.020, s.e. = 0.004) and as a percentage of the estimate of certainty bias under the equal-certainty counterfactual (52 percent versus 22 percent).

Figure 6.5: Effect of partisan certainty gap on partisan belief differences.

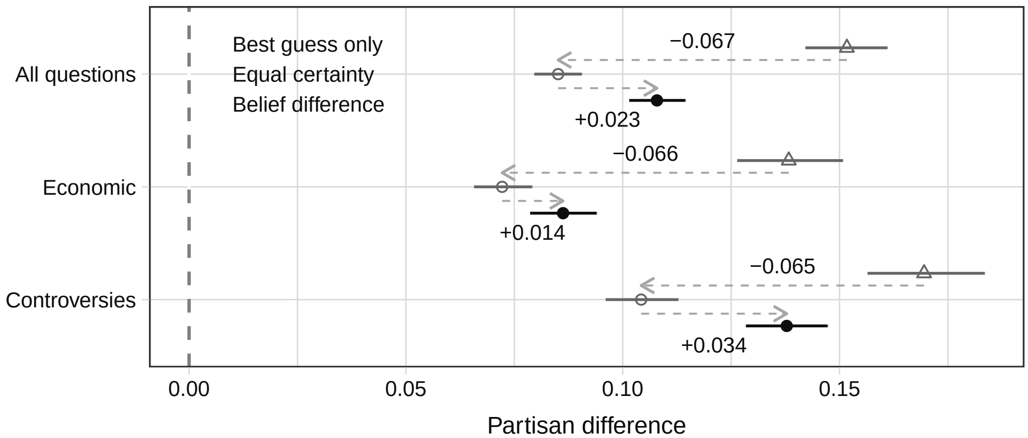

Differences between Questions

Variation in the two contributors to certainty bias — generalized uncertainty and the partisan certainty gap — also explains the heterogeneity that exists among individual questions. For each question in both studies, Figure 6.6 plots the same information that is displayed in Figures 6.4 and 6.5, split by political party and the partisan valence of the question. The two leftmost panels display the partisan certainty gap (i.e., the information from Figure 6.4), while the rightmost panel displays the three estimates of the partisan belief difference (i.e., the information from Figure 6.5).

The question-by-question results crystallize the broad-based role that differential knowledge of convenient truths can play in driving partisan differences. Visually, this is evident in the left panel, which plots the partisan difference in certainty among respondents who answered correctly. In the top row, the blue diamonds tend to be to the right of the red squares, indicating that Democrats are more certain when the truth is more convenient to Democrats. In the bottom row, the red squares tend to be to the right of the blue diamonds, indicating the opposite. Of the 39 point estimates, the congenial party is more certain in 36, including 27 that are statistically significant. Consistent with the pooled results above, this means that condition 2a above almost always holds.

By contrast, partisan convenience is a less consistent predictor of degree of belief in falsehoods that would be convenient if they were true. The middle panel plots the partisan difference in certainty among respondents who answered incorrectly. The party with greater incentive to endorse the falsehood is more certain of the incorrect answer in 25 of the 39 cases; of these, just seven are statistically significant.[6] Again consistent with the pooled results, this suggests that condition 2b

[6]This is driven by Republicans' tendency to be more certain of incorrect answers to economic questions, regardless of their partisan valence. Visually, this is evident from the red squares' tendency to be to the right of the blue diamonds in both the top and bottom rows.

holds inconsistently.

To show how these factors translate into partisan belief differences, the rightmost panel displays the same information that is displayed in Figure 6.5, but for each individual question. The entire figure is sorted by the solid black dots, which represent the partisan belief difference. The hollow grey dots representing the partisan difference in best guesses jut out inconsistently to the right, representing the variable influence of certainty bias on estimates of partisan belief differences. Meanwhile, the hollow grey circles represent the partisan belief difference that would realize if both parties were equally certain. When these sit atop the black circles, it indicates that the partisan certainty gap plays a small role; when the black circles are substantially to the right, the certainty gap plays a larger role. Outside the third panel, the numbers display the effect of the partisan certainty gap on partisan differences (i.e., the difference between the two sets of circles), with the ranking among all questions in parentheses.

Together, variation in average certainty and the partisan certainty explain the substantial variation in certainty bias across questions. The questions whose partisan differences in best guesses were the same or smaller than the partisan belief difference — Clinton's emails, trends in the unemployment rate, Trump's "grab them" comment, and Trump's payment to Stormy Daniels — have relatively high certainty (left two panels) and sizeable partisan certainty gaps (right panel). On these questions, the partisan certainty gap inflates the partisan belief difference by about 0.04 to 0.06. By contrast, the questions on which certainty bias was largest — concerning inflation, wage growth, GDP growth, and health insurance — have low certainty (left two panels) and smaller partisan certainty gap (right panel).

The question-level results add context to the finding that certainty bias is less influential, and more strongly counteracted by partisan certainty gaps, on questions about politicized controversies. This pattern is clearest on the nineteen questions for which the truth was congenial to Democrats. Among these questions, the partisan certainty gap had its eight largest effects on the eight questions about politicized controversies. In order, these concerned Trump's claim that Article II of the Constitution allows him to do whatever he wants; Trump's tape recorded about sexually assaulting women; a Trump lawyer's payoff of adult film star Stormy Daniels; the false claim, promoted heavily by Trump, that his predecessor in office was born outside the U.S. and hence ineligible to be president; and the false claim by Republican opinion leaders that Republicans were not allowed to select witnesses at his impeachment hearings.[7] Partisans' tendency to know more about controversies

[7]This survey was fielded while the proceedings were still in the Democrat-controlled House of Representatives, and before the proceedings moved to the Republican-controlled Senate.

Figure 6.6: The partisan certainty gap and its effect on belief differences.

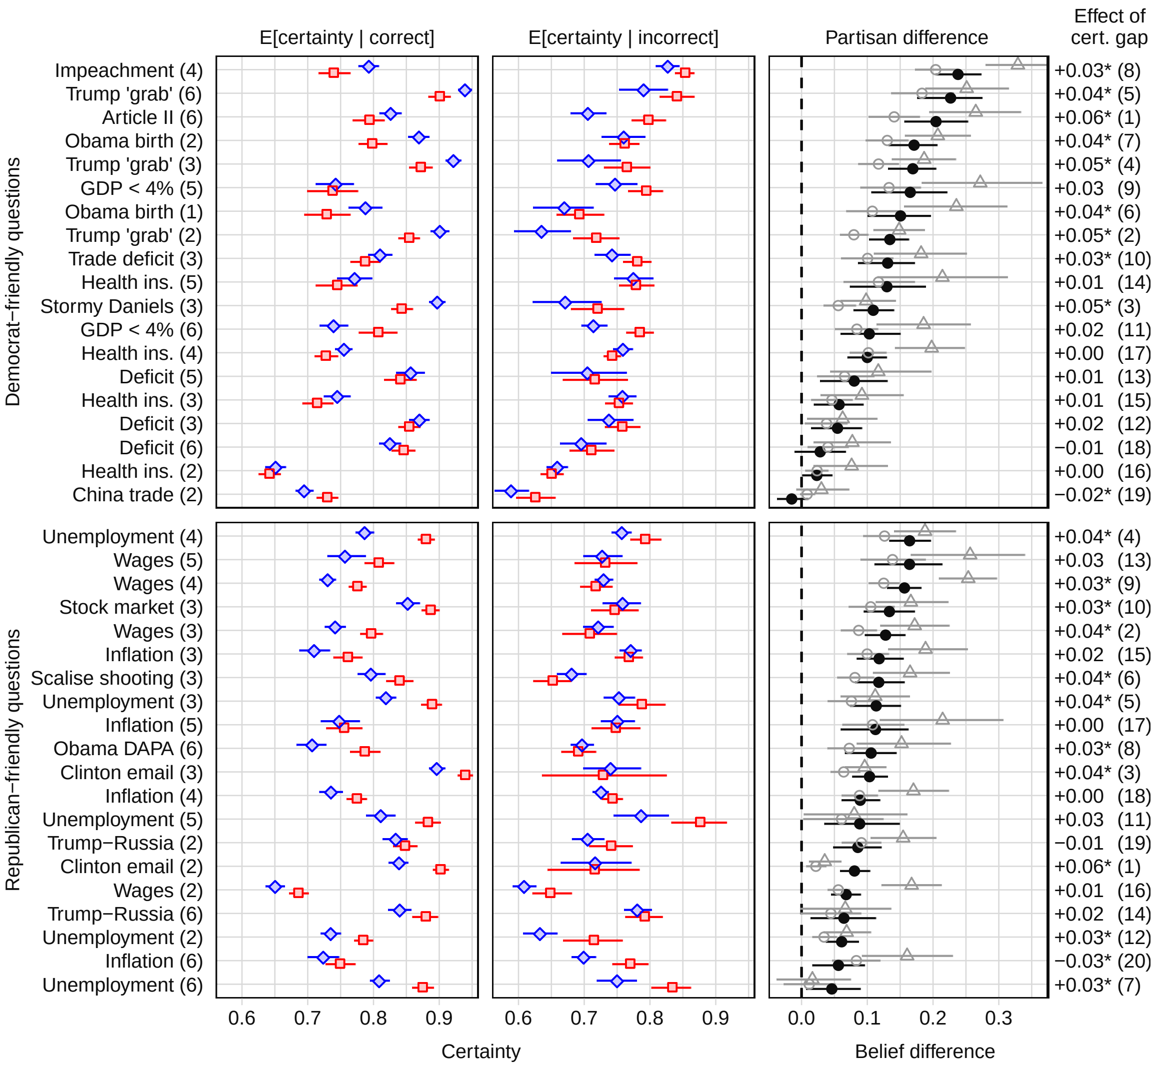

Note: In the left and middle panels, blue diamonds are average certainty among Democrats; red squares, among Republicans (i.e., the same symbols used in Figures 6.2 and 6.3). In the right panel, grey hollow triangles are the partisan difference in best guesses; grey hollow circles, the partisan belief difference in the equal-certainty counterfactual; and black solid circles, the partisan belief difference (i.e., the same symbols used in Figure 6.5). Outside the right panel, the effect of the partisan certainty gap (i.e., the difference between the hollow and solid dots) is printed, with a * indicating a statistically significant difference at the 0.05 level (two-tailed). The numbers in parentheses are the rank order of the effect of the partisan certainty gap.

that serve de-legitimize the other side, and less about democratic transgressions on one's own side, systematically boosts partisan belief differences on these matters.

Relative to the economic questions, the controversy questions conform especially well to a key pattern seen in Figure 6.4: that correct answers are stated with greater confidence than incorrect answers, regardless of party. At the same time, the two exceptions to this rule highlight the substantive value of measuring certainty rather than assuming that rules of thumb always hold.

The exception to the first half of this formulation, "correct answers are stated with greater confidence," concerns witnesses at Trump's impeachment trial. Here, among both Democrats and Republicans, respondents who answered incorrectly were more confident than respondents who answered correctly. In fact, for both Democrats and Republicans, this question saw the highest confidence in incorrect answers out of any of the 39 questions. This suggests that Republican opinion leaders were successful in convincing a broad swath of the public that House Democrats had rigged the process against the President to a larger extent than was actually the case. Though the typical partisan difference is better-characterized as a knowledge gap, in this case it appeared to be Democrats' relatively greater avoidance of misleading information or heuristics that drove the partisan difference.

The exception to the second half of the formulation, "regardless of party," comes on the question on which the partisan certainty gap has its largest impact on partisan belief differences: Trump's claim of unlimited Article II power. Here, the average Republican who answered incorrectly was 0.09 points more certain than the average Democrat who answered incorrectly. This works out to 45 percent more certain on the 0.5 to 1 scale.[8] The typical incorrect response was still not an outright misperception — the Article II question generally saw lower certainty than the other controversy questions, ranking 11th out of 14 in overall average certainty. Instead, Republicans' relatively greater confidence in their inference that Trump would not make such a claim to power counteracted partisans' general lack of knowledge about this specific fact.

Among the twenty questions that concerned truths that were convenient for Republicans, the tendency for the partisan certainty gap's largest effects to be concentrated among controversy questions was less pronounced. Ranked in terms of the certainty gap's effect on the partisan belief difference, controversy questions made up five of the top eight and just two of the bottom twelve. The controversy questions sitting at the top of this list were true events that are politically inconvenient for Democrats: 2016 presidential candidate Clinton's use of a private email server to send and receive classified information, a Bernie Sanders volunteer's shooting of a Republican member of

[8] (0.798-0.705)/(0.705-0.5)=0.448

Congress and others, and President Obama's reversal of his position that the Deferred Action for Parents of Americans (DAPA) program would amount to "ignoring the law." Least-affected by partisan certainty gaps was a question about the false claim that Robert Mueller's special counsel investigation found that Trump had personally colluded with Russia to influence the 2016 election.

The Republican-congenial questions included the most notable exception to the finding that removing for certainty bias reduces estimates partisan belief differences. On the question about Clinton's emails, the partisan certainty gap boosted the belief difference from 0.02 to 0.08 in the second survey and from 0.06 to 0.10 in the third. Both of these differences still lag behind the overall average difference of 0.108. Examining only the partisan difference in best guesses would suggest that Democrats overwhelmingly "know" that Clinton used the private server. Removing certainty bias reveals that the partisan divide over this controversy is closer to the middle of the pack than prevailing practices would suggest.

These examples demonstrate that incorporating uncertainty is not simply about shrinking partisan differences and disparaging the role of misperceptions. Instead, it is about painting a more nuanced and accurate picture of where belief differences are largest, and where they are best-characterized as products of knowledge gaps or misperceptions. Through their bias toward seeing the same problems around every corner, prevailing practices dull the observer's sense of where these problems are most severe.

Implications

The major implication of this chapter is that current research projects distort partisan differences in factual beliefs, and do so in a manner that is distinct from more-studied concerns about expressive responding. The net result is to substantially change the picture of partisan belief differences.

To give this repeated metaphor some more life, Figure 6.7 plots two such pictures: one that treats best guesses as a measure of belief (left panel) and one based on the probabilistic measure of beliefs (right panel). In each panel, the x-axis is the average for Democrats and the y-axis is the average for Republicans. As above, questions for which the truth is convenient to Democrats are plotted as blue diamonds, while questions whose incorrect answer is congenial to Republicans are plotted red squares. For good measure, the general political awareness questions analyzed in previous chapters are also included. These naturally fall along the thin, black 45-degree line that would indicate perfect correspondence between Democrats' and Republicans' beliefs.

Figure 6.7: Partisan belief differences by question.

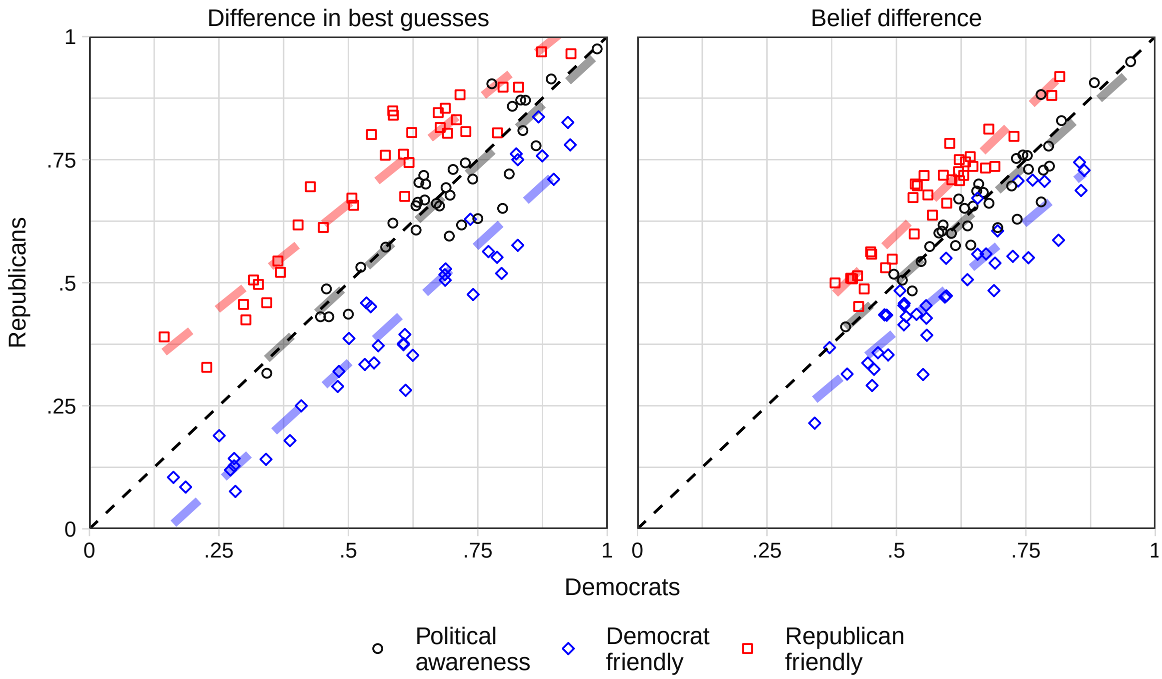

Three striking patterns appear in this figure. First, estimating the partisan belief difference rather than the partisan difference in best guesses substantially shrinks the differences that exist. Relative to the left panel, the points in the right panel almost uniformly shrink toward the 45-degree line, representing the typical — but not universal — tendency for certainty bias to exaggerate the partisan belief difference. The switch to the probabilistic measure preserves the tendency for Democrats' and Republicans' beliefs to be highly correlated across questions.

Second, displaying the data this way highlights the substantial common component to Democrats' and Republicans' beliefs. Even as the points' tendency to fall above or below the 45-degree line illustrates the fact that Democrats and Republicans are more likely to express accurate beliefs when the truth is convenient for their side, the high correlation within each group highlights the fact that Democrats are also more likely to know things that Republicans know, and be ignorant of things that Republicans do not know. This is easy not to see in most analysis of partisans' factual beliefs, which tends to obsess over differences in means without pausing to dwell on the fact that a 10 percentage point difference implies a 90 percentage point overlap.[9]

Third, comparing the two panels, one notices that the switch from best guesses to beliefs does more to shift the points away from the left side of the x-axis than from the right side. This reflects

[9] A 10 percentage point difference is about double the typical partisan difference in best guesses (Roush and Sood 2020).

the finding that certainty-biased measures exaggerate the role of misperceptions by more than they exaggerate the role of knowledge gaps.

Although the finding that partisan differences are better-attributed to knowledge gaps than misperceptions is an outlier in both academic and public discourse over partisan belief differences, it is consistent with a great deal of existing evidence. For example, scholars have found that evidence that misperceptions tend to be bipartisan when they exist (Graham 2020), that media coverage predicts greater knowledge of both convenient and inconvenient truths (Jerit and Barabas 2012, Table 1), and that partisan differences in factual beliefs are larger among elites than among the general public primarily due to greater knowledge of convenient truths (Flynn, Lee, Nyhan and Reifler 2019). False claims are a real threat when people believe them, but for any given partisan difference, ignorance of inconvenient truths is more consistent with the evidence than belief in lies.

Not visible in Figure 6.7 is the final major takeaway from this chapter's findings: systematic differences across question categories. The difference in results between questions about the economy and about politicized controversies demonstrated that accounting for uncertainty is necessary not only for understanding the magnitude and nature of partisan perceptual differences, but for understanding the contours of those differences across questions. By examining a wider range of topics than is typically featured in studies of partisan belief differences, this chapter showed that the certainty bias in conventional measures is not content-neutral. Instead, certainty bias has a substantially larger effect on estimates of differences in economic perceptions than on estimates of differences over politicized controversies.

The finding that the degree of certainty bias is not content-neutral suggests that measuring respondents' uncertainty about their beliefs is necessary to fully understand the challenges that the United States and other democracies face today. Seen through the lens of the traditional American politics focus on policy and performance, disagreement over issues like Obama's birthplace comes as an odd curiosity that may not affect the fundamentals. Yet comparative research on democratic backsliding in polarized societies highlights the destructive role of delegitimizing the other side and turning a blind eye one's own side's violations of democratic norms. Seen in this light, rumors and controversies that undermine the legitimacy of one's opponents are a sign of decay that lays the groundwork for more severe steps to consolidate power. As Mickey, Levitsky and Way (2017) note, "Parties that view their rivals as illegitimate are more likely to resort to extreme measures to weaken them" (27; also see Levitsky and Ziblatt 2018; McCoy and Somer 2019). The current president's claims that his predecessor (Obama) was born outside the country and should have been constitutionally barred from holding office, or that his 2016 opponent (Clinton) deserves to be imprisoned for her use of a private email server, are not merely oddities — they are political weapons deployed as part of a case that the opposition party is not a legitimate competitor. In the context of democratic breakdowns around the world, these divides and U.S. politicians' increasingly brazen attempts to tilt the electoral playing field are two stanzas in a familiar tune. Though the American politics literature's traditional focus on economic performance was era-appropriate, observers of contemporary politics have recognized a growing threat to democratic stability in the United States and around the world.

In theory, the voting public is not defenseless in such situations: just as voters who want a strong economy can punish incumbents who preside over weak performance, voters who value democracy can punish incumbents who violate democratic norms at the polls. Nalepa et al. (2019) show that partisan differences in perceptions of the incumbent's actions should be expected to weaken this accountability: to the extent that partisan voters remain blind to their own side's misdeeds, they cannot punish those misdeeds at the polls. This article's finding of a partisan certainty gap over instances of executive aggrandizement — Trump's claim of unlimited Article II power, the impeachment proceedings against Trump, and Obama's willingness to cross his own lines on the legality of executive actions to protect unauthorized immigrants — suggests such gaps may exist in practice, further undermining the capacity of polarized publics to hold elected leaders accountable for attempts to rig the system in their favor (Graham and Svolik 2020; Svolik 2020).

More could be said, and tested, about these links than is within the scope of this chapter. What has been shown here is that standard surveys overstate the magnitude of partisan differences, misdiagnose the nature of these differences, and hide the influence of forces that drive important variation across topics. Given public opinion surveys' growing centrality in portraits of politics from the academy to public-facing journalism, this is no mean set of flaws. At a time when concerns over polarization, misinformation, and democratic erosion are running higher in the United States than at any point in recent memory, prevailing survey practices manage to both exaggerate and distort portraits of these problems.

Chapter 7

Conclusion

This volume opens with the observation that while a great deal of scholarship embraces a threshold notion of belief, the only sense in which empirical scholarship proves that respondents generally believe their answers to survey questions is probabilistic. The empirical chapters then provide concrete, new evidence as to how this distinction applies to measures of misperceptions and partisan belief differences. These chapters show that the uncertainty in individuals' probabilistic beliefs is not simply a nuisance to be shrugged off. Instead, face value interpretations of survey data paint a distorted picture of the prevalence of misperceptions, the size and nature of partisan belief differences, and the substantive topics on which these pathologies are most prevalent.

Perhaps disappointing relative to the aspirations of pioneering pollsters like George Gallup, these findings lessons for better-interpreting the findings from existing research, for understanding the implications for citizen competence, and for how public-facing and academic survey researchers should collect and present their data.

First, misperceptions are neither as widespread or as deeply held as they appear in surveys. Respondents are quite willing to guess, even when an option to say DK is provided; when such an option is available, only a subset of the least certain respondents take it. The fact that partisanship, or some other attitude, predicts endorsement of the false claim does not generally indicate that those from the group who endorse it more often believe it more strongly. Instead, incorrect answers are about equally strongly believed, regardless of partisanship.

Second, partisan belief differences are generally smaller than they appear. Survey practices that only elicit respondents' best guesses are not neutral vis-a-vis the respondents' uncertainty — instead, they treat respondents as if they were completely certain. Consequently, surveys exaggerate partisan belief differences for about 40 percent, for reasons that are conceptually and empirically distinct from the more-studied issue of expressive responding.

Third, when partisan belief differences exist, they are generally better-attributed to differential knowledge of convenient and inconvenient truths. This suggests that partisans' selective exposure to information is systematically leading people to encounter more information that is good for their side, without encountering equal volumes of information that are good for the other side. Of course, selective exposure can also lead to greater exposure to false claims. But as a rule, a partisan belief difference is not evidence that this has occurred.

Fourth, standard surveys do not just exaggerate the extent of misperceptions and partisan belief differences, but also dull the observer's sense of where they are most pronounced. Because respondents' uncertainty about their beliefs varies across questions, some misperceptions may be about as common as they look, while others may be dramatically over-stated. When it comes to partisan belief differences, partisan differences in certainty can even cause standard practices to *under*-estimate the extent of the divide. The "knowledge gaps" framing for partisan differences has exceptions too.

Together, all of these takeaways suggest that the underlying problem identified in this volume is not that misperceptions or partisan belief differences are a total illusion. These are real problems that, as we saw in the introduction, sometimes appear to have real consequences. Instead, the problem is that surveys lead researchers to see the same problems around every corner. With little independent basis to verify who believes or does not believe their response, and in what sense they do or do not believe it, researchers have an enormous amount of freedom in interpreting the nature of incorrect answers and partisan divides. Too often, this freedom is used to support the most pessimistic possible interpretation of data.

For the larger issue of citizen competence, this all suggests that partisan divides over matters of fact are not as grave or systematic a threat as surveys generally make them out to be. The public's beliefs are characterized by a larger dose of ignorance, and correspondingly less systematic belief in falsehoods, than conventional wisdom suggests. This still leaves us a long way away from the ideal of a perfectly-informed citizen. Yet scholars have known for decades that this ideal is not realistically attainable. Understanding the nature of the public's inevitable ignorance is crucial for understanding the public's capacity to hold politicians accountable for their performance and actions. The conclusions that partisans remain systematically ignorant of inconvenient truths, and merely give credence to rather than fully accept false claims, leave us in a second-best world. But observers should be a bit less pessimistic about the nature of our second-best world than survey research often suggests.

Survey researchers have an obligation to try to get this right. At present, survey measures of misperceptions and partisan belief differences help feed profoundly pessimistic narratives about the public's capacity for democratic citizenship. Yet to the extent that researchers mismeasure misperceptions, it is they who are encouraging members of the public to believe "destructive, provably wrong stuff" (Rampell 2016) about their fellow citizens.

In light of these findings, researchers who want to do better can take three concrete steps. First, researchers can actively incorporate measures of respondents' confidence *in their answers* into portraits of the public's beliefs. DK options simply don't cut it. Second, researchers can limit their interpretations of public opinion data to those that can be verified independently of our theories of who is likely to believe what. These hunches may often be right, but this volume showed how often they can be wrong. Third, researchers who want to take their interpretation beyond the limitations raised in this volume should affirmatively prove that their favorite question, or measurement technology, captures the sort of belief or belief difference they are imagining. Just as Chapter 4 deployed this volume's strategies in a horse race between two measurement technologies, researchers can use the strategies here to prove that they can do better than the current state of the art.

Researchers who do not abide by these recommendations risk doing substantial social harm. Today, when new rumors and false claims gain prominence, polls soon follow, with a resultant message that some double digit percentage of people has heard and accepted the crazy claim — lumping together those who express uncertainty about their incorrect answer and those who (dubiously, as this volume showed) claim to fully believe it (e.g., Rampell 2016; Schaeffer 2020). Given the role that familiarity with a rumor plays in the willingness to endorse or believe it (Berinsky 2017*b*), such coverage may have the effect of helping the rumor spread. At very least, it promotes a profoundly pessimistic message about the public's capacity for rational thought and democratic citizenship, without proper evidence that this message is true. The findings in this volume suggest that such accounts should be viewed as a form of misinformation.

Until the day when someone affirmatively proves that some measurement technology captures misperceptions as they have traditionally been defined, researchers should substantially scale back their messaging about belief in falsehoods and apparently-consequent partisan belief differences. We live in a time in which political divisions run high and many question the public's fitness as democratic citizens. Even as this makes the most pessimistic possible interpretations of survey data seem more plausible, it raises the stakes for researchers who care about their social impact. Anyone throwing fuel on this fire ought to be sure they have their measurement right.

Bibliography

Achen, Christopher H. 1975. "Mass Political Attitudes and the Survey Response." *The American Political Science Review* 69(4):1218–1231.

Ahler, Douglas J. 2014. "Self-Fulfilling misperceptions of public polarization." *Journal of Politics* 76(3):607–620.

Ahler, Douglas J. and Guarav Sood. 2018. "The Parties in Our Heads: Misperceptions about Party Composition and Their Consequences." *The Journal of Politics* 80(3).

Alesina, Alberto, Armando Miano and Stefanie Stantcheva. 2020. "The Polarization of Reality." *AEA Papers and Proceedings* 110:324–328.

Allen, Franklin. 1987. "Discovering Personal Probabilities When Utility Functions Are Unknown." *Management Science* 33(4):542–544.

Alvarez, R Michael and Charles H Franklin. 1994. "Uncertainty and Political Perceptions." *The Journal of Politics* 56(3):671.

Ansolabehere, Stephen, Jonathan Rodden and James M. Snyder. 2008. "The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting." *American Political Science Review* 102(2):215–232.

Ansolabehere, Stephen, M. Meredith and Erik Snowberg. 2013. "Asking About Numbers: Why and How." *Political Analysis* 21(1):48–69.

Atkeson, Lonna Rae and Alex N. Adams. 2018. Mixing Survey Modes and Its Implications. In *The Oxford Handbook of Polling and Survey Methods*, edited by Lonna Rae Atkeson and Michael R. Alvarez. Oxford University Press.

Bartels, Larry M. 2002. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24(2):117–150.

Berelson, B.R., P.F. Lazarsfeld and W.N. McPhee. 1954. *Voting*. Chicago: The University of Chicago Press.

Berinsky, Adam J. 1999. "The Two Faces of Public Opinion." *The American Journal of Political Science* 43(4):1209–1230.

Berinsky, Adam J. 2017*a*. "Measuring Public Opinion with Surveys." *Annual Review of Political Science* 20(1):309–329.

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

Berinsky, Adam J. 2018. "Telling the Truth about Believing the Lies? Evidence for the Limited Prevalence of Expressive Survey Responding." *The Journal of Politics* 80(1):211–224.

Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-In-Differences Estimates?" *Quarterly Journal of Economics* 119(1):249–275.

Bishop, George F., Robert W. Oldendick, Alfred J. Tuchfarber and Stephen E. Bennett. 1980. "Pseudo-Opinions on Public Affairs." *Public Opinion Quarterly* 44(2):198–209.

Bullock, John G, Alan S Gerber, Seth J Hill and Gregory A Huber. 2015. "Partisan Bias in Factual Beliefs about Politics." *Quarterly Journal of Political Science* 10:1–60.

Bullock, John G and Gabriel Lenz. 2019. "Partisan Bias in Surveys." *Annual Review of Political Science* 22:325–342.

Clifford, Scott and Jennifer 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, Pamela Johnston, Stanley Feldman and Kathleen Knight. 1986. "Judging Inflation and Unemployment: The Origins of Retrospective Evaluations." *The Journal of Politics* 48(03):565.

Converse, Philip E. 1964. The nature of belief systems in mass publics. In *Ideology and Discontent*, edited by D. Apter. New York: Free Press.

Converse, Philip E. 1970. Attitudes and Non-Attitudes: Continuation of a Dialogue. In *The Quantitative Analysis of Social Problems*, edited by Edward R. Tufte. Reading, MA: Addison-Wesley pp. 168–189.

Converse, Philip E. 1974. "Comment: The Status of Nonattitudes." *The American Political Science Review* 68(2):650–660.

Converse, Philip E. 2000. "Assessing the Capacity of Mass Electorates." *Annual Review of Political Science* 3:331–53.

Converse, Phillip E. 1990. Popular representation and the distribution of information. In *Information and Democratic Processes*, edited by John A. Ferejohn and James H. Kuklinski. Urbana: University of Illinois Press pp. 369–388.

Delli Carpini, Michael X. and Scott Keeter. 1993. "Measuring political knowledge: Putting first things first." *American Journal of Political Science* 37(4):1179–1206.

Delli Carpini, Michael X. and Scott Keeter. 1996. *What Americans Know About Politics and Why it Matters*. New Haven: Yale University Press.

Diem, Susan J., John D. Lantos and James A. Tulsky. 1996. "Cardiopulmonary Resuscitation on Television — Miracles and Misinformation." *New England Journal of Medicine* 334:1578–1582.

Dodd, Stuart C. and Kaare Svalastoga. 1952. "On estimating latent from manifest undecidedness: the "don't know" percent as a warning of instability among the knowers." *Educational and Psychological Measurement* 12(3):467–471.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.

Ducharme, Wesley M. and Michael L. Donnell. 1973. "Intrasubject comparison of four response modes for 'subjective probability' assessment." *Organizational Behavior and Human Performance* 10(1):108–17.

Eagly, Alice H. and Shelly Chaiken. 1993. *The Psychology of Attitudes*.

Ecker, Ullrich K. H., Stephan Lewandowsky, Ee Pin Chang and Rekha Pillai. 2014. "The Effects of Subtle Misinformation in News Headlines." *Journal of Experimental Psychology Applied* 20(4):323–35.

Erikson, Robert S. 1979. "The SRC Panel Data and Mass Political Attitudes." *British Journal of Political Science* 9(1):89–114.

Evans, Geoffrey and Mark Pickup. 2010. "Reversing the Causal Arrow: The Political Conditioning of Economic Perceptions in the 2000–2004 U.S. Presidential Election Cycle." *The Journal of Politics* 72(04):1236–1251.

Evans, Geoffrey and Robert Andersen. 2006. "The political conditioning of economic perceptions." *Journal of Politics* 68(1):194–207.

Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven: Yale University Press.

Flynn, D. J., Brendan Nyhan and Jason Reifler. 2017. "The Nature and Origins of Misperceptions: Understanding False and Unsupported Beliefs About Politics." *Political Psychology* 38(682758):127–150.

Flynn, D.J. 2016. "The Scope and Correlates of Political Misperceptions in the Mass Public." (Unpublished manuscript).

Flynn, D.J., Nathan Lee, Brendan Nyhan and Jason Reifler. 2019. "More Accurate Yet More Polarized? Comparing the Factual Beliefs of Government Officials and the Public." (Unpublished manuscript).

Foley, Richard. 1992. "The Epistemology of Belief and the Epistemology of Degrees of Belief." *American Philosophical Quarterly* 29(2):111–124.

Gaines, Brian J, James H Kuklinski, Paul J Quirk, Buddy Peyton and Jay Verkuilen. 2007. "Same Facts, Different Interpretations: Partisan Motivation and Opinion on Iraq.".

Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78(1):35–71.

Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23(6):121–123.

Gilens, Martin. 1999. *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*. Chicago: University of Chicago Press.

Gilens, Martin. 2012. Citizen Competence and Democratic Governance. In *New Directions in Public Opinion*, edited by Adam J. Berinsky. 1 edited by New York: Routledge pp. 52–76.

Glaeser, Edward L. and Gergely Ujhelyi. 2010. "Regulating misinformation." *Journal of Public Economics* 94(3-4):247–257.

Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73(20):247–259.

Graham, Matthew H. 2020. "Self-Awareness of Political Knowledge." *Political Behavior* 42:305–326.

Graham, Matthew H and Milan W Svolik. 2020. "Democracy in America? Partisanship, Polarization, and the Robustness of Support for Democracy in the United States." *The American Political Science Review* 114(2):392–409.

Greenland, Sander, Stephen J. Senn, Kenneth J. Rothman, John B. Carlin, Charles Poole, Steven N. Goodman and Douglas G. Altman. 2016. "Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations." *European Journal of Epidemiology* 31(4):337–350.

Grigorieff, Alexis, Christopher Roth and Diego Ubfal. 2020. "Does Information Change Attitudes Towards Immigrants? Representative Evidence from Survey Experiments." *Demography* 57:1117–1143.

Grossman, Guy, Soojong Kim, Jonah M. Rexer and Harsha Thiurmuthy. 2020. "Political partisanship influences behavioral responses to governors' recommendations for COVID-19 prevention in the United States." *Proceedings of the National Academy of Sciences* (forthcoming).

Guess, Andrew and Alexander Coppock. 2018. "Does Counter-Attitudinal Information Cause Backlash? Results from Three Large Survey Experiments." *British Journal of Political Science* forthcomin.

Guess, Andrew M. 2020. "(Almost) Everything in Moderation: New Evidence on Americans' Online Media Diets." *American Journal of Political Science* (forthcoming).

Guess, Andrew M., Brendan Nyhan and Jason Reifler. 2020. "Exposure to untrustworthy websites in the 2016 US election." *Nature Human Behavior* .

Hainmueller, Jens, Daniel J Hopkins and Teppei Yamamoto. 2015. "Causal Inference in Conjoint Analysis: Understanding Multi-Dimensional Choices via Stated Preference Experiments." *Poltical Analysis* 22(1):1–30.

Hill, Seth J. 2017. "Learning Together Slowly: Bayesian Learning about Political Facts." *The Journal of Politics* 79(4):1403–1418.

Hochschild, Jennifer L. and Katherine Levine Einstein. 2015. *Do Facts Matter?* Norman: University of Oklahoma Press.

Hofstetter, C. Richard, David C. Barker, James T. Smith, Gina M. Zari and Thomas A. Ingrassia. 1999. "Information, Misinformation, and Political Talk Radio." *Political Research Quarterly* 52(2):353–369.

Holt, Charles A. and Angela M. Smith. 2016. "Belief Elicitation with a Synchronized Lottery Choice Menu That Is Invariant to Risk Attitudes." *American Economic Journal: Microeconomics* 8(1):110–39.

Hopkins, Daniel J., John Sides and Jack Citrin. 2019. "The muted consequences of correct information about immigration." *Journal of Politics* 81(1):315–320.

Huber, Franz. 2009. Belief and Degrees of Belief. In *Degrees of Belief*, edited by Franz Huber and Christoph Schmidt-Petri.

Ichikawa, Jonathan Jenkins and Matthias Steup. 2018. The Analysis of Knowledge. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Jacobson, Gary C. 2010. "Perception, memory, and partisan polarization on the Iraq War." *Political Science Quarterly* 125(1):31–56.

Jerit, Jennifer and Jason Barabas. 2012. "Partisan Perceptual Bias and the Information Environment." *The Journal of Politics* 74(3):672–684.

Jerit, Jennifer, Tine Paulsen and Joshua Aaron Tucker. 2020. "Confident and Skeptical: What Science Misinformation Patterns Can Teach Us About the COVID-19 Pandemic." *SSRN Electronic Journal* pp. 1–26.

Jerit, Jennifer and Yangzi Zhao. 2020. "Political Misinformation." *Annual Review of Political Science* 23:77–96.

Joyce, James. 2003. Bayes' Theorem. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Kaplan, Mark. 1996. *Decision Theory as Philosophy.* Cambridge University Press.

Key, V.O. 1966. *Public Opinion and American Democracy.*

Koriat, Asher, Sarah Lichtenstein and Baruch Fischhoff. 1980. "Reasons for confidence." *Journal of Experimental Psychology: Human Learning and Memory* 6(2):107–118.

Krosnick, Jon A. and Matthew K. Berent. 1993. "Comparisons of Party Identification and Policy Preferences: The Impact of Survey Question Format." *American Journal of Political Science* 37(3):941–64.

Krosnick, Jon A., Neil Malhotra and Urja Mittal. 2014. "Public misunderstanding of political facts: How question wording affected estimates of partisan differences in birtherism." *Public Opinion Quarterly* 78(1):147–165.

Kuklinski, James H, Paul J Quirk, David W Schwieder and Robert F Rich. 1998. "'Just the Facts, Ma'am': Political Facts and Public Opinion Source." *The Annals of the American Academy of Political and Social Science* 560(November):143–154.

Kuklinski, James H, Paul J Quirk, Jennifer Jerit, David Schwieder and Robert F Rich. 2000. "Misinformation and the currency of democratic citizenship." *The Journal of Politics* 62(3):790–816.

Kull, Steven, Clay Ramsay and Evan Lewis. 2003. "Misperceptions, the media, and the Iraq War." *Political Science Quarterly* 118(4):569–598.

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

Lang, Corey and Shanna Pearson-Merkowitz. 2015. "Partisan sorting in the United States, 1972–2012: New evidence from a dynamic analysis." *Political Geography* 48(1):119–129.

Langley, G.H. 1933. "Belief." *Philosophy* 8(29):66–76.

Laurison, Daniel. 2015. "The Willingness to State an Opinion: Inequality, Don't Know Responses, and Political Participation." *Sociological Forum* 30(4).

Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, Jonathan L. Zittrain, David M. J. Lazer and Jonathan L. Zittrain. 2018. "The science of fake news." *Science* 359(6380):1094–1096.

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

Levitsky, Steven and Daniel Ziblatt. 2018. *How Democracies Die*. New York: Crown.

Lewandowsky, Stephan, Ullrich K. H. Ecker, Colleen M. Seifert, Norbert Schwarz and John Cook. 2012. "Misinformation and Its Correction." *Psychological Science in the Public Interest* 13(3):106–131.

Lewandowsky, Stephan, Werner G.K. Stritzke, Klaus Oberauer and Michael Morales. 2005. "Memory for Fact, Fiction, and Misinformation: The Iraq War 2003." *Psychological Science* 16(3):190–195.

Lichtenstein, Sarah and Baruch Fischhoff. 1977. "Do those who know more also know more about how much they know?" *Organizational Behavior and Human Performance* 20(2):159–183.

Lopez, Jesse and D Sunshine Hillygus. 2018. "Why So Serious?: Survey Trolls and Political Misinformation." (Unpublished manuscript).

Luna, Karlos and Beatriz Martín-Luengo. 2012. "Con fi dence – Accuracy Calibration with General Knowledge and Eyewitness Memory Cued Recall Questions." 295(August 2011):289–295.

Luskin, Robert C., Guarav Sood and Joshua Blank. 2018. "Misinformation about Misinformation: Of Headlines and Survey Design." (Unpublished manuscript).

Luskin, Robert C and John G Bullock. 2011. "'Don't Know' Means 'Don't Know': DK Responses and the Public's Level of Political Knowledge." *Journal of Politics* .

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

Marietta, Morgan, David C. Barker and Todd Bowser. 2015. "Fact-checking polarized politics: Does the fact-check industry provide consistent guidance on disputed realities?" *Forum (Germany)* 13(4):577–596.

McCoy, Jennifer and Murat Somer. 2019. "Toward a Theory of Pernicious Polarization and How it Harms Democracies: Comparative Evidence and Possible Remedies." *The ANNALS of the American Academy of Political and Social Science* 681(1):234–269.

Mickey, Robert, Steven Levitsky and Lucan A. Way. 2017. "Is America Still Safe for Democracy?" *Foreign Affairs* 96(3):20.

Mondak, Jeffery J. 2000. "Reconsidering the Measurement of Political Knowledge." *Political Analysis* 8(1):57–82.

Mondak, Jeffery J and Belinda Creel Davis. 2001. "Asked and Answered: Knowledge Levels When We Will Not Take 'Don't Know' for an Answer." *Political Behavior* 23(3):199–224.

Mondak, Jeffrey J. 2001. "Developing Valid Knowledge Scales." *American Journal of Political Science* 45(1):224–238.

Mondak, Jeffrey J. and Mary R. Anderson. 2004. "The Knowledge Gap: A Reexamination of Gender-Based Differences in Political Knowledge." *The Journal of Politics* 66(2):492–512.

Mutz, Diana. 1998. *Impersonal Influence.* London: Cambridge University Press.

Nadeau, Richard and Richard G Niemi. 1995. "Knowledge Questions in Surveys." *Public Opinion Quarterly* 59:323–346.

Nalepa, Monika, Georg Vanberg and Caterina Chiopris. 2019. "Authoritarian Backsliding." *Unpublished Manuscript* .

Nyhan, Brendan. 2010. "Why the "Death Panel" Myth Wouldn't Die: Misinformation in the Health Care Reform Debate." *The Forum* 8(1):Article 5.

Nyhan, Brendan. 2020. "Facts and Myths about Misperceptions." *Journal of Economic Perspectives* 34(October):220–236.

Nyhan, Brendan and Jason Reifler. 2010. "When corrections fail: The persistence of political misperceptions." *Political Behavior* 32(2):303–330.

Oaksford, Mike and Nick Chater. 2007. *Bayesian Rationality: The Probabilistic Approach to Human Reasoning.* London: Oxford University Press.

Oliver, Eric and Thomas Wood. 2016. "A new poll shows 52% of Republicans actually think Trump won the popular vote." *The Washington Post* (December 18, 2016).

Oliver, J. Eric and Thomas J. Wood. 2014. "Conspiracy Theories and the Paranoid Style of Mass Opinion." *American Journal of Political Science* 58(4):952–966.

Oskamp, S. 1962. "The relationship of clinical experience and training methods to several criteria of clinical prediction." *Psychological Monographs: General and Applied* 76(28):1–27.

Page, Benjamin I. and Robert Y. Shapiro. 1992. *The Rational Public.* Chicago: University of Illinois Press.

Parker-Stephen, Evan. 2013. "Tides of disagreement: How reality facilitates (and inhibits) partisan public opinion." *Journal of Politics* 75(4):1077–1088.

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

Peterson, Erik and Shanto Iyengar. 2020. "Partisan Gaps in Political Information and Information-Seeking Behavior: Motivated Reasoning or Cheerleading." *The American Journal of Political Science* (forthcoming).

Prasad, Monica, Andrew J. Perrin, Kieran Bezila, Steve G. Hoffman, Kate Kindleberger, Kim Manturuk and Ashleigh Smith Powers. 2009. "'There must be a reason': Osama, Saddam, and inferred justification." *Sociological Inquiry* 79(2):142–162.

Prior, Markus. 2005. "News vs. Entertainment: How Increasing Media Choice Widens Gaps in Political Knowledge and Turnout." *American Journal of Political Science* 49(3):577–592.

Prior, Markus, Gaurav Sood and Kabir Khanna. 2015. "You Cannot be Serious. The Impact of Accuracy Incentives on Partisan Bias in Reports of Economic Perceptions." *Quarterly Journal of Political Science* 10(July):489–518.

Przeworski, Adam. 1999. Minimalist Conception of Democracy: A Defense. In *Democracy's Value*, edited by Ian Shapiro and Casiano Hacker-Cordon. Cambridge: Cambridge University Press pp. 23–55.

Rampell, Catherine. 2016. "Americans — especially but not exclusively Trump voters — believe crazy, wrong things." *The Washington Post* (December 28, 2016).

Ross, Janell. 2015. "How come 53 percent of Republicans think the unemployment rate has risen under Obama?" *The Washington Post* (November 20, 2015).

Roush, Carolyn E. and Guarav Sood. 2020. "A Gap in Our Understanding? Reconsidering the Evidence for Partisan Knowledge Gaps." (Unpublished manuscript).

Schaeffer, Katherine. 2020. "A look at the Americans who believe there is some truth to the conspiracy theory that COVID-19 was planned." *Pew Research Center* (July 24, 2020).

Schaffner, Brian F. and Samantha Luks. 2018. "Misinformation or Expressive Responding?" *Public Opinion Quarterly* 82(1):135–147.

Schuman, Howard and Stanley Presser. 1981. *Questions and Answers in Attitude Surveys.* San Diego: SAGE Publications.

Schwitzgibel, Eric. 2015. Belief. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Scotto, Thomas J., Jason Reifler, David Hudson and Jennifer 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.

Smith, Robin. 2017. Aristotle's Logic. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Smith, Tom W. 1985. Nonattitudes: A Review and Evaluation. In *Surveying Subjective Phenomena*, edited by Charles F. Turner Martin and Elizabeth. Russell Sage Foundation pp. 215–254.

Stalnaker, Robert. 1984. *Inquiry.* Cambridge: MIT Press.

Stanig, Piero. 2013. "Political polarization in retrospective economic evaluations during recessions and recoveries." *Electoral Studies* 32(4):729–745.

Strack, Fritz and Leonard L Martin. 1987. Thinking, Judging, and Communicationg: A Process Account of Context Effects in Attitudes Surveys. In *Social information processing and survey methodology*, edited by Hans J. Hippler, Norbert Schwarz and Seymour Sudman. New York: Springer-Verlag pp. 123–148.

Stroud, Natalie Jomini. 2010. "Polarization and partisan selective exposure." *Journal of Communication* 60(3):556–576.

Sturgis, Patrick, Nick Allum and Patten Smith. 2008. "An experiment on the measurement of political knowledge in surveys." *Public Opinion Quarterly* 72(1):90–102.

Sudman, Seymour, Norman M Bradburn and Norbert Schwarz. 1996. *Thinking About Answers: The Application of Cognitive Processes to Survey Methodology.*

Svolik, Milan W. 2020. "When Polarization Trumps Civic Virtue: Partisan Conflict and the Subversion of Democracy by Incumbents." *Quarterly Journal of Political Science* 15(1):3–31.

Swire, Briony, Adam J. Berinsky, Stephan Lewandowsky and Ullrich K. H. Ecker. 2017. "Processing political misinformation: comprehending the Trump phenomenon." *Royal Society Open Science* 4:1–21.

Taber, Charles S. and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science* 50(3):755–769.

Thorson, Emily. 2015. "Belief Echoes: The Persistent Effects of Corrected Misinformation." *Political Communication* 4609(November):1–21.

Tourangeau, Robert, Lance J. Rips and Kenneth Rasinski. 2000. *The Psychology of Survey Response.* Cambridge University Press.

Trautmann, Stefan T. and Gijs van de Kuilen. 2015. "Belief Elicitation: A Horse Race Among Truth Serums." *The Economic Journal* 125(589):2116–2135.

Vick, Steven. 2002. *Degrees of Belief.* Reston, VA: American Society of Civil Engineers.

Weatherson, Brian. 2005. "Can we do without pragmatic encroachment?" *Philosophical Perspectives* 19(1):417–443.

Weber, Nathan and Neil Brewer. 2003. "The effect of judgment type and confidence scale on confidence-accuracy calibration in face recognition." *The Journal of applied psychology* 88(3):490–499.

Wedgwood, Ralph. 2012. "Outright Belief." *Dialectica* 66(3):309–329.

Wood, Thomas and Ethan Porter. 2018. "The Elusive Backfire Effect: Mass Attitudes ' Steadfast Factual Adherence." *Political Behavior* .

Zaller, John. 1992. *The Nature and Origins of Mass Opinion.* Cambridge: Cambridge University Press.

Zaller, John and Stanley Feldman. 1992. "A Simple Theory of the Survey Response: Answering Questions versus Revealing Preferences." *American Journal of Political Science* 36(3):579–616.

Appendix A

Appendix to Chapter 3

A.1 Supplemental Results

Estimating uncertainty for question-level relationships

The empirical pattern of interest in the chapter is the question-to-question relationship between the rate of DK responses and average certainty among everyone else. To estimate uncertainty about these statistics, it would not be appropriate to assume that the questions are a random sample of some larger population of questions. Instead, what is needed is a procedure that can estimate the statistical error that comes from randomly selecting a sample of respondents to answer the set of questions. Given the selected set of questions, what is the estimated sampling distribution of each statistic for our population of Lucid respondents?

To obtain uncertainty estimates that answer this question, I use the block bootstrap. This procedure is commonly used to compute the standard errors of estimators when each unit of analysis contributes more than one observation (e.g., in panel data (Bertrand et al. 2004) and conjoint experiments (Hainmueller et al. 2015)). To estimate the sampling distribution of each statistic, I (1) resampled respondents with replacement 10,000 times and (2) calculated the statistic in each resample, and (3) recorded each result in a vector. Standard errors are the standard deviation of this vector; 95 percent certainty intervals, the 2.5th and 97.5th percentiles.

Regression tables

For Study 1, Table A.1 displays summary statistics for the question-to-question relationship between the percent saying DK and the average certainty of those who did answer. The intercept and slope columns are α and β from the OLS regression

$$(\text{Average certainty})_q = \alpha + \beta(\%\ \text{DK})_q + \epsilon_q,$$

where q indexes questions. The R^2 columns are R^2 from this same regression. Block bootstrapped 95 percent certainty intervals in parentheses (see above).

Table A.1: Regression Test, Study 1

	Model 1: % DK	Model 2: ln(% DK)	Difference
(Intercept)	0.546	0.323	
	(0.537, 0.555)	(0.308, 0.338)	
DK	-0.409	-0.063	
	(-0.448, -0.371)	(-0.069, -0.057)	
R^2	0.300	0.521	-0.222
	(0.253, 0.348)	(0.449, 0.584)	(-0.272, -0.159)

For Study 2, Table A.2 displays summary statistics for the question-to-question relationship between the percent saying DK and the average certainty of those who did answer. The first row displays the overall relationship for Study 2 and the subsequent rows display the relationship within question category. The intercept and slope columns are α and β from the OLS regression

$$(\text{Average certainty})_q = \alpha + \beta(\%\ \text{DK})_q + \epsilon_q,$$

where q indexes questions. The R^2 columns are R^2 from this same regression. Block bootstrapped 95 percent certainty intervals in parentheses (see above).

Table A.2: Regression Summary, Study 2

	Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
Study 2	0.910	-0.291	0.725	-0.072	0.500	0.592	-0.092
	(0.906, 0.914)	(-0.306, -0.277)	(0.719, 0.731)	(-0.076, -0.069)	(0.463, 0.538)	(0.556, 0.627)	(-0.108, -0.075)
Favorability ratings	0.912	-0.250	0.747	-0.058	0.913	0.928	-0.015
	(0.906, 0.918)	(-0.269, -0.231)	(0.736, 0.757)	(-0.064, -0.054)	(0.883, 0.939)	(0.897, 0.954)	(-0.050, 0.021)
Knowledge (civic trivia)	0.912	-0.411	0.698	-0.077	0.727	0.747	-0.020
	(0.902, 0.921)	(-0.465, -0.358)	(0.678, 0.719)	(-0.088, -0.065)	(0.642, 0.802)	(0.643, 0.832)	(-0.062, 0.025)
Knowledge (public figures)	0.966	-0.350	0.763	-0.071	0.863	0.804	0.060
	(0.958, 0.974)	(-0.380, -0.323)	(0.750, 0.775)	(-0.078, -0.064)	(0.800, 0.917)	(0.751, 0.855)	(0.019, 0.102)
Knowledge (statistics)	0.878	-0.378	0.685	-0.071	0.670	0.638	0.032
	(0.866, 0.891)	(-0.439, -0.324)	(0.665, 0.703)	(-0.083, -0.059)	(0.552, 0.784)	(0.520, 0.751)	(-0.009, 0.070)
Policy (attitudes)	0.946	-0.317	0.757	-0.072	0.835	0.840	-0.005
	(0.942, 0.951)	(-0.340, -0.294)	(0.747, 0.766)	(-0.077, -0.066)	(0.795, 0.870)	(0.794, 0.883)	(-0.032, 0.021)
Policy (perceptions)	0.862	-0.259	0.691	-0.075	0.820	0.843	-0.023
	(0.850, 0.873)	(-0.293, -0.223)	(0.676, 0.706)	(-0.085, -0.065)	(0.739, 0.889)	(0.764, 0.909)	(-0.056, 0.013)

Table A.3: Comparison of Model Fit for Table 3.6.

Model	R^2
Bivariate	0.49960 (0.46258, 0.53840)
Intercept	0.89112 (0.87447, 0.90729)
Slope	0.90505 (0.88914, 0.92018)
Bivariate vs. Intercept	0.39152 (0.35673, 0.42771)
Bivariate vs. Slope	0.40546 (0.36864, 0.44162)
Intercept vs. Slope	0.01394 (0.00841, 0.02100)

Note: This table tests the difference in R^2 for the three columns of Table 3.6.

Table A.4 displays the same information as the "Study 2" row of Table A.2, split by demographic group. The same basic result holds for all groups. Further assurance that the results are not driven by between-person differences in certainty, see the within-person analysis below.

Table A.4: Regression Test by Demographic Characteristics, Study 2

		Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
Trait	Category	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
Age	18-29	0.888 (0.880, 0.897)	-0.300 (-0.332, -0.270)	0.692 (0.681, 0.704)	-0.079 (-0.086, -0.072)	0.458 (0.390, 0.525)	0.565 (0.495, 0.637)	-0.108 (-0.135, -0.079)
	30-39	0.904 (0.896, 0.913)	-0.301 (-0.336, -0.267)	0.716 (0.703, 0.730)	-0.073 (-0.081, -0.065)	0.486 (0.408, 0.567)	0.572 (0.496, 0.650)	-0.087 (-0.120, -0.050)
	40-49	0.912 (0.902, 0.921)	-0.284 (-0.321, -0.245)	0.743 (0.728, 0.758)	-0.062 (-0.071, -0.054)	0.462 (0.372, 0.548)	0.504 (0.416, 0.590)	-0.042 (-0.087, 0.016)
	50-64	0.915 (0.909, 0.922)	-0.259 (-0.285, -0.235)	0.754 (0.744, 0.765)	-0.060 (-0.066, -0.054)	0.455 (0.389, 0.522)	0.522 (0.460, 0.581)	-0.066 (-0.100, -0.032)
	65+	0.924 (0.917, 0.932)	-0.273 (-0.305, -0.239)	0.773 (0.758, 0.787)	-0.050 (-0.057, -0.043)	0.459 (0.373, 0.533)	0.448 (0.375, 0.517)	0.011 (-0.044, 0.075)
Education	Associate's, degree	0.902 (0.888, 0.915)	-0.257 (-0.304, -0.208)	0.752 (0.733, 0.770)	-0.054 (-0.065, -0.044)	0.378 (0.261, 0.490)	0.401 (0.300, 0.496)	-0.023 (-0.077, 0.038)
	Bachelor's degree	0.910 (0.904, 0.918)	-0.299 (-0.324, -0.274)	0.724 (0.713, 0.735)	-0.072 (-0.078, -0.066)	0.510 (0.446, 0.571)	0.587 (0.526, 0.646)	-0.077 (-0.107, -0.045)
	Did not finish high school	0.896 (0.876, 0.914)	-0.206 (-0.282, -0.132)	0.769 (0.731, 0.806)	-0.046 (-0.066, -0.027)	0.299 (0.133, 0.464)	0.245 (0.101, 0.410)	0.054 (-0.033, 0.133)
	Graduate or professional	0.904 (0.894, 0.915)	-0.249 (-0.290, -0.212)	0.753 (0.736, 0.769)	-0.057 (-0.067, -0.048)	0.370 (0.279, 0.471)	0.413 (0.325, 0.507)	-0.044 (-0.085, 0.001)
	High school graduate	0.908 (0.900, 0.915)	-0.290 (-0.318, -0.261)	0.727 (0.714, 0.738)	-0.070 (-0.076, -0.064)	0.446 (0.376, 0.506)	0.533 (0.464, 0.596)	-0.087 (-0.117, -0.052)
	Some college, no degree	0.905 (0.897, 0.913)	-0.266 (-0.295, -0.235)	0.739 (0.726, 0.752)	-0.063 (-0.070, -0.056)	0.449 (0.379, 0.518)	0.512 (0.436, 0.583)	-0.063 (-0.099, -0.027)
Gender	Female	0.902	-0.289	0.716	-0.073	0.466	0.567	-0.101

Table A.4: Regression Results by Demographic Characteristics, Study 2 (continued)

Trait	Category	Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
		Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
		(0.897, 0.908)	(-0.308, -0.268)	(0.708, 0.724)	(-0.078, -0.069)	(0.419, 0.514)	(0.523, 0.611)	(-0.121, -0.079)
	Male	0.917	-0.291	0.740	-0.067	0.521	0.584	-0.062
		(0.912, 0.922)	(-0.311, -0.271)	(0.731, 0.749)	(-0.072, -0.063)	(0.467, 0.575)	(0.532, 0.632)	(-0.086, -0.034)
Hispanic	No	0.910	-0.293	0.725	-0.072	0.502	0.593	-0.091
		(0.906, 0.915)	(-0.309, -0.278)	(0.718, 0.731)	(-0.076, -0.069)	(0.461, 0.540)	(0.554, 0.631)	(-0.109, -0.073)
	Yes	0.900	-0.264	0.739	-0.061	0.402	0.438	-0.035
		(0.888, 0.912)	(-0.308, -0.222)	(0.720, 0.756)	(-0.072, -0.050)	(0.301, 0.498)	(0.340, 0.534)	(-0.076, 0.012)
Income	0 to 25	0.909	-0.287	0.730	-0.069	0.465	0.532	-0.067
		(0.902, 0.916)	(-0.314, -0.260)	(0.719, 0.742)	(-0.075, -0.062)	(0.399, 0.529)	(0.465, 0.596)	(-0.097, -0.033)
	100 to 200	0.899	-0.242	0.761	-0.049	0.357	0.399	-0.042
		(0.888, 0.909)	(-0.282, -0.197)	(0.743, 0.779)	(-0.059, -0.040)	(0.249, 0.451)	(0.303, 0.486)	(-0.089, 0.012)
	200+	0.902	-0.251	0.759	-0.051	0.380	0.414	-0.034
		(0.891, 0.911)	(-0.291, -0.207)	(0.742, 0.776)	(-0.060, -0.042)	(0.280, 0.472)	(0.322, 0.500)	(-0.082, 0.017)
	25 to 50	0.905	-0.281	0.722	-0.072	0.466	0.559	-0.093
		(0.897, 0.912)	(-0.308, -0.253)	(0.711, 0.734)	(-0.078, -0.065)	(0.399, 0.532)	(0.489, 0.621)	(-0.123, -0.059)
	50 to 75	0.910	-0.283	0.736	-0.066	0.437	0.508	-0.071
		(0.902, 0.919)	(-0.315, -0.251)	(0.723, 0.748)	(-0.073, -0.059)	(0.356, 0.512)	(0.430, 0.578)	(-0.105, -0.035)
	75 to 100	0.910	-0.272	0.743	-0.063	0.457	0.514	-0.056
		(0.899, 0.921)	(-0.314, -0.233)	(0.724, 0.761)	(-0.073, -0.052)	(0.357, 0.558)	(0.409, 0.609)	(-0.106, 0.001)
	Missing	0.892	-0.242	0.753	-0.049	0.306	0.289	0.016
		(0.875, 0.911)	(-0.309, -0.177)	(0.726, 0.782)	(-0.064, -0.034)	(0.176, 0.440)	(0.154, 0.434)	(-0.060, 0.097)
Party	Democrat	0.908	-0.317	0.724	-0.066	0.506	0.598	-0.092
		(0.903, 0.914)	(-0.338, -0.298)	(0.714, 0.732)	(-0.071, -0.062)	(0.454, 0.552)	(0.548, 0.645)	(-0.121, -0.061)
	Independent	0.911	-0.290	0.695	-0.101	0.434	0.482	-0.048
		(0.898, 0.927)	(-0.329, -0.252)	(0.678, 0.712)	(-0.114, -0.088)	(0.343, 0.525)	(0.395, 0.575)	(-0.074, -0.021)

Table A.4: Regression Results by Demographic Characteristics, Study 2 (continued)

		Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
Trait	Category	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
	Republican	0.905 (0.900, 0.910)	-0.245 (-0.267, -0.224)	0.760 (0.751, 0.768)	-0.052 (-0.056, -0.048)	0.423 (0.361, 0.483)	0.522 (0.467, 0.574)	-0.099 (-0.128, -0.068)
Race	Asian or Pacific Is.	0.910 (0.895, 0.926)	-0.302 (-0.366, -0.241)	0.750 (0.725, 0.772)	-0.054 (-0.068, -0.042)	0.387 (0.258, 0.510)	0.368 (0.255, 0.479)	0.019 (-0.038, 0.081)
	Black	0.903 (0.893, 0.913)	-0.287 (-0.334, -0.243)	0.724 (0.706, 0.741)	-0.069 (-0.078, -0.059)	0.451 (0.354, 0.542)	0.536 (0.448, 0.623)	-0.084 (-0.128, -0.040)
	Other	0.900 (0.887, 0.912)	-0.241 (-0.283, -0.195)	0.748 (0.730, 0.767)	-0.059 (-0.069, -0.049)	0.395 (0.281, 0.501)	0.468 (0.354, 0.567)	-0.073 (-0.118, -0.023)
	White	0.909 (0.904, 0.913)	-0.284 (-0.301, -0.268)	0.729 (0.722, 0.736)	-0.070 (-0.074, -0.066)	0.479 (0.434, 0.519)	0.563 (0.521, 0.604)	-0.085 (-0.102, -0.064)

Within-subject analysis

The main text explains the relationship between the prevalence of DK responses and average certainty among the "knowers" using a combination of an individual-level factor (a latent threshold for answering questions) and population-level factor (a substantial common component to peoples' knowledge and ignorance). A reasonable reader could wonder whether it is truly variation in the respondents' certainty levels that drives the results. If some other factor affects both respondents' overall average certainty and their propensity to say DK, perhaps the key results could arise simply because more and less certain people say DK at different times.

To rule out the possibility that alternative explanations related to respondents' baseline level of certainty can explain the results, this appendix reproduces many of the key results from Study 2 using only within-respondent variation in certainty. To remove between-person variation in certainty, I calculated each respondent's average certainty level on the questions they answered, then subtracted it from the certainty level the respondent chose alongside each answer. Everything else in the within-subject analysis identical to the main calculations.

Study 1: Figure 1 using only within-respondent variation in certainty

Figure A.1: Percent DK versus average certainty among other respondents, 1993-2000 ANES.

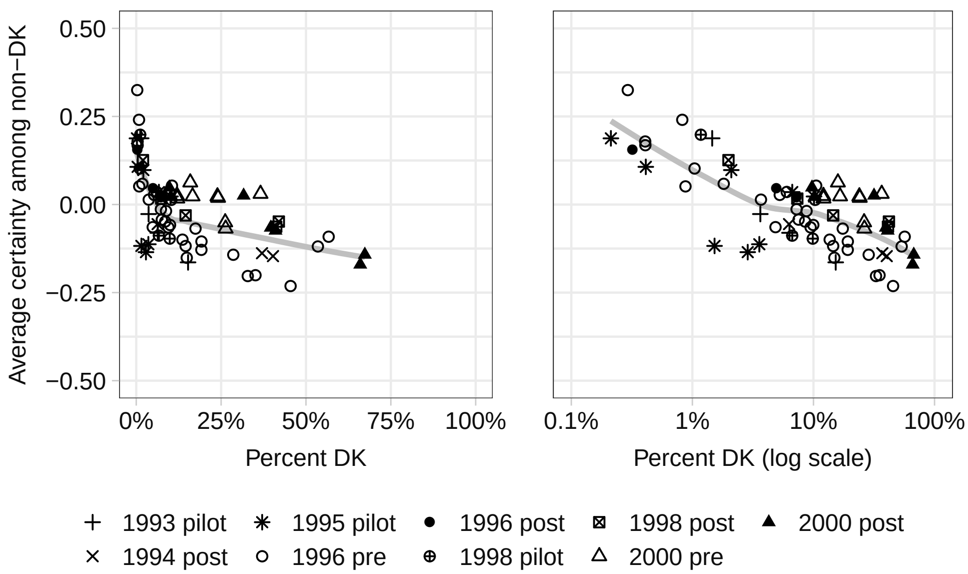

Note: This figure is identical to Figure 3.2, except that it only uses within-respondent variation in certainty.

Study 1: regression using only within-respondent variation in certainty

Table A.5: Regression Test, Within-Respondent, Study 1

	Model 1: % DK	Model 2: ln(% DK)	Difference
(Intercept)	0.043 (0.038, 0.048)	-0.178 (-0.189, -0.165)	
DK	-0.416 (-0.445, -0.389)	-0.062 (-0.067, -0.056)	
R^2	0.325 (0.293, 0.360)	0.519 (0.450, 0.576)	-0.194 (-0.247, -0.130)

Note: This table is identical to Table A.1, except that it only uses within-respondent variation in certainty.

Study 2: Figure 3 using only within-respondent variation in certainty

Figure A.2: Percent DK versus within-respondent variation in certainty, study 2.

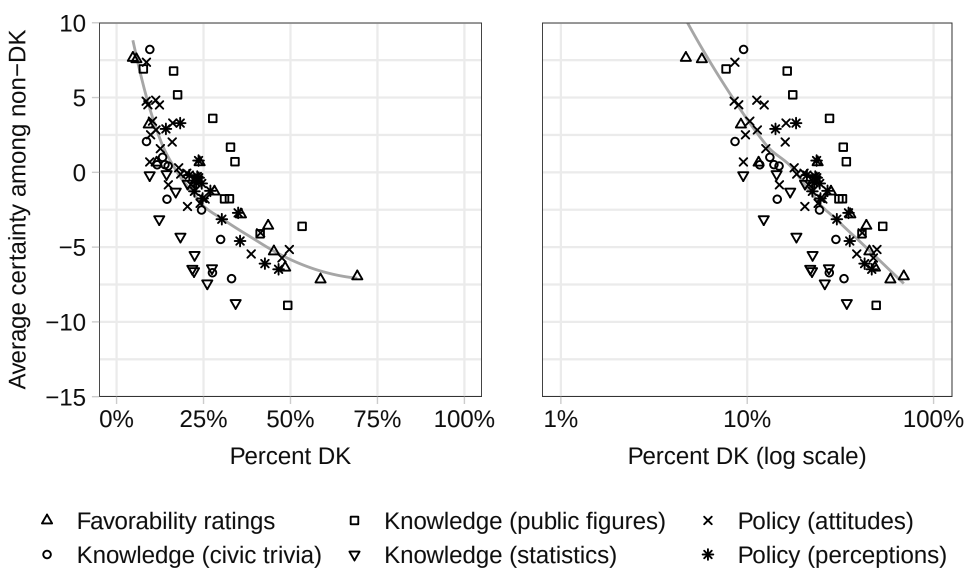

Note: This figure is identical to Figure 3.3, except that it only uses within-respondent variation in certainty.

Study 2: all regression tables using only within-respondent variation in certainty

Table A.6: Regression Summary, Within-Respondent, Study 2

	Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
Study 2	0.052	-0.243	-0.103	-0.060	0.543	0.640	-0.097
	(0.049, 0.054)	(-0.254, -0.233)	(-0.107, -0.099)	(-0.063, -0.058)	(0.513, 0.572)	(0.610, 0.670)	(-0.113, -0.082)
Favorability ratings	0.061	-0.227	-0.091	-0.055	0.876	0.944	-0.068
	(0.056, 0.066)	(-0.243, -0.210)	(-0.098, -0.084)	(-0.059, -0.051)	(0.839, 0.911)	(0.919, 0.966)	(-0.102, -0.032)
Knowledge (civic trivia)	0.070	-0.432	-0.155	-0.081	0.747	0.769	-0.021
	(0.061, 0.079)	(-0.484, -0.385)	(-0.173, -0.137)	(-0.091, -0.070)	(0.668, 0.816)	(0.668, 0.850)	(-0.062, 0.021)
Knowledge (public figures)	0.120	-0.363	-0.091	-0.074	0.874	0.831	0.043
	(0.113, 0.128)	(-0.392, -0.338)	(-0.102, -0.081)	(-0.082, -0.068)	(0.823, 0.922)	(0.786, 0.874)	(0.006, 0.083)
Knowledge (statistics)	0.032	-0.365	-0.154	-0.068	0.675	0.644	0.031
	(0.022, 0.044)	(-0.421, -0.313)	(-0.173, -0.137)	(-0.080, -0.057)	(0.565, 0.778)	(0.534, 0.756)	(-0.011, 0.070)
Policy (attitudes)	0.057	-0.255	-0.100	-0.060	0.771	0.830	-0.059
	(0.053, 0.061)	(-0.275, -0.236)	(-0.108, -0.093)	(-0.064, -0.055)	(0.726, 0.812)	(0.780, 0.876)	(-0.087, -0.029)
Policy (perceptions)	0.066	-0.294	-0.128	-0.084	0.870	0.882	-0.012
	(0.057, 0.075)	(-0.325, -0.263)	(-0.140, -0.115)	(-0.094, -0.075)	(0.810, 0.921)	(0.816, 0.932)	(-0.040, 0.019)

Note: This table is identical to Table A.2, except that it only uses within-respondent variation in certainty.

Table A.7: Regression Test for Differences in Intercepts and Slopes, Within Respondent, Study 2

Term	Model Bivariate	 Variable Intercepts	 Variable Intercepts & Slopes
(Intercept)	0.052 (0.049, 0.054)	0.080 (0.075, 0.084)	0.061 (0.056, 0.066)
DK	-0.243 (-0.254, -0.233)	-0.284 (-0.294, -0.274)	-0.227 (-0.243, -0.210)
Knowledge (civic trivia)		-0.037 (-0.043, -0.031)	0.009 (-0.002, 0.020)
Knowledge (public figures)		0.019 (0.013, 0.025)	0.059 (0.051, 0.068)
Knowledge (statistics)		-0.064 (-0.070, -0.058)	-0.029 (-0.041, -0.016)
Policy (attitudes)		-0.017 (-0.022, -0.012)	-0.005 (-0.011, 0.002)
Policy (perceptions)		-0.016 (-0.022, -0.010)	0.005 (-0.006, 0.016)
DK × Knowledge (civic trivia)			-0.205 (-0.259, -0.154)
DK × Knowledge (public figures)			-0.137 (-0.171, -0.107)
DK × Knowledge (statistics)			-0.138 (-0.198, -0.086)
DK × Policy (attitudes)			-0.028 (-0.055, -0.005)
DK × Policy (perceptions)			-0.068 (-0.105, -0.030)
R^2	0.543 (0.513, 0.572)	0.818 (0.795, 0.839)	0.853 (0.831, 0.873)

Note: This table is identical to Table 3.6, except that it only uses within-respondent variation in certainty.

Table A.8: Comparison of Model Fit for Table A.7.

Model	R^2
Bivariate	0.54316 (0.51253, 0.57244)
Intercept	0.81840 (0.79535, 0.83931)
Slope	0.85267 (0.83092, 0.87258)
Bivariate vs. Intercept	0.27523 (0.24774, 0.30316)
Bivariate vs. Slope	0.30951 (0.27935, 0.33832)
Intercept vs. Slope	0.03428 (0.02389, 0.04578)

Note: This table is identical to Table A.3, except that it only uses within-respondent variation in certainty.

Table A.9: Regression Test by Demographic Characteristics, Within Respondent, Study 2

		Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
Trait	Category	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
Age	18-29	0.051 (0.044, 0.058)	-0.257 (-0.282, -0.234)	-0.116 (-0.125, -0.107)	-0.067 (-0.073, -0.061)	0.498 (0.436, 0.557)	0.595 (0.534, 0.653)	-0.097 (-0.125, -0.067)
	30-39	0.049 (0.042, 0.055)	-0.239 (-0.266, -0.214)	-0.100 (-0.110, -0.090)	-0.058 (-0.064, -0.052)	0.500 (0.427, 0.573)	0.577 (0.503, 0.647)	-0.078 (-0.109, -0.045)
	40-49	0.051 (0.045, 0.057)	-0.230 (-0.256, -0.205)	-0.087 (-0.097, -0.076)	-0.051 (-0.058, -0.045)	0.503 (0.430, 0.577)	0.561 (0.479, 0.639)	-0.058 (-0.103, -0.010)
	50-64	0.048 (0.044, 0.053)	-0.218 (-0.238, -0.199)	-0.087 (-0.095, -0.079)	-0.050 (-0.055, -0.046)	0.481 (0.423, 0.543)	0.546 (0.486, 0.606)	-0.065 (-0.099, -0.028)
	65+	0.050 (0.044, 0.055)	-0.236 (-0.260, -0.211)	-0.079 (-0.090, -0.068)	-0.042 (-0.048, -0.036)	0.519 (0.445, 0.586)	0.487 (0.408, 0.561)	0.031 (-0.024, 0.097)
Education	Associate's, degree	0.043 (0.034, 0.052)	-0.207 (-0.241, -0.173)	-0.076 (-0.090, -0.062)	-0.042 (-0.051, -0.034)	0.407 (0.308, 0.504)	0.410 (0.315, 0.508)	-0.003 (-0.065, 0.061)
	Bachelor's degree	0.049 (0.044, 0.054)	-0.239 (-0.259, -0.220)	-0.100 (-0.108, -0.093)	-0.057 (-0.062, -0.053)	0.529 (0.471, 0.586)	0.604 (0.549, 0.662)	-0.075 (-0.103, -0.048)
	Did not finish high school	0.035 (0.021, 0.048)	-0.160 (-0.215, -0.113)	-0.064 (-0.089, -0.042)	-0.036 (-0.050, -0.023)	0.276 (0.146, 0.415)	0.236 (0.110, 0.376)	0.040 (-0.027, 0.112)
	Graduate or professional	0.047 (0.039, 0.054)	-0.214 (-0.245, -0.185)	-0.081 (-0.095, -0.069)	-0.048 (-0.056, -0.041)	0.375 (0.299, 0.458)	0.405 (0.323, 0.496)	-0.030 (-0.070, 0.010)
	High school graduate	0.049 (0.044, 0.054)	-0.233 (-0.254, -0.211)	-0.096 (-0.105, -0.088)	-0.056 (-0.061, -0.051)	0.479 (0.412, 0.540)	0.570 (0.507, 0.629)	-0.091 (-0.120, -0.058)
	Some college, no degree	0.050 (0.045, 0.056)	-0.237 (-0.260, -0.213)	-0.097 (-0.107, -0.087)	-0.056 (-0.062, -0.050)	0.515 (0.452, 0.575)	0.580 (0.512, 0.645)	-0.066 (-0.104, -0.029)
Gender	Female	0.050 (0.047, 0.054)	-0.244 (-0.259, -0.228)	-0.106 (-0.112, -0.100)	-0.061 (-0.065, -0.057)	0.537 (0.494, 0.578)	0.631 (0.588, 0.672)	-0.094 (-0.115, -0.073)
	Male	0.051	-0.235	-0.093	-0.055	0.522	0.598	-0.076

Table A.9: Regression Results by Demographic Characteristics, Study 2 (continued)

		Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
Trait	Category	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
		(0.047, 0.055)	(-0.251, -0.219)	(-0.099, -0.087)	(-0.059, -0.051)	(0.476, 0.567)	(0.550, 0.642)	(-0.100, -0.049)
Hispanic	No	0.051	-0.241	-0.101	-0.059	0.538	0.632	-0.094
		(0.048, 0.054)	(-0.253, -0.229)	(-0.106, -0.097)	(-0.062, -0.056)	(0.504, 0.573)	(0.601, 0.665)	(-0.111, -0.078)
	Yes	0.048	-0.232	-0.094	-0.054	0.465	0.508	-0.043
		(0.040, 0.056)	(-0.263, -0.201)	(-0.107, -0.082)	(-0.062, -0.045)	(0.379, 0.545)	(0.422, 0.587)	(-0.085, 0.006)
Income	0 to 25	0.050	-0.237	-0.099	-0.057	0.501	0.589	-0.088
		(0.045, 0.055)	(-0.257, -0.218)	(-0.107, -0.091)	(-0.063, -0.052)	(0.440, 0.560)	(0.526, 0.647)	(-0.118, -0.058)
	100 to 200	0.039	-0.206	-0.077	-0.041	0.390	0.420	-0.031
		(0.032, 0.047)	(-0.237, -0.173)	(-0.091, -0.063)	(-0.049, -0.033)	(0.304, 0.471)	(0.331, 0.508)	(-0.076, 0.020)
	200+	0.041	-0.211	-0.079	-0.043	0.405	0.436	-0.031
		(0.035, 0.049)	(-0.243, -0.181)	(-0.092, -0.066)	(-0.050, -0.036)	(0.322, 0.486)	(0.355, 0.517)	(-0.079, 0.017)
	25 to 50	0.050	-0.236	-0.101	-0.059	0.520	0.598	-0.078
		(0.045, 0.055)	(-0.256, -0.214)	(-0.109, -0.093)	(-0.064, -0.054)	(0.465, 0.579)	(0.542, 0.657)	(-0.109, -0.048)
	50 to 75	0.048	-0.227	-0.090	-0.052	0.460	0.521	-0.061
		(0.042, 0.054)	(-0.250, -0.202)	(-0.100, -0.081)	(-0.058, -0.046)	(0.391, 0.529)	(0.449, 0.590)	(-0.099, -0.024)
	75 to 100	0.050	-0.220	-0.085	-0.050	0.435	0.480	-0.045
		(0.042, 0.057)	(-0.250, -0.189)	(-0.097, -0.072)	(-0.058, -0.043)	(0.349, 0.519)	(0.386, 0.567)	(-0.093, 0.008)
	Missing	0.044	-0.217	-0.080	-0.043	0.349	0.327	0.022
		(0.033, 0.058)	(-0.268, -0.170)	(-0.100, -0.059)	(-0.055, -0.031)	(0.242, 0.459)	(0.193, 0.459)	(-0.061, 0.108)
Party	Democrat	0.051	-0.262	-0.103	-0.055	0.539	0.652	-0.113
		(0.047, 0.054)	(-0.279, -0.247)	(-0.109, -0.096)	(-0.059, -0.052)	(0.494, 0.582)	(0.607, 0.695)	(-0.137, -0.088)
	Independent	0.065	-0.236	-0.110	-0.081	0.451	0.488	-0.037
		(0.054, 0.075)	(-0.265, -0.204)	(-0.122, -0.097)	(-0.091, -0.070)	(0.373, 0.534)	(0.412, 0.566)	(-0.062, -0.010)
	Republican	0.042	-0.211	-0.082	-0.045	0.481	0.593	-0.112
		(0.039, 0.046)	(-0.227, -0.194)	(-0.088, -0.077)	(-0.048, -0.042)	(0.430, 0.531)	(0.542, 0.642)	(-0.141, -0.081)

Table A.9: Regression Results by Demographic Characteristics, Study 2 (continued)

		Model 1: % DK		Model 2: ln(% DK)		Comparison of R^2		
Trait	Category	Intercept	Slope	Intercept	Slope	Model 1	Model 2	Difference
Race	Asian or Pacific Is.	0.047 (0.035, 0.058)	-0.240 (-0.285, -0.194)	-0.080 (-0.098, -0.063)	-0.043 (-0.053, -0.034)	0.391 (0.288, 0.494)	0.376 (0.272, 0.478)	0.015 (-0.044, 0.080)
	Black	0.047 (0.039, 0.055)	-0.230 (-0.262, -0.199)	-0.097 (-0.109, -0.085)	-0.055 (-0.063, -0.048)	0.450 (0.366, 0.531)	0.542 (0.462, 0.621)	-0.092 (-0.134, -0.047)
	Other	0.044 (0.035, 0.052)	-0.210 (-0.242, -0.180)	-0.088 (-0.101, -0.075)	-0.050 (-0.059, -0.042)	0.440 (0.347, 0.530)	0.506 (0.408, 0.602)	-0.065 (-0.113, -0.015)
	White	0.050 (0.047, 0.053)	-0.234 (-0.248, -0.222)	-0.098 (-0.103, -0.092)	-0.057 (-0.060, -0.054)	0.519 (0.483, 0.555)	0.604 (0.567, 0.641)	-0.085 (-0.105, -0.068)

Note: This table is identical to Table A.4, except that it only uses within-respondent variation in certainty.

Testing the common knowledge assumption

The threshold model's second key assumption was that despite heterogeneity in political interest and news consumption, there is also a substantial common component to the topics on which the public possesses or lacks a basis to answer survey questions. This section uses data from study 1 to present a two-part validation of this assumption. The first portion of the analysis tests directly for a common component of knowledge, and the second shows that the main results hold within these same patterns.

Measurement

Both portions of the analysis split the results by interest in politics and frequency of news consumption. The political interest question asks, "Some people don't pay much attention to political campaigns. How about you? Would you say that you were very much interested, somewhat interested, or not much interested in following the political campaigns this year?" As the survey questions used in the analysis concerned candidate attributes, candidate policy positions, and respondent policy positions, this wording is apt. Although political interest was not asked in the 1993 and 1995 pilot surveys, most of the respondents also participated in the 1992 and 1994 pre-election surveys, which included the political interest questions. The political interest variables used were V923101 (1993 pilot), V940124 (1994 post and 1995 pilot), V960201 (1996 pre), V961001 (1996 post), V98P101 (1998 pilot), V980201 (1998 post), V000301 (2000 pre), and V001201 (2000 post).

Among the surveys examined, the best coverage on news consumption is a question that asks, "How many days in the past week did you read a daily newspaper?" This question was available usually asked in pre-election surveys, which provides some assurance that campaign effects do not affect the comparability of the measure across years. The newspaper readership questions used are V923203 (1993 pilot), V940125 (1994 post and 1995 pilot), V960246 (1996 pre and post), V98P103 (1998 pilot), V980202 (1998 post), and V000305 (2000 pre and post).

Analysis

First, across levels of both political interest and news consumption, there is a strong common component to respondents' rates of DK responding and average certainty. Figure A.3 presents correlograms using the three-level measure of interest in political campaigns. Figures A.4 and A.5 show that this relationship is even across levels of news consumption.

Figure A.3: Correlograms by interest in political campaigns.

(a) DK rate

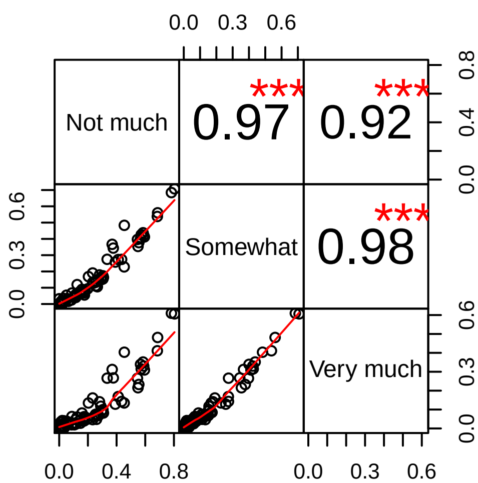

(b) Average certainty, non-DK respondents

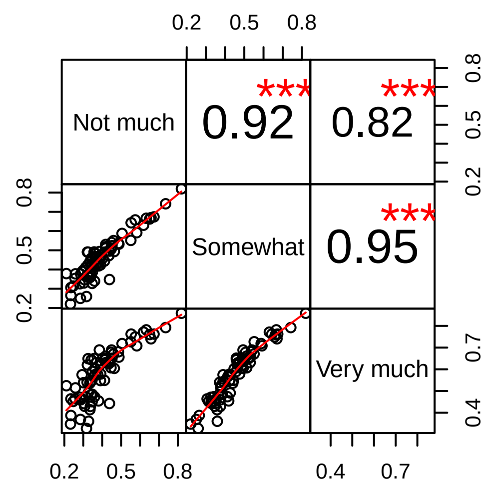

Note: This figure presents pairwise correlations of the rate of DK responding, and of average certainty, across three levels of political interest. Each point represents a survey question. The top-right set of panels prints the correlations and the bottom-left set of panels visualizes them. This figure was created in `R` using `PerformanceAnalysis::chart.Correlation`.

Figure A.4: Correlogram of DK rate by frequency of newspaper readership.

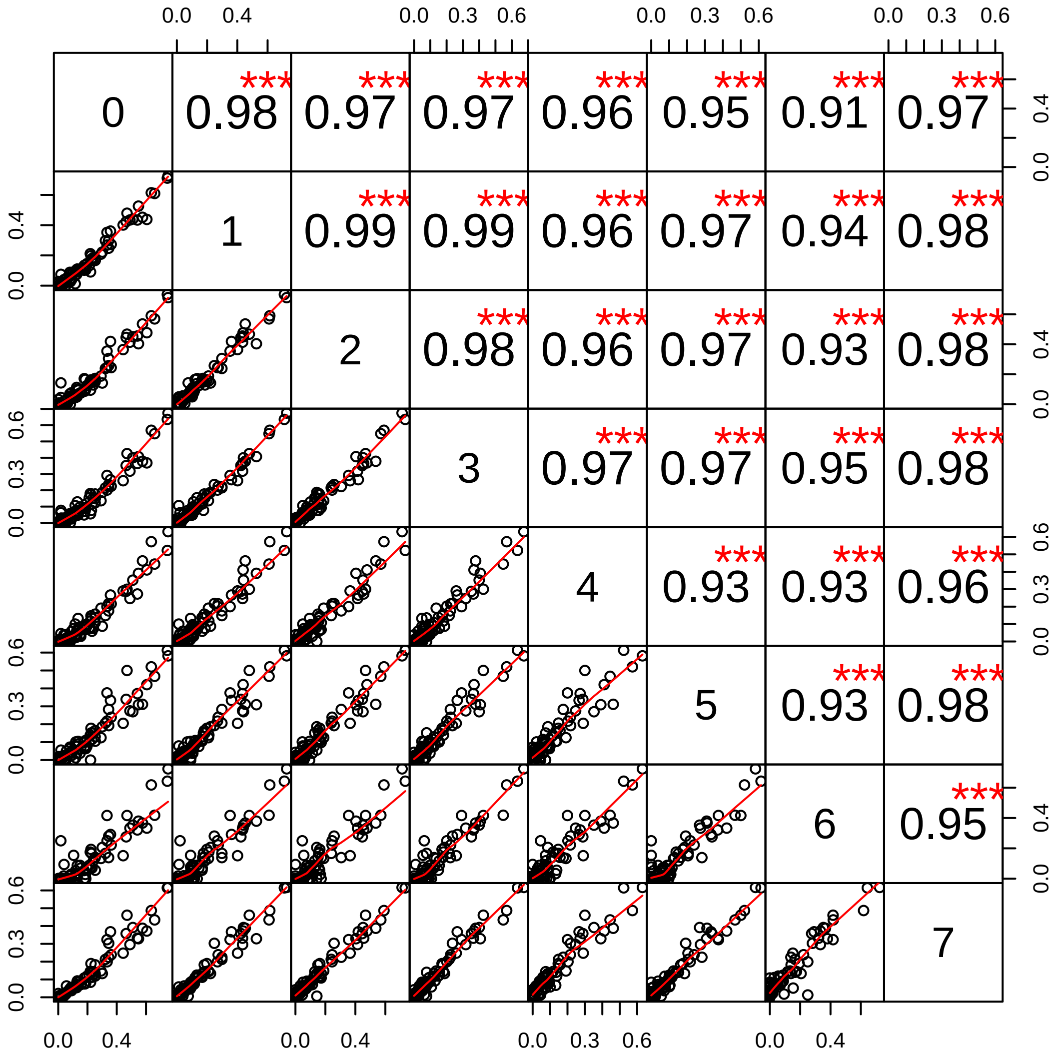

Note: This figure presents pairwise correlations of the rate of DK responding according to the number of days the respondent reported reading a daily newspaper in the past week. Each point represents a survey question. The top-right set of panels prints the correlations and the bottom-left set of panels visualizes them. This figure was created in `R` using `PerformanceAnalysis::chart.Correlation`.

Figure A.5: Correlogram of average certainty by frequency of newspaper readership.

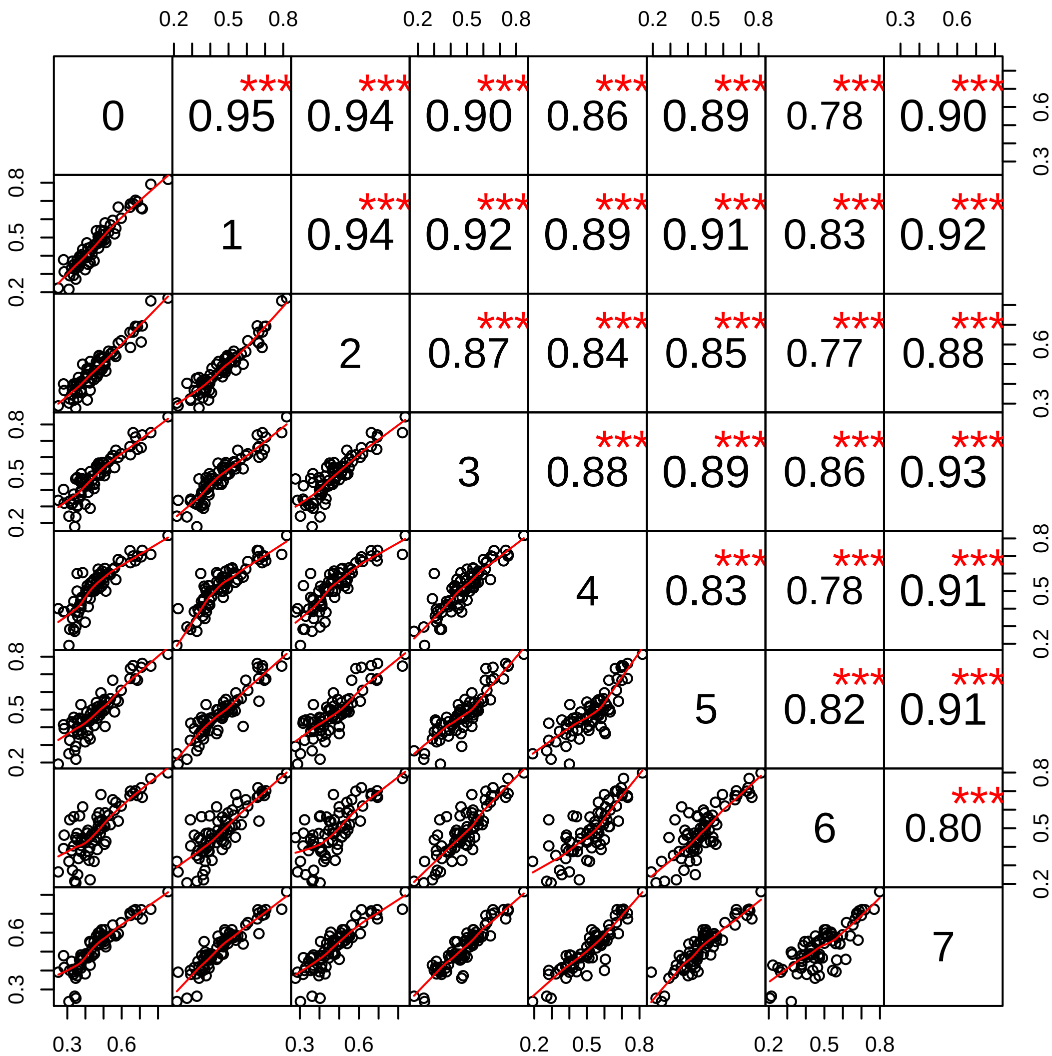

Note: This figure presents pairwise correlations of the rate of average certainty according to the number of days the respondent reported reading a daily newspaper in the past week. Each point represents a survey question. The top-right set of panels prints the correlations and the bottom-left set of panels visualizes them. This figure was created in `R` using `PerformanceAnalysis::chart.Correlation`.

Second, I show that within levels of political interest and news consumption, the paper's main results hold. Figure A.6 plots the relationship separately for the three levels of political interest, while figure A.7 conducts the same analysis using newspaper readership terciles.

Figure A.6: Percent DK versus average certainty among other respondents, by level of political interest, 1993-2000 ANES.

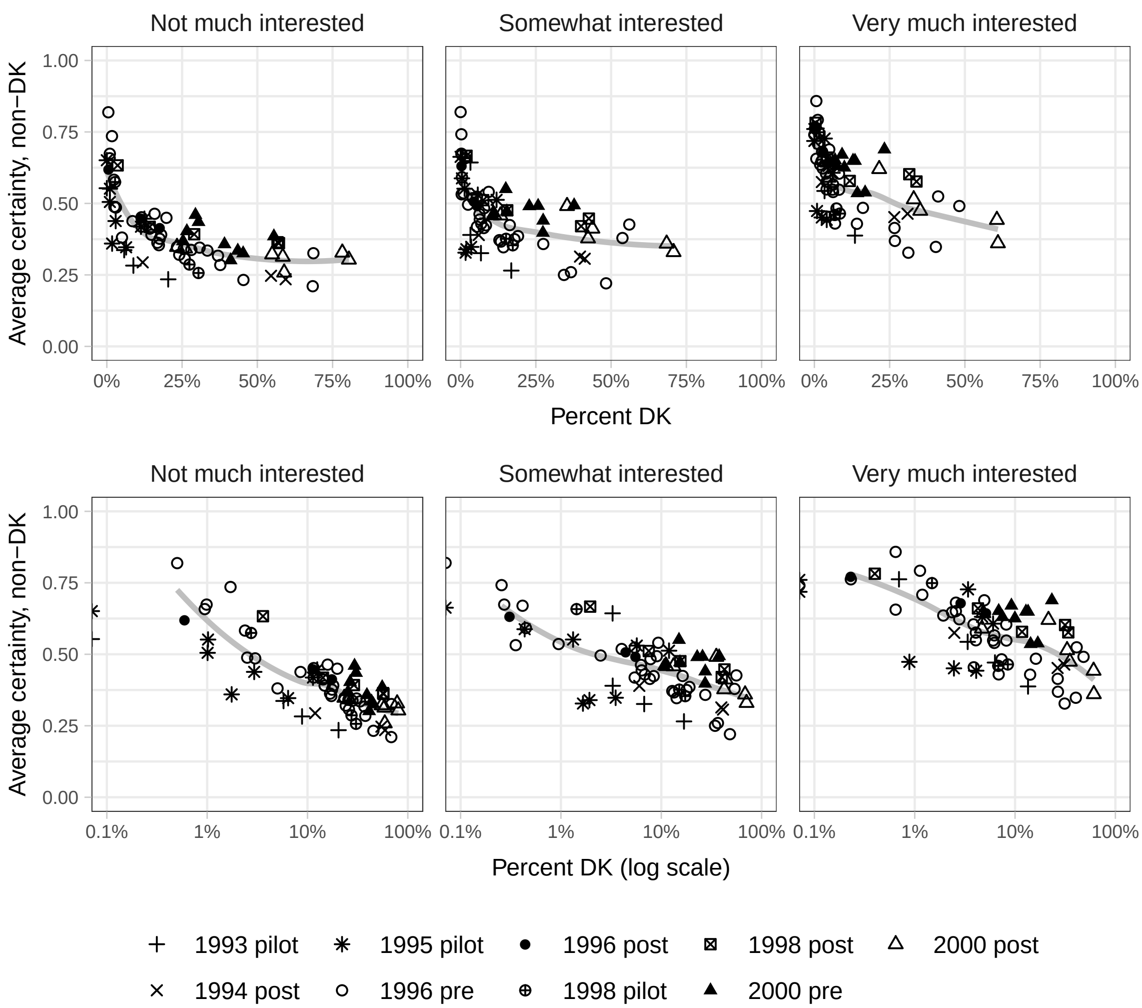

Note: This figure is identical to Figure 3.2 in the main text, but with the results split by the respondent's level of political interest. The top row corresponds to the left panel, and the bottom row corresponds to the right panel.

Figure A.7: Percent DK versus average certainty among other respondents, by number of days read a daily newspaper in the past week, 1993-2000 ANES.

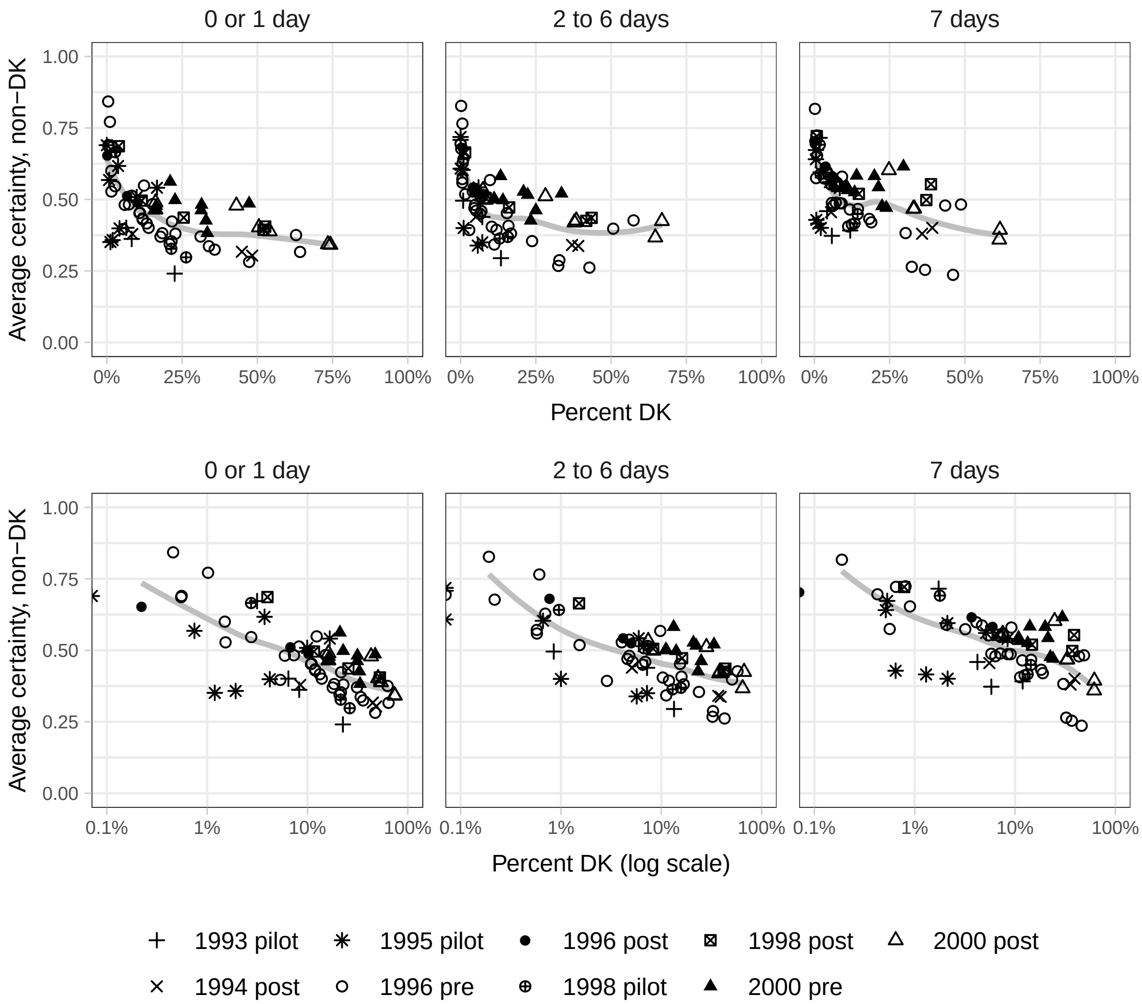

Note: This figure is identical to Figure 3.2 in the main text, but with the results split by the number of days the respondent read a daily newspaper in the past week. The top row corresponds to the left panel, and the bottom row corresponds to the right panel.

Estimating the certainty level of DK respondents

The main text uses an estimator of the certainty levels that DK respondents would have stated if they had answered the question. This section proves this estimator.

Let $E[C|\cdot]$ be average certainty conditional on some other variable. Let $DK \in \{0, 1\}$ be an indicator variable for choosing DK. Let $A \in \{0, 1\}$ be an indicator of whether a DK response is allowed.

Assume that allowing a DK response option does not affect respondents' certainty in their inference about the question; instead, it only affects how they express their level of certainty. Formally, this assumption is that

$$E[C|A = 1] = E[C|A = 0] \tag{A.1}$$

The quantity of interest is $E[C|DK = 1, A = 1]$: the average certainty level of DK respondents when DK responses are allowed.

Begin by rewriting $E[C|A = 1]$, i.e. average certainty when DK is allowed.

$$\begin{aligned} E[C|A = 1] = {} & E[C|DK = 0, A = 1]Pr[DK = 0|A = 1] + \\ & E[C|DK = 1, A = 1]Pr[DK = 1|A = 1] \end{aligned} \tag{A.2}$$

$$E[C|DK = 1, A = 1] = \frac{E[C|A = 1] - E[C|DK = 0, A = 1]Pr[DK = 0|A = 1]}{Pr[DK = 1|A = 1]} \tag{A.3}$$

Now use (A.1) to substitute the first term in the numerator, and note that $Pr[DK = 0|A = 1] = 1 - Pr[DK = 1|A = 1]$.

$$E[C|DK = 1, A = 1] = \frac{E[C|A = 0] - E[C|DK = 0, A = 1](1 - Pr[DK = 1|A = 1])}{Pr[DK = 1|A = 1]} \tag{A.4}$$

This is identical to expression (3.1) in the main text, which substitutes brief verbal explanations for the variables (e.g., "said DK" instead of $DK = 1$, "DK allowed" instead of $A = 1$).

To estimate statistical uncertainty about this estimator, I used the block bootstrap, which is described in section A.1.

Simulation study

This section briefly supplements the main studies with a simulation analysis based on the February 2020 Lucid and March 2020 MTurk surveys. These questions were similar in format to Study 2's knowledge questions, except that DK responses were not allowed and refusals were strongly discouraged.

Approach

The intuition behind these simulations is that the oft-noted pressure to produce a response in interviewer-administered surveys may be thought of as setting a relatively low certainty threshold for answering the question, while the relative ease of saying DK online can be viewed as setting the threshold higher. An opinion filter can be viewed the same way: a lack of a filter sets a low threshold, while using a filter that tells people not to answer unless they are sure sets a higher threshold.

If this intuition is correct, performing simulations that compare a low and a high threshold should reproduce both the patterns seen in this study, as well as other patterns observed across these survey modes. I focus on previous studies that allow comparisons between DK rates on the phone versus online (Atkeson and Adams 2018) and with and without an opinion filter (Bishop et al. 1980).

To capture this, I compare two thresholds: a low threshold, 51 percent certainty, and a higher threshold, 60 percent. The higher threshold is set at 60 percent to match this paper's finding that the average online respondent who says DK would have stated 60 percent certainty. The 51 percent threshold is an arbitrarily lower number, representing the intuition that one will guess on the phone if one even has the slightest idea about the question. Below, I relax the assumption that all respondents have the same threshold.

Results

Figure A.8a plots the percentage of responses that the model predicts would be converted into DKs. With both a low and a higher threshold, greater certainty means fewer DKs. However, because high-threshold technologies cast a wider net, they "convert" uncertainty into DK responses at a higher rate. This gives rise to the more general property observed by both Atkeson and Adams (2018) and Bishop et al. (1980): the more DKs are captured by a low-threshold measure, the more an increase in the threshold will increase the DK rate.
To show this more directly, Figure A.8b more directly reproduces the originally observed relationship, plotting the "conversion rate" (y-axis) against the DK rate on the lower-threshold technology. The conversion rate is the percentage of substantive responses that switch to DK when the threshold increases. Mathematically, this is given by the expression

$$\frac{E[DK|\text{higher threshold} - E[DK|\text{lower threshold}]}{1 - E[DK|\text{lower threshold}]}.$$

Just as in the original studies, Figure A.8b suggests that when the DK rate using the lower-threshold technology is higher, a larger share of substantive responses are eliminated by a threshold-increasing design choice.

Low thresholds' tendency to "convert" responses into DKs at a lower rate can also explain why the shape of the "we don't know" means "they're not sure" relationship differed between Study 1 and Study 2. Figure A.8c plots the predicted relationship at a low threshold and a high threshold; the x- and y-axes have the same interpretation as Figures 2, 3, 4, and 5.

The low-threshold simulation predicts that low-threshold technologies will see a sharp drop-off at lower rates of DK. This is consistent with Study 1, which used interviewer-administered surveys.

Figure A.8: Predicted differences between survey technologies under the threshold model.

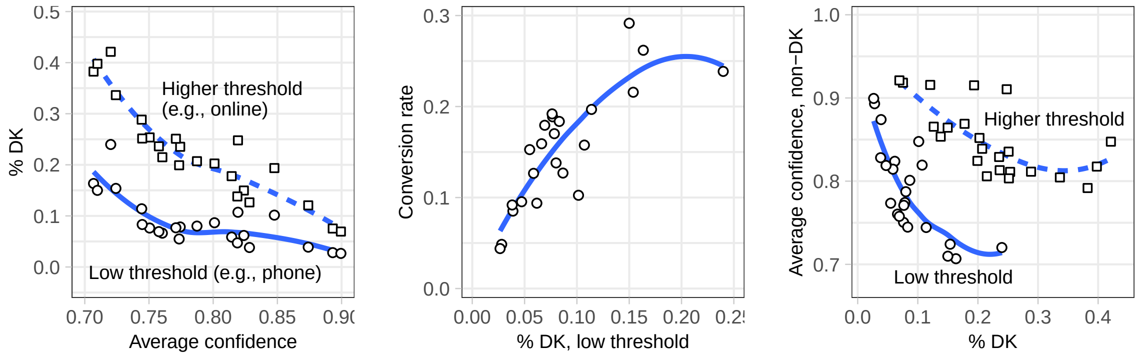

(a) Predicted DK rate by average certainty.

(b) Predicted difference in DK conversion rate, low vs. higher threshold.

(c) Predicted "we don't know" means "they're not sure" relationship.

Meanwhile, the higher threshold sees a more gradual relationship, consistent with the online surveys analyzed in Study 2.

Varying the Threshold

Though the threshold model is easier to illustrate under the assumption that all respondents use the same certainty threshold, this is not likely to be the case. The main text provides two reasons to think that thresholds might vary: the discussion of Mondak's findings on differences in guessing behavior, and the finding that DK responses filter out a few responses that would have been stated with 70 percent certainty or higher.

In one sense, it does not matter if the threshold varies; to whatever extent it does or does not, the "we don't know" means "they're not sure" relationship exists. But to build better understanding, it is worth checking robustness to heterogeneity in the threshold.

To examine how heterogeneity in the certainty threshold might affect the results, I re-created each panel in Figure A.8 under the following eleven alternative assumptions:

- **Small random differences.** Uniformly distributed noise from -5 to 5 is added to respondent's thresholds.
- **Large random differences.** Uniformly distributed noise from -10 to 10 is added to respondent's thresholds.
- **Large random differences, reversed.** The same uniformly distributed noise is added, but is multiplied by -1.
- **Small knowledge-based.** The threshold varies as a function of the number of questions the respondent answered correctly. The more knowledge, the higher the threshold. Each respondent's threshold is equal to the overall threshold (51 or 60) plus 50 percent of the difference between their number correct and the average number.
- **Small knowledge-based, reversed.** Same as above, but reversed. Less knowledge means a higher the threshold.
- **Large knowledge-based.** The threshold varies as a function of the number of questions the respondent answered correctly. The more knowledge, the higher the threshold. Each respondent's threshold is equal to the overall threshold (51 or 60) plus 100 percent of the difference between their number correct and the average number.

- **Large knowledge-based, reversed.** Same as above, but reversed. The less knowledge, the higher the threshold.
- **Small certainty-based.** The threshold varies as a function of the respondent's average certainty level. The more certainty, the higher the threshold. Each respondent's threshold is equal to the overall threshold (51 or 60) plus 50 percent of the difference between their certainty and the average.
- **Small certainty-based, reversed.** Same as above, but reversed. Less certainty means a higher the threshold.
- **Large certainty-based.** The threshold varies as a function of the respondent's average certainty level. The more certainty, the higher the threshold. Each respondent's threshold is equal to the overall threshold (51 or 60) plus 100 percent of the difference between their certainty and the average.
- **Large certainty-based, reversed.** Same as above, but reversed. The less certainty, the higher the threshold.

The figures below reproduce each panel of Figure A.8 under these assumptions. In all cases, the relationship is robust to heterogeneity in the threshold. The final figure shows that the "we don't know" means "they're not sure" relationship should still be expected even when the threshold varies and only within-respondent variation in certainty is used.

Figure A.9: Predicted DK rate by average certainty, variable threshold scenarios.

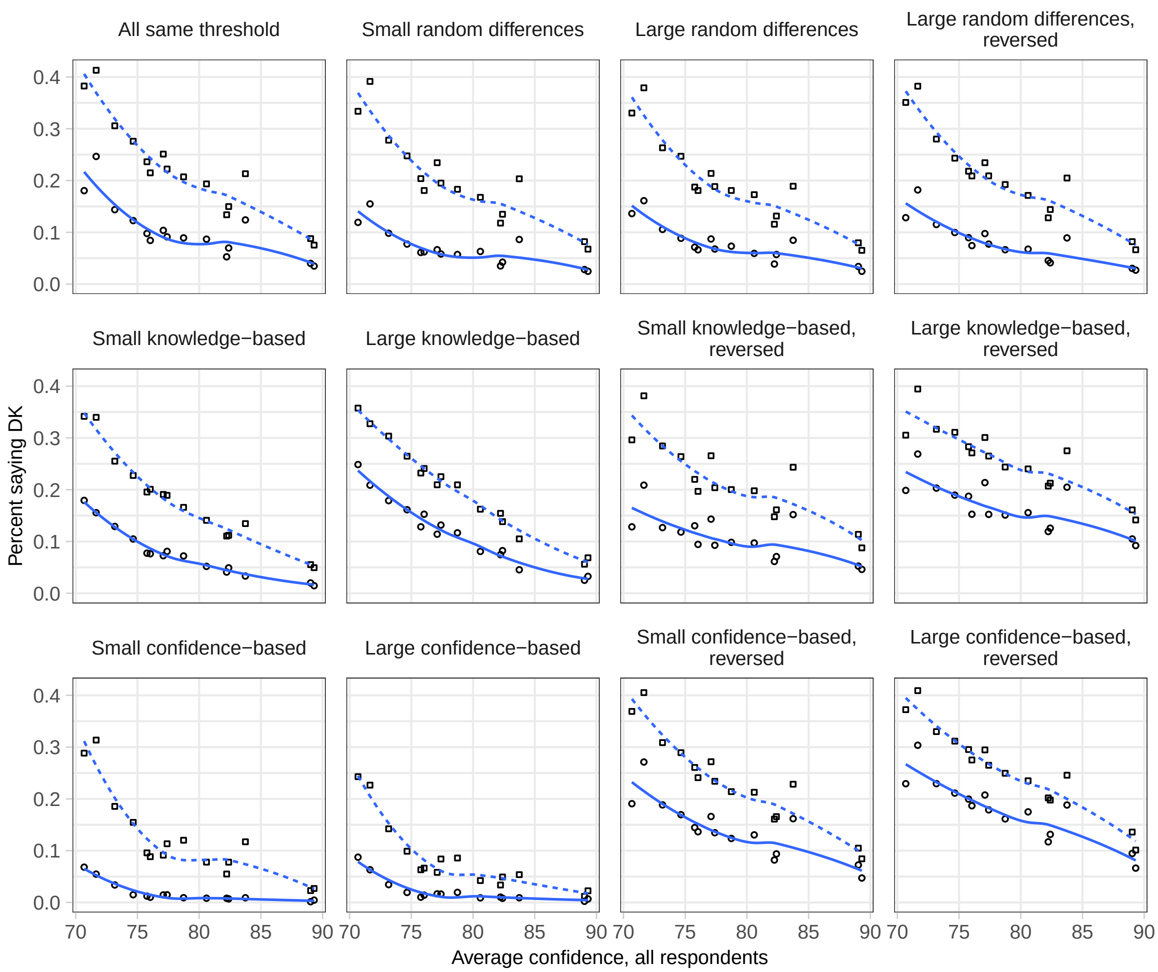

Figure A.10: Predicted difference in DK conversion rate, low vs. higher threshold, variable threshold scenarios.

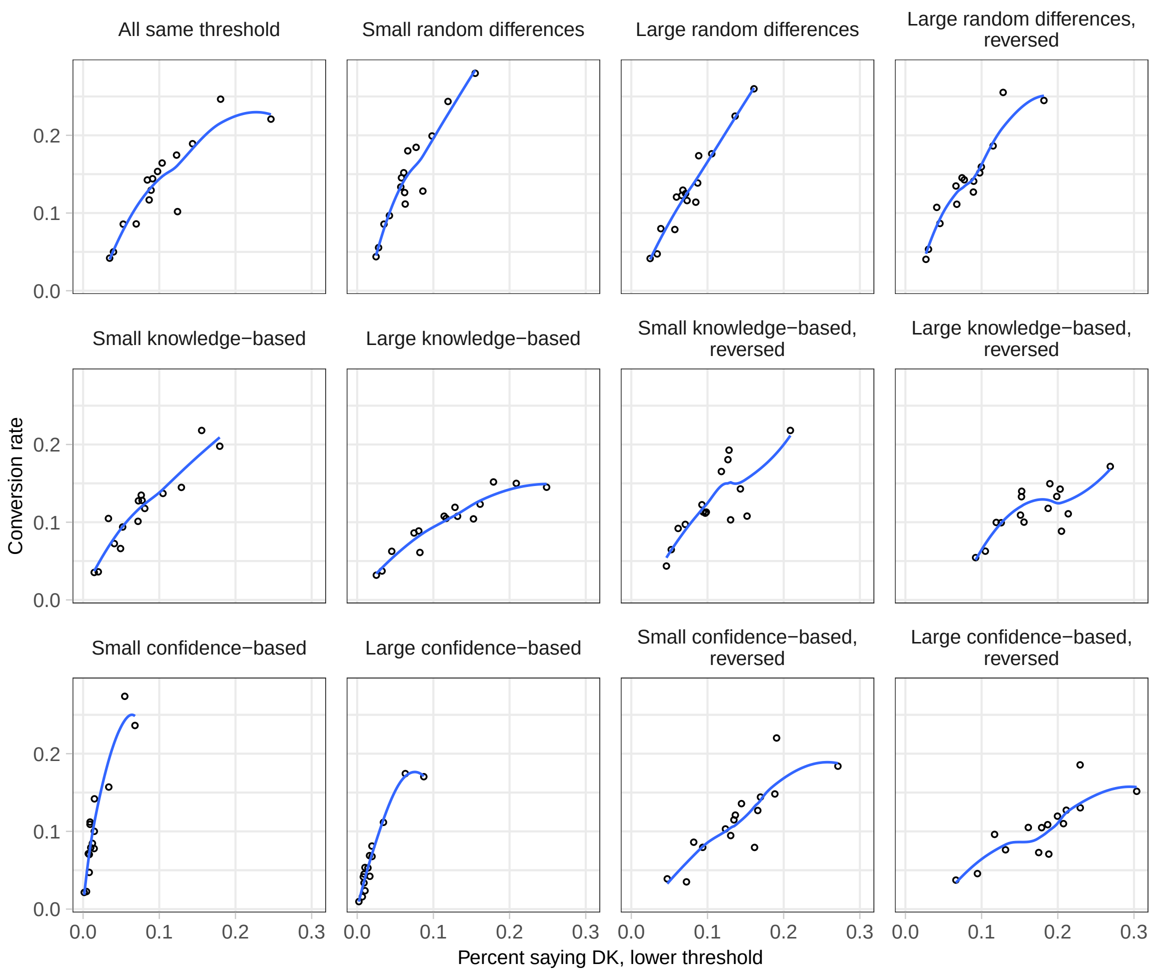

Figure A.11: Predicted "we don't know" means "they're not sure" relationship, variable threshold scenarios.

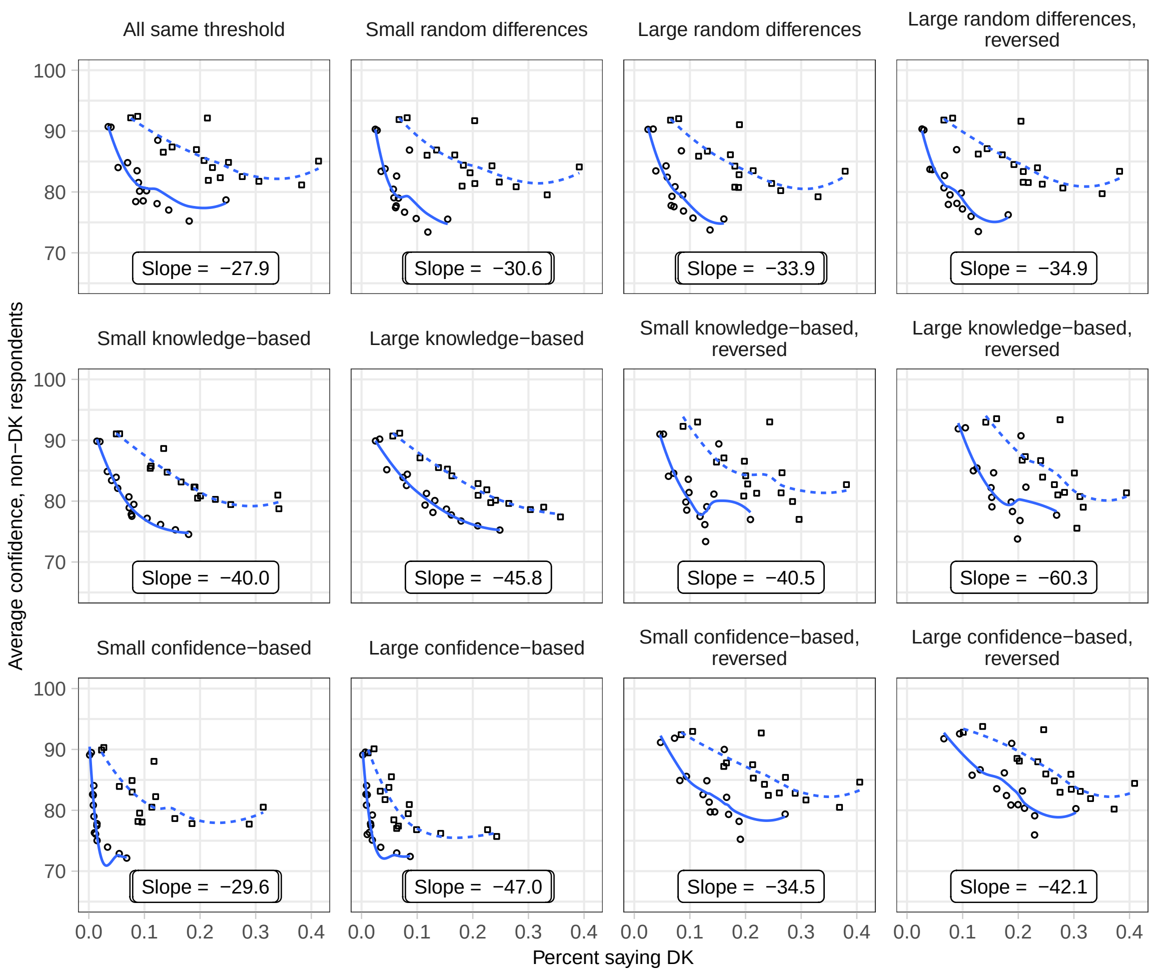

Figure A.12: Predicted "we don't know" means "they're not sure" relationship, within-respondent variation only, variable threshold scenarios.

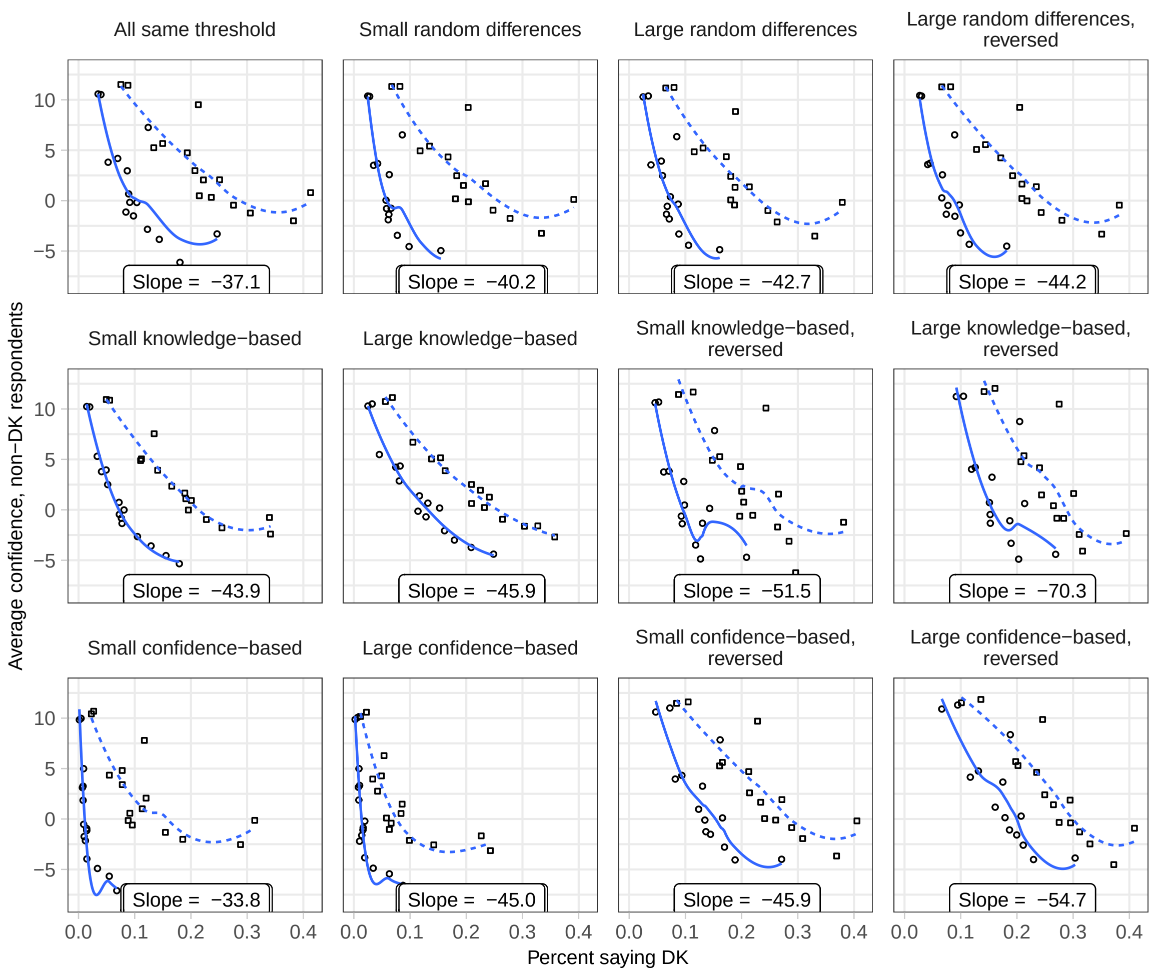

Reanalysis of Atkeson and Adams (2018) and Bishop et al. (1980)

The main text briefly discusses the consistency between the certainty threshold model and a result seen in both Atkeson and Adams (2018) and Bishop et al. (1980): in response to a change in survey technologies that can be interpreted as raising the threshold, a larger proportion of substantive responses should be "converted" to DK on questions that were high on DK to begin with. The previous section's simulation study demonstrated this prediction (Figures A.8b and A.10). This section confirms that the relationship existed in the original studies through a reanalysis of their data.

Bishop et al. (1980) conducted a split ballot experiment that featured seven attitudinal questions. Though they do not directly report the relationship between the DK rate and the conversion rate, it can be backed out of their Tables 1 and 2. Figure A.13a displays the results. Just as noted in the main text, a higher percentage DK without an opinion filter predicts a higher conversion rate due to the opinion filter.

Atkeson and Adams (2018) describe a survey that was conducted partly over the phone and partly online. They report that on questions that were high-DK on the phone, there was an especially large increase in the DK rate with the move online. I contacted the authors and obtained cross-tabulations for all of the questions related to perceptions of election integrity and voter fraud.

Figure A.13b displays the results. Just as noted in the main text, a higher percentage DK on the phone predicts a higher conversion rate with the move online.

Figure A.13: Reanalysis of previous studies.

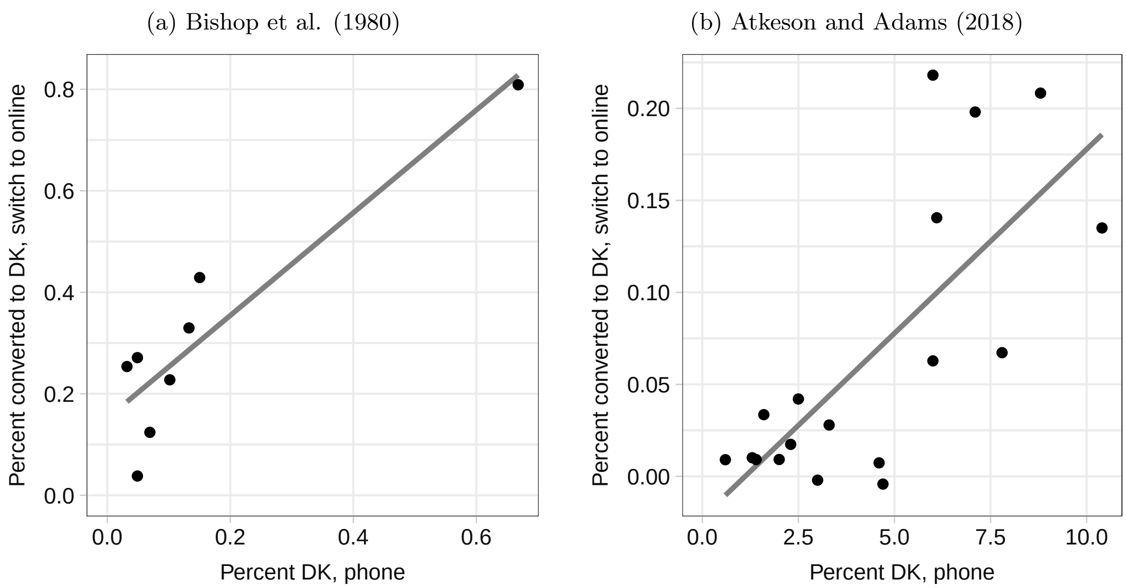

A.2 Survey Information

Study 1 question text

Table A.10 lists all of the ANES questions used in the analysis. The source file for the 1992-1997 surveys was the 1992-1997 merged file. The weight variables used to analyze the election-year surveys were V923008, V940004, and V960003. The source file for the 1998 post-election survey was the 1998 time series file. The weight variable used to analyze this survey was V980002. The source file for the 1998 pilot was the 1998 pilot file. No weight variable appeared in this file. The source file for the 2000 pre- and post-election surveys was the 2000 time series file. The weight variables used to analyze this survey were WT00PRE and WT00PO.

Table A.10: List of ANES questions

Survey	Initial Q	Certainty Q	Description
1993 Pilot	937204	937208	Liberal/conservative self-placement
1993 Pilot	937209	937210	* Clinton liberal/conservative placement
1993 Pilot	937214	937216	* Clinton liberal/conservative placement
1993 Pilot	937220	937221	Perot liberal/conservative placement
1993 Pilot	937220	937221	House incumbent liberal/conservative placement
1994 Post	940841	940842	Clinton liberal/conservative placement
1994 Post	940843	940844	House Dem. candidate liberal/conservative placement
1994 Post	940843	940844	House Rep. candidate liberal/conservative placement
1995 Pilot	952072	952073	Clinton strong leader
1995 Pilot	952074	952075	Clinton moral
1995 Pilot	952076	952077	Dole strong leader
1995 Pilot	952078	952079	Dole moral
1995 Pilot	952190a	952191	Environmental regulation self-placement
1995 Pilot	952194a	952195	Clinton environmental regulation placement
1995 Pilot	952204a	952205	Senator #1 environmental regulation placement
1995 Pilot	952108a	952109	Senator #2 environmental regulation placement
1996 Pre	960365	960367	Liberal/conservative self-placement
1996 Pre	960369	960370	Clinton liberal/conservative placement
1996 Pre	960371	960372	Dole liberal/conservative placement
1996 Pre	960373	960374	Perot liberal/conservative placement
1996 Pre	960375	960376	House Dem. candidate liberal/conservative placement
1996 Pre	960377	960378	House Rep. candidate liberal/conservative placement
1996 Pre	960423	960424	Clinton moral
1996 Pre	960430	960431	Clinton gets things done
1996 Pre	960432	960433	Dole moral
1996 Pre	960439	960440	Dole gets things done
1996 Pre	960441	960442	Perot moral
1996 Pre	960448	960449	Perot gets things done
1996 Pre	960450	960451	Spending/services self-placement
1996 Pre	960453	960454	Clinton spending/services placement
1996 Pre	960455	960456	Dole spending/services placement
1996 Pre	960457	960458	Perot spending/services placement
1996 Pre	960463	960464	Defense spending self-placement
1996 Pre	960466	960467	Clinton defense spending placement
1996 Pre	960469	960470	Dole defense spending placement
1996 Pre	960472	960473	Perot defense spending placement

Table A.10: List of ANES questions (continued)

Survey	Initial Q	Certainty Q	Description
1996 Pre	960487	960488	Aid to blacks self-placement
1996 Pre	960490	960491	Clinton aid to blacks placement
1996 Pre	960492	960493	Dole aid to blacks placement
1996 Pre	960494	960495	Perot aid to blacks placement
1996 Pre	960503	960504	Abortion self-placement
1996 Pre	960506	960507	Clinton abortion placement
1996 Pre	960509	960510	Dole abortion placement
1996 Pre	960512	960513	Perot abortion placement
1996 Pre	960523	960524	Environment/jobs self-placement
1996 Pre	960526	960527	Clinton environment/jobs placement
1996 Pre	960529	960530	Dole environment/jobs placement
1996 Pre	960532	960533	Perot environment/jobs placement
1996 Post	961269	961271	Liberal/conservative self-placement
1996 Post	961273	961274	Clinton liberal/conservative placement
1996 Post	961275	961276	Dole liberal/conservative placement
1998 Pilot	98P291	98P293	* Liberal/conservative self-placement
1998 Pilot	98P294	98P296	* Liberal/conservative self-placement
1998 Pilot	98P297	98P298	Governor candidate #1 liberal/conservative placement
1998 Pilot	98P299	98P300	Governor candidate #2 liberal/conservative self-placement
1998 Post	980399	980400	Liberal/conservative self-placement
1998 Post	980403	980404	Clinton liberal/conservative placement
1998 Post	980405	980406	Gore liberal/conservative placement
1998 Post	980407	980408	House Dem. candidate liberal/conservative placement
1998 Post	980409	980410	House Rep. candidate liberal/conservative placement
2000 Pre	000463	000464a	Gore liberal/conservative placement
2000 Pre	000473	000474a	Bush liberal/conservative placement
2000 Pre	000483	000484a	Buchanan liberal/conservative placement
2000 Pre	000696	000697	Gore abortion placement
2000 Pre	000698	000699	Bush abortion placement
2000 Pre	000735	000736	Gore gun control placement
2000 Pre	000739	000740	Bush gun control placement
2000 Pre	000783	000790	Gore environment placement
2000 Pre	000784	000791	Bush environment placement
2000 Post	001372	001373	Gore liberal/conservative placement
2000 Post	001374	001375	Bush liberal/conservative placement
2000 Post	001376	001377	Buchanan liberal/conservative placement
2000 Post	001378a	001379a	House Dem. candidate liberal/conservative placement
2000 Post	001379b	001379b	House indep. candidate #1 liberal/conservative placement
2000 Post	001380a	001381a	House Rep. candidate liberal/conservative placement
2000 Post	001380b	001381b	House indep. candidate #2 liberal/conservative placement
2000 Post	001405a	001406a	House Dem. candidate abortion placement
2000 Post	001405b	001406b	House indep. candidate #1 abortion placement
2000 Post	001407a	001408a	House Rep. candidate abortion placement
2000 Post	001407b	001408b	House indep. candidate #2 abortion placement

* denotes split ballot questions that were combined for data analysis.

Study 2 question text

Full text of all questions, Survey 1

Form 1

1. What job or political office does Mike Pence hold? [Vice President, Senate Majority Leader]
2. What job or political office does Mike Pompeo hold? [Secretary of Defense, Secretary of State]
3. What country does Angela Merkel lead? [Germany, Austria]
4. What country does Viktor Orban lead? [Turkey, Hungary]
5. In which chamber of Congress must revenue bills originate? [House of Representatives, Senate]
6. What Senate procedure allows budget changes with a simple majority vote? [Filibuster, Reconciliation]
7. How many votes in Congress are required to override a presidential veto? [A 2/3 majority of both the House and Senate, A simple majority of both the House and Senate]
8. Which party currently has the most members in the U.S. House of Representatives? [Democrats, Republicans]
9. Which party currently has the most members in the U.S. Senate? [Democrats, Republicans]
10. Each month, the Bureau of Labor Statistics estimates the *unemployment rate*. [break] Over the past year, did the unemployment rate increase or decrease? [Increased (more unemployment now), Decreased (less unemployment now)]
11. The amount of money people earn at their jobs is often measured using the *median real wage*. "Median" means the person right in the middle and "real" means adjusted for inflation. [break] Over the past year, did the median real wage increase or decrease? [Increased (higher wages now), Decreased (lower wages now)]
12. The rate of *inflation* measures how quickly prices are rising. [break] Over the past year, has inflation been higher or lower than the historical average (since 1945)? [Higher than average, Lower than average]
13. True or false? *The United States has never overthrown Iran's government.* [True, False]
14. Over the past year, did the percentage of Americans who have health insurance increase or decrease? [Increased (higher percentage has insurance now), Decreased (lower percentage has insurance now)]
15. True or false? *While she was Secretary of State, Hillary Clinton used a private email server to send and receive classified information.* [True, False]
16. True or false? *In 2017, a former Bernie Sanders campaign volunteer shot Republican congressman Steve Scalise and three other people during a practice for the House Republican baseball team.* [True, False]
17. True or false? *During the 2016 presidential campaign, Michael Cohen paid adult film actress Stormy Daniels to keep quiet about an alleged sexual encounter with President Trump.* [True, False]
18. True or false? *Before becoming president, Donald Trump was tape recorded saying that he kisses women and grabs them between the legs without their consent.* [True, False]

Form 2

1. When the U.S. buys more products from other countries than it sells to other countries, the U.S. has a *trade deficit.* [break] Between May 2018 and May 2019, did the U.S. trade deficit with other countries increase or decrease? [Increased, Decreased]
2. The value of the stock market is often measured using the Dow Jones Industrial Average (the Dow). [break] Over the past year, has the value of the stock market increased or decreased? [Increased, Decreased]
3. Most years, the U.S. national government spends more than it collects in taxes. In these years, the government has an *annual budget deficit.* [break] From 2017 to 2018, did the budget deficit increase or decrease? [Increased, Decreased]
4. Each year, how many immigrants are allowed to become legal permanent residents of the United States? [More than 1 million, Less than 1 million]
5. Over the past few decades, has the percentage of U.S. residents who are immigrants increased or decreased? [Increased, Decreased]
6. The pay gap between blacks and whites is often measured as the difference in their average hourly earnings. [break] Over the past forty years, has the pay gap between black and white men gotten larger or smaller? [Larger (less equal now), Smaller (more equal now)]
7. The pay gap between men and women is often measured using the difference in their median hourly earnings. [break] Over the past forty years, has the pay gap between men and women gotten larger or smaller? [Larger (less equal now), Smaller (more equal now)]
8. What job or political office does Kevin McCarthy hold? [House Minority Leader, Secretary of Defense]
9. What job or political office does John Roberts hold? [White House Chief of Staff, Chief Justice of the Supreme Court]
10. What job or political office does William Barr hold? [Attorney General, Secretary of State]
11. What country does Vladimir Putin lead? [Russia, Ukraine]
12. What country does Recep Erdogan lead? [Hungary, Turkey]
13. For how many years is a United States Senator elected – that is, how many years are there in one full term of office for a U.S. Senator? [2 years, 6 years]
14. On which of the following does the U.S. federal government currently spend the least? [Medicare, Foreign Aid]
15. Which political party is generally considered to be more conservative? [The Democratic Party, The Republican Party]
16. Which branch of government decides whether laws are constitutional? [Congress, Judiciary]
17. Under federal law, are lesbian, gay, and bisexual people currently protected from sexual orientation discrimination? [Yes, No]
18. Under federal law, are transgender people currently protected from gender identity discrimination? [Yes, No]

Full text of all questions, Survey 2

Form 1: CCES questions
Response options: Yes, No, Don't know

1. Do you support or oppose this proposal? *Always allow a woman to obtain an abortion as a matter of choice.*
2. Do you support or oppose this proposal? *Permit abortion ONLY in case of rape, incest or when the woman's life is in danger.*
3. Do you support or oppose this proposal? *Ban abortions after the 20th week of pregnancy.*
4. Do you support or oppose this proposal? *Allow employers to decline coverage of abortions in insurance plans.*
5. Do you support or oppose this proposal? *Prohibit the expenditure of funds authorized or appropriated by federal law for any abortion.*
6. Do you support or oppose this proposal? *Make abortions illegal in all circumstances.*
7. Do you support or oppose this proposal? *Increase spending on border security by $25 billion, including building a wall between the U.S. and Mexico.*
8. Do you support or oppose this proposal? *Provide legal status to children of immigrants who are already in the United States and were brought to the United States by their parents. Provide these children the option of citizenship in 10 years if they meet citizenship requirements and commit no crimes.*
9. Do you support or oppose this proposal? *Withhold federal funds from any local police department that does not report to the federal government anyone they identify as an illegal immigrant.*
10. Do you support or oppose this proposal? *Reduce legal immigration by 50 percent.*
11. Do you support or oppose this proposal? *Increase the number of visas for overseas workers to work in the US.*
12. Do you support or oppose this proposal? *Send to prison any person who has been deported from the United States and reenters the United States.*

Form 2: Additional policy attitudes (designed to vary in the percent saying DK)
Response options: Yes, No, Don't know

1. Do you support or oppose repealing the Affordable Care Act of 2010, also known as Obamacare?
2. Do you support or oppose repealing the "Cadillac tax" on high-cost employer-sponsored insurance plans?
3. Do you support or oppose allow Americans to purchase cheaper prescription drugs from Canada?
4. Do you support or oppose prohibiting drug companies from paying rebates to pharmacy benefit managers?
5. Do you support or oppose requiring banks that make risky investments to hold enough assets to cover their potential losses?
6. Do you support or oppose increasing set-aside requirements for banks engaged in derivatives trading?

7. Do you support or oppose allowing same-sex couples to marry?
8. Do you support or oppose discrimination protections under federal law for LGBTQ people?
9. Do you have a generally favorable or unfavorable opinion of the 2010 health reform law?
10. Do you have a generally favorable or unfavorable opinion of the 2015 education reform law?
11. Do you have a generally favorable or unfavorable opinion of the 2018 tax cut law?
12. Do you have a generally favorable or unfavorable opinion of the recent budget agreement in Congress?

Form 3: Favorability toward politicians
Response options: Yes, No, Don't know

1. Do you have a generally favorable or unfavorable opinion of Donald Trump?
2. Do you have a generally favorable or unfavorable opinion of Barack Obama?
3. Do you have a generally favorable or unfavorable opinion of Joe Biden?
4. Do you have a generally favorable or unfavorable opinion of Bernie Sanders?
5. Do you have a generally favorable or unfavorable opinion of Elizabeth Warren?
6. Do you have a generally favorable or unfavorable opinion of Kamala Harris?
7. Do you have a generally favorable or unfavorable opinion of Pete Buttigieg?
8. Do you have a generally favorable or unfavorable opinion of Kirsten Gillibrand?
9. Do you have a generally favorable or unfavorable opinion of Cory Booker?
10. Do you have a generally favorable or unfavorable opinion of Amy Klobuchar?
11. Do you have a generally favorable or unfavorable opinion of Steve Bullock?
12. Do you have a generally favorable or unfavorable opinion of Mike Gravel?

Form 4: Perceptions of policy
Response options: Yes, No, Don't know

1. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *All United States residents would have health insurance coverage.*
2. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *Taxes for most people would increase.*
3. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *Individuals and employers would continue to pay health insurance premiums.*
4. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *People would continue to pay deductibles and co-pays when they use health care services.*
5. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *People with insurance through their jobs would be able to keep their current plans.*
6. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *People who buy their own insurance would be able to keep their current plans.*

7. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *Private health insurance companies would still be the primary way Americans get health coverage.*
8. Do you think this would happen under a national health care plan, sometimes called Medicare For All? *Doctors and hospitals would be paid less.*
9. Which statement best describes the Green New Deal? [Response options: A set of goals, proposed in Congress as a resolution; A complete plan, proposed in Congress as a bill]
10. Does the proposed Green New Deal include this feature? *Guarantee every American a job.*
11. Does the proposed Green New Deal include this feature? *Guarantee every American access to clean water.*
12. Does the proposed Green New Deal include this feature? *Build a more sustainable food system.*
13. Does the proposed Green New Deal include this feature? *Government investments in renewable energy companies.*
14. Does the proposed Green New Deal include this feature? *Ban all oil drilling in the United States by 2030.*
15. Does the proposed Green New Deal include this feature? *Ban all gasoline powered cars by 2040.*
16. Does the proposed Green New Deal include this feature? *Create a new immigrant visa for victims of global climate change.*

Appendix B

Appendix to Chapter 4

These tables below present the estimates plotted in the main text figures. When the respondents' certainty levels were binned, the "certainty" column displays the midpoint of the bin.

Table B.1: Estimates plotted in Figure 4.1

Survey	Certainty	Estimate	SE	CI
Lucid, Sep. 2017	1.000	0.367	0.012	(0.343, 0.391)
	2.000	0.439	0.014	(0.411, 0.467)
	3.000	0.508	0.013	(0.483, 0.533)
	4.000	0.655	0.016	(0.624, 0.687)
	5.000	0.849	0.011	(0.827, 0.871)
Lucid, July 2019	0.500	0.544	0.009	(0.527, 0.561)
	0.600	0.576	0.012	(0.551, 0.600)
	0.700	0.580	0.012	(0.557, 0.604)
	0.800	0.658	0.011	(0.637, 0.680)
	0.900	0.747	0.011	(0.726, 0.767)
	1.000	0.923	0.004	(0.915, 0.930)
MTurk, Oct. 2019	0.500	0.537	0.041	(0.456, 0.617)
	0.550	0.441	0.039	(0.364, 0.518)
	0.650	0.537	0.044	(0.450, 0.625)
	0.750	0.653	0.025	(0.604, 0.702)
	0.850	0.682	0.033	(0.617, 0.748)
	0.945	0.795	0.036	(0.724, 0.866)
	1.000	0.958	0.008	(0.943, 0.973)
Lucid, Feb. 2020	0.500	0.517	0.042	(0.433, 0.601)
	0.550	0.551	0.028	(0.495, 0.607)
	0.650	0.613	0.033	(0.547, 0.679)
	0.750	0.613	0.033	(0.547, 0.679)
	0.850	0.618	0.036	(0.547, 0.690)
	0.945	0.700	0.029	(0.643, 0.757)
	1.000	0.905	0.014	(0.877, 0.933)
MTurk, Mar. 2020	0.500	0.474	0.029	(0.418, 0.530)
	0.550	0.528	0.032	(0.466, 0.590)
	0.650	0.531	0.037	(0.457, 0.604)
	0.750	0.564	0.032	(0.501, 0.626)
	0.850	0.573	0.030	(0.514, 0.631)
	0.945	0.755	0.027	(0.702, 0.808)
	1.000	0.952	0.007	(0.937, 0.966)

Table B.2: Estimates plotted in Figure 4.2

Survey	Outcome	Certainty	Estimate	SE	CI
MTurk, Mar. 2020	Response stability	0.500	0.532	0.046	(0.442, 0.623)
		0.550	0.573	0.044	(0.485, 0.661)
		0.650	0.712	0.058	(0.595, 0.829)
		0.750	0.605	0.058	(0.489, 0.720)
		0.850	0.719	0.040	(0.639, 0.799)
		0.945	0.890	0.029	(0.833, 0.947)
		1.000	0.983	0.006	(0.972, 0.994)
	Belief stability	0.500	0.495	0.016	(0.464, 0.527)
		0.550	0.509	0.020	(0.469, 0.550)
		0.650	0.575	0.036	(0.503, 0.648)
		0.750	0.546	0.038	(0.470, 0.621)
		0.850	0.667	0.031	(0.606, 0.728)
		0.945	0.829	0.026	(0.777, 0.880)
		1.000	0.973	0.006	(0.962, 0.984)
MTurk, Oct. 2019	Response stability	0.500	0.616	0.038	(0.541, 0.691)
		0.550	0.553	0.042	(0.469, 0.636)
		0.650	0.657	0.043	(0.570, 0.743)
		0.750	0.721	0.023	(0.675, 0.767)
		0.850	0.776	0.032	(0.713, 0.839)
		0.945	0.887	0.027	(0.833, 0.942)
		1.000	0.978	0.006	(0.966, 0.989)
	Belief stability	0.500	0.542	0.022	(0.498, 0.586)
		0.550	0.542	0.023	(0.496, 0.588)
		0.650	0.608	0.026	(0.556, 0.660)
		0.750	0.660	0.017	(0.626, 0.693)
		0.850	0.737	0.025	(0.687, 0.786)
		0.945	0.870	0.024	(0.822, 0.918)
		1.000	0.969	0.006	(0.957, 0.981)

Table B.3: Estimates plotted in Figure 4.3

Survey	Certainty	Estimate	SE	CI
MTurk, Oct. 2019	0.500	0.590	0.009	(0.572, 0.608)
	0.550	0.595	0.011	(0.573, 0.618)
	0.650	0.721	0.026	(0.669, 0.773)
	0.750	0.727	0.015	(0.698, 0.756)
	0.850	0.830	0.020	(0.790, 0.870)
	0.945	0.904	0.023	(0.858, 0.950)
	1.000	0.940	0.010	(0.920, 0.961)
Lucid, Feb. 2020	0.500	0.601	0.017	(0.566, 0.636)
	0.550	0.611	0.017	(0.577, 0.645)
	0.650	0.638	0.020	(0.599, 0.678)
	0.750	0.682	0.020	(0.643, 0.722)
	0.850	0.710	0.021	(0.668, 0.752)
	0.945	0.767	0.019	(0.729, 0.804)
	1.000	0.853	0.013	(0.827, 0.879)
MTurk, Mar. 2020	0.500	0.547	0.013	(0.521, 0.573)
	0.550	0.594	0.015	(0.564, 0.624)
	0.650	0.602	0.017	(0.568, 0.636)
	0.750	0.649	0.017	(0.615, 0.683)
	0.850	0.682	0.017	(0.648, 0.716)
	0.945	0.765	0.019	(0.727, 0.803)
	1.000	0.936	0.007	(0.923, 0.949)

Table B.4: Estimates plotted in Figure 4.4

Outcome	Scale	Certainty	Estimate	SE	CI
Accuracy	All in one	0.500	0.563	0.064	(0.434, 0.691)
		0.550	0.427	0.051	(0.326, 0.529)
		0.650	0.500	0.059	(0.382, 0.618)
		0.750	0.642	0.031	(0.581, 0.703)
		0.850	0.721	0.049	(0.621, 0.820)
		0.945	0.882	0.046	(0.789, 0.976)
		1.000	0.952	0.012	(0.929, 0.975)
	Branching	0.500	0.520	0.053	(0.415, 0.625)
		0.550	0.462	0.061	(0.340, 0.583)
		0.650	0.583	0.066	(0.450, 0.717)
		0.750	0.671	0.043	(0.587, 0.756)
		0.850	0.661	0.043	(0.575, 0.747)
		0.945	0.750	0.047	(0.656, 0.844)
		1.000	0.965	0.010	(0.946, 0.985)
Response stability	All in one	0.500	0.594	0.057	(0.479, 0.708)
		0.550	0.510	0.056	(0.399, 0.622)
		0.650	0.595	0.060	(0.475, 0.714)
		0.750	0.738	0.028	(0.683, 0.793)
		0.850	0.794	0.049	(0.696, 0.892)
		0.945	0.902	0.043	(0.814, 0.990)
		1.000	0.981	0.007	(0.967, 0.995)
	Branching	0.500	0.630	0.050	(0.529, 0.731)
		0.550	0.615	0.063	(0.488, 0.742)
		0.650	0.733	0.061	(0.609, 0.858)
		0.750	0.693	0.042	(0.610, 0.776)
		0.850	0.766	0.042	(0.683, 0.849)
		0.945	0.880	0.035	(0.810, 0.950)
		1.000	0.974	0.009	(0.955, 0.993)
Belief stability	All in one	0.500	0.532	0.041	(0.448, 0.616)
		0.550	0.528	0.028	(0.471, 0.584)
		0.650	0.578	0.035	(0.507, 0.650)
		0.750	0.670	0.020	(0.631, 0.710)
		0.850	0.756	0.034	(0.686, 0.826)
		0.945	0.889	0.035	(0.817, 0.962)
		1.000	0.972	0.008	(0.956, 0.988)
	Branching	0.500	0.549	0.026	(0.497, 0.601)
		0.550	0.561	0.038	(0.484, 0.638)
		0.650	0.640	0.039	(0.560, 0.719)
		0.750	0.641	0.031	(0.579, 0.704)
		0.850	0.728	0.033	(0.663, 0.794)
		0.945	0.861	0.031	(0.799, 0.924)
		1.000	0.967	0.009	(0.948, 0.985)
Revealed belief	All in one	0.500	0.598	0.017	(0.564, 0.633)
		0.550	0.608	0.016	(0.576, 0.640)
		0.650	0.714	0.034	(0.645, 0.783)
		0.750	0.712	0.017	(0.678, 0.745)
		0.850	0.858	0.034	(0.787, 0.930)
		0.945	0.922	0.046	(0.822, 1.022)
		1.000	0.935	0.015	(0.906, 0.965)
	Branching	0.500	0.585	0.010	(0.564, 0.605)
		0.550	0.579	0.015	(0.549, 0.609)
		0.650	0.729	0.041	(0.644, 0.814)
		0.750	0.762	0.028	(0.705, 0.818)
		0.850	0.816	0.025	(0.766, 0.865)
		0.945	0.897	0.026	(0.843, 0.951)
		1.000	0.945	0.015	(0.916, 0.975)

Appendix C

Appendix to Chapter 5

These tables below present the estimates plotted in the main text figures. When the respondents' certainty levels were binned, the "certainty" column displays the midpoint of the bin.

Table C.1: Estimates plotted in Figure 5.2

Question	Party	Answer	Certainty	Estimate	SE	CI	N
Clinton email	All respondents	Correct	0.500	0.808	0.043	(0.721, 0.895)	34
			0.550	0.767	0.053	(0.659, 0.876)	29
			0.650	0.772	0.046	(0.675, 0.869)	20
			0.750	0.828	0.036	(0.754, 0.902)	37
			0.850	0.872	0.019	(0.834, 0.911)	70
			0.945	0.938	0.012	(0.914, 0.963)	82
			1.000	0.960	0.007	(0.946, 0.975)	194
		Incorrect	0.500	0.558	0.171	(0.084, 1.032)	5
			0.550	0.280	0.180	(-2.007, 2.567)	2
			0.650	0.230	0.198	(-0.400, 0.860)	4
			0.750	0.211	0.104	(-0.035, 0.458)	8
			0.850	0.330	0.184	(-0.256, 0.916)	4
			0.945	0.495	0.405	(-4.651, 5.641)	2
			1.000	0.350	0.325	(-1.050, 1.750)	3
	Democrat	Correct	0.500	0.733	0.067	(0.592, 0.874)	18
			0.550	0.738	0.059	(0.616, 0.860)	25
			0.650	0.723	0.070	(0.569, 0.878)	12
			0.750	0.815	0.048	(0.716, 0.914)	24
			0.850	0.866	0.028	(0.810, 0.923)	39
			0.945	0.898	0.026	(0.845, 0.951)	34
			1.000	0.942	0.014	(0.914, 0.969)	79
		Incorrect	0.500	0.558	0.171	(0.084, 1.032)	5
			0.550	0.280	0.180	(-2.007, 2.567)	2
			0.650	0.033	0.033	(-0.110, 0.177)	3
			0.750	0.337	0.198	(-0.292, 0.967)	4
			0.850	0.200	0.200	(-2.341, 2.741)	2
			0.945	0.900	0.000	(NA, NA)	1
			1.000	0.350	0.325	(-1.050, 1.750)	3
	Republican	Correct	0.500	0.893	0.042	(0.803, 0.983)	16
			0.550	0.950	0.029	(0.859, 1.041)	4
			0.650	0.845	0.040	(0.751, 0.939)	8
			0.750	0.852	0.056	(0.731, 0.974)	13
			0.850	0.880	0.027	(0.825, 0.935)	31
			0.945	0.967	0.008	(0.952, 0.983)	48
			1.000	0.973	0.008	(0.957, 0.989)	115
		Incorrect	0.500			(NA, NA)	0
			0.550			(NA, NA)	0
			0.650	0.820	0.000	(NA, NA)	1
			0.750	0.085	0.031	(-0.014, 0.184)	4
			0.850	0.460	0.360	(-4.114, 5.034)	2
			0.945	0.090	0.000	(NA, NA)	1
			1.000			(NA, NA)	0
Obama birth certificate	All respondents	Correct	0.500	0.614	0.058	(0.496, 0.731)	37
			0.550	0.535	0.063	(0.407, 0.663)	34
			0.650	0.636	0.066	(0.501, 0.772)	33

Table C.1: Estimates plotted in Figure 5.2 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			0.750	0.752	0.066	(0.618, 0.887)	28
			0.850	0.758	0.061	(0.632, 0.883)	31
			0.945	0.879	0.035	(0.808, 0.950)	53
			1.000	0.886	0.025	(0.837, 0.935)	129
		Incorrect	0.500	0.592	0.063	(0.461, 0.723)	26
			0.550	0.578	0.074	(0.425, 0.730)	24
			0.650	0.312	0.142	(-0.043, 0.668)	8
			0.750	0.504	0.099	(0.294, 0.715)	19
			0.850	0.618	0.079	(0.454, 0.782)	21
			0.945	0.564	0.099	(0.352, 0.776)	17
			1.000	0.461	0.069	(0.320, 0.601)	34
	Democrat	Correct	0.500	0.713	0.092	(0.516, 0.911)	15
			0.550	0.729	0.083	(0.546, 0.912)	12
			0.650	0.746	0.092	(0.548, 0.944)	14
			0.750	0.730	0.081	(0.560, 0.901)	20
			0.850	0.808	0.069	(0.664, 0.952)	19
			0.945	0.932	0.031	(0.868, 0.996)	34
			1.000	0.907	0.029	(0.850, 0.964)	84
		Incorrect	0.500	0.419	0.124	(0.120, 0.717)	9
			0.550	0.409	0.136	(0.096, 0.722)	9
			0.650	0.310	0.424	(-5.081, 5.701)	3
			0.750	0.295	0.136	(-0.138, 0.728)	4
			0.850	0.566	0.139	(0.226, 0.905)	7
			0.945	0.133	0.120	(-0.248, 0.513)	4
			1.000	0.284	0.085	(0.105, 0.463)	17
	Republican	Correct	0.500	0.545	0.073	(0.394, 0.696)	22
			0.550	0.429	0.079	(0.266, 0.593)	22
			0.650	0.556	0.091	(0.365, 0.747)	19
			0.750	0.807	0.112	(0.542, 1.073)	8
			0.850	0.678	0.116	(0.423, 0.934)	12
			0.945	0.783	0.078	(0.620, 0.946)	19
			1.000	0.847	0.046	(0.755, 0.940)	45
		Incorrect	0.500	0.684	0.066	(0.543, 0.824)	17
			0.550	0.679	0.078	(0.511, 0.846)	15
			0.650	0.314	0.131	(-0.049, 0.677)	5
			0.750	0.560	0.121	(0.297, 0.823)	15
			0.850	0.644	0.098	(0.431, 0.856)	14
			0.945	0.697	0.097	(0.482, 0.912)	13
			1.000	0.637	0.093	(0.440, 0.834)	17
Obama DAPA reversal	All respondents	Correct	0.500	0.489	0.035	(0.417, 0.561)	23
			0.550	0.485	0.049	(0.385, 0.585)	31
			0.650	0.571	0.045	(0.480, 0.662)	32
			0.750	0.438	0.052	(0.332, 0.544)	28
			0.850	0.632	0.051	(0.528, 0.737)	26
			0.945	0.649	0.075	(0.487, 0.810)	15
			1.000	0.692	0.107	(0.460, 0.925)	13
		Incorrect	0.500	0.480	0.028	(0.424, 0.537)	54
			0.550	0.514	0.029	(0.456, 0.572)	53
			0.650	0.526	0.040	(0.444, 0.608)	33
			0.750	0.507	0.040	(0.427, 0.587)	48
			0.850	0.614	0.055	(0.503, 0.726)	32
			0.945	0.550	0.115	(0.293, 0.807)	11
			1.000	0.799	0.136	(0.466, 1.131)	7
	Democrat	Correct	0.500	0.427	0.070	(0.263, 0.592)	8
			0.550	0.443	0.073	(0.289, 0.597)	17
			0.650	0.473	0.075	(0.309, 0.637)	13
			0.750	0.402	0.069	(0.255, 0.550)	17
			0.850	0.596	0.112	(0.338, 0.853)	9
			0.945	0.585	0.133	(0.243, 0.927)	6
			1.000	0.402	0.218	(-0.203, 1.007)	5
		Incorrect	0.500	0.563	0.033	(0.496, 0.631)	27

Table C.1: Estimates plotted in Figure 5.2 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			0.550	0.563	0.035	(0.490, 0.635)	29
			0.650	0.602	0.041	(0.517, 0.688)	20
			0.750	0.545	0.051	(0.440, 0.649)	31
			0.850	0.659	0.056	(0.543, 0.775)	24
			0.945	0.674	0.126	(0.377, 0.971)	8
			1.000	0.799	0.136	(0.466, 1.131)	7
	Republican	Correct	0.500	0.457	0.057	(0.311, 0.602)	6
			0.550	0.596	0.064	(0.451, 0.741)	10
			0.650	0.686	0.062	(0.552, 0.821)	14
			0.750	0.544	0.101	(0.306, 0.782)	8
			0.850	0.652	0.053	(0.540, 0.763)	17
			0.945	0.688	0.105	(0.440, 0.935)	8
			1.000	0.883	0.057	(0.744, 1.022)	7
		Incorrect	0.500	0.367	0.048	(0.263, 0.471)	14
			0.550	0.440	0.057	(0.319, 0.561)	17
			0.650	0.424	0.074	(0.261, 0.587)	12
			0.750	0.468	0.082	(0.285, 0.652)	11
			0.850	0.436	0.149	(0.071, 0.800)	7
			0.945	0.080	0.080	(-0.936, 1.096)	2
			1.000			(NA, NA)	0
Trump Article II claim	All respondents	Correct	0.500	0.587	0.058	(0.463, 0.710)	18
			0.550	0.595	0.044	(0.504, 0.686)	30
			0.650	0.678	0.041	(0.593, 0.762)	30
			0.750	0.631	0.039	(0.551, 0.710)	43
			0.850	0.793	0.033	(0.727, 0.860)	45
			0.945	0.780	0.047	(0.684, 0.877)	33
			1.000	0.912	0.034	(0.843, 0.981)	48
		Incorrect	0.500	0.518	0.053	(0.406, 0.629)	19
			0.550	0.410	0.058	(0.290, 0.530)	23
			0.650	0.428	0.057	(0.307, 0.549)	19
			0.750	0.486	0.074	(0.332, 0.640)	20
			0.850	0.543	0.064	(0.412, 0.675)	26
			0.945	0.574	0.099	(0.365, 0.784)	18
			1.000	0.759	0.061	(0.634, 0.883)	29
	Democrat	Correct	0.500	0.686	0.093	(0.458, 0.913)	7
			0.550	0.647	0.060	(0.519, 0.775)	17
			0.650	0.699	0.052	(0.588, 0.809)	15
			0.750	0.671	0.046	(0.578, 0.765)	27
			0.850	0.809	0.039	(0.729, 0.888)	33
			0.945	0.766	0.059	(0.644, 0.889)	22
			1.000	0.936	0.032	(0.871, 1.001)	42
		Incorrect	0.500	0.370	0.070	(0.190, 0.550)	6
			0.550	0.335	0.091	(0.133, 0.538)	11
			0.650	0.354	0.112	(0.080, 0.628)	7
			0.750	0.395	0.117	(0.130, 0.660)	10
			0.850	0.461	0.122	(0.180, 0.742)	9
			0.945	0.361	0.146	(0.016, 0.706)	8
			1.000	0.390	0.150	(-0.254, 1.034)	3
	Republican	Correct	0.500	0.552	0.081	(0.343, 0.761)	6
			0.550	0.547	0.076	(0.371, 0.723)	9
			0.650	0.693	0.076	(0.524, 0.861)	12
			0.750	0.575	0.098	(0.353, 0.797)	10
			0.850	0.729	0.076	(0.554, 0.904)	9
			0.945	0.794	0.102	(0.560, 1.029)	9
			1.000	0.893	0.074	(0.573, 1.214)	3
		Incorrect	0.500	0.713	0.079	(0.509, 0.917)	6
			0.550	0.478	0.086	(0.284, 0.672)	10
			0.650	0.441	0.069	(0.284, 0.598)	10
			0.750	0.574	0.096	(0.354, 0.795)	9
			0.850	0.600	0.090	(0.395, 0.805)	10
			0.945	0.825	0.107	(0.572, 1.078)	8
			1.000	0.789	0.065	(0.654, 0.924)	24

Table C.1: Estimates plotted in Figure 5.2 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
Trump Russia collusion	All respondents	Correct	0.500	0.512	0.037	(0.438, 0.586)	56
			0.550	0.562	0.045	(0.471, 0.653)	47
			0.650	0.481	0.051	(0.377, 0.585)	36
			0.750	0.616	0.034	(0.548, 0.685)	79
			0.850	0.676	0.038	(0.600, 0.752)	71
			0.945	0.706	0.036	(0.634, 0.777)	101
			1.000	0.863	0.018	(0.828, 0.898)	249
		Incorrect	0.500	0.644	0.045	(0.553, 0.735)	35
			0.550	0.647	0.048	(0.550, 0.743)	40
			0.650	0.614	0.051	(0.510, 0.718)	35
			0.750	0.610	0.049	(0.511, 0.708)	38
			0.850	0.669	0.046	(0.577, 0.761)	49
			0.945	0.689	0.056	(0.575, 0.803)	36
			1.000	0.380	0.084	(0.205, 0.555)	24
	Democrat	Correct	0.500	0.484	0.059	(0.361, 0.608)	20
			0.550	0.495	0.077	(0.333, 0.657)	18
			0.650	0.505	0.071	(0.358, 0.651)	22
			0.750	0.596	0.047	(0.501, 0.692)	43
			0.850	0.670	0.047	(0.575, 0.764)	45
			0.945	0.670	0.055	(0.560, 0.781)	47
			1.000	0.810	0.030	(0.750, 0.869)	106
		Incorrect	0.500	0.738	0.049	(0.637, 0.839)	21
			0.550	0.637	0.060	(0.514, 0.760)	30
			0.650	0.653	0.057	(0.536, 0.770)	26
			0.750	0.642	0.064	(0.509, 0.775)	22
			0.850	0.675	0.058	(0.556, 0.794)	34
			0.945	0.762	0.071	(0.612, 0.912)	20
			1.000	0.410	0.118	(0.152, 0.668)	14
	Republican	Correct	0.500	0.517	0.066	(0.380, 0.653)	24
			0.550	0.607	0.064	(0.475, 0.739)	24
			0.650	0.464	0.075	(0.300, 0.628)	13
			0.750	0.642	0.052	(0.536, 0.748)	33
			0.850	0.695	0.071	(0.547, 0.842)	22
			0.945	0.739	0.050	(0.637, 0.840)	49
			1.000	0.901	0.022	(0.858, 0.945)	135
		Incorrect	0.500	0.545	0.078	(0.365, 0.724)	11
			0.550	0.670	0.058	(0.533, 0.807)	8
			0.650	0.501	0.110	(0.249, 0.754)	9
			0.750	0.604	0.079	(0.432, 0.775)	14
			0.850	0.608	0.085	(0.417, 0.798)	12
			0.945	0.598	0.090	(0.400, 0.795)	16
			1.000	0.410	0.147	(0.063, 0.757)	8
Trump said 'grab them'	All respondents	Correct	0.500	0.794	0.036	(0.721, 0.868)	34
			0.550	0.764	0.044	(0.674, 0.854)	35
			0.650	0.729	0.059	(0.609, 0.850)	27
			0.750	0.831	0.033	(0.764, 0.898)	50
			0.850	0.823	0.032	(0.760, 0.886)	79
			0.945	0.884	0.021	(0.842, 0.926)	123
			1.000	0.947	0.009	(0.929, 0.965)	367
		Incorrect	0.500	0.413	0.059	(0.291, 0.535)	29
			0.550	0.522	0.058	(0.402, 0.642)	26
			0.650	0.418	0.086	(0.234, 0.601)	16
			0.750	0.442	0.076	(0.286, 0.599)	25
			0.850	0.544	0.076	(0.387, 0.701)	28
			0.945	0.604	0.071	(0.457, 0.750)	28
			1.000	0.623	0.072	(0.478, 0.769)	35
	Democrat	Correct	0.500	0.823	0.077	(0.649, 0.997)	10
			0.550	0.896	0.031	(0.831, 0.962)	17
			0.650	0.806	0.068	(0.661, 0.951)	15
			0.750	0.911	0.028	(0.855, 0.968)	32
			0.850	0.825	0.045	(0.734, 0.916)	40

Table C.1: Estimates plotted in Figure 5.2 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			0.945	0.964	0.012	(0.940, 0.989)	60
			1.000	0.984	0.005	(0.974, 0.994)	246
		Incorrect	0.500	0.411	0.127	(0.117, 0.705)	9
			0.550	0.453	0.106	(0.216, 0.689)	11
			0.650	0.292	0.179	(-0.277, 0.862)	4
			0.750	0.303	0.116	(0.036, 0.571)	9
			0.850	0.251	0.141	(-0.094, 0.597)	7
			0.945	0.676	0.191	(0.144, 1.208)	5
			1.000	0.262	0.131	(-0.041, 0.565)	9
	Republican	Correct	0.500	0.803	0.036	(0.728, 0.879)	23
			0.550	0.643	0.078	(0.477, 0.809)	16
			0.650	0.642	0.106	(0.406, 0.878)	11
			0.750	0.713	0.083	(0.533, 0.893)	13
			0.850	0.806	0.051	(0.701, 0.910)	34
			0.945	0.801	0.039	(0.723, 0.880)	60
			1.000	0.853	0.030	(0.794, 0.912)	100
		Incorrect	0.500	0.434	0.078	(0.267, 0.601)	17
			0.550	0.582	0.077	(0.413, 0.751)	12
			0.650	0.459	0.099	(0.240, 0.678)	12
			0.750	0.521	0.096	(0.316, 0.725)	16
			0.850	0.641	0.084	(0.463, 0.818)	19
			0.945	0.587	0.088	(0.402, 0.773)	19
			1.000	0.753	0.077	(0.592, 0.913)	24

Table C.2: Estimates plotted in Figures 5.3 and 5.4

Question	Party	Answer	Certainty	Estimate	SE	CI	N
Budget deficit increased	All respondents	Correct	0.500	0.571	0.029	(0.511, 0.630)	38
			0.550	0.589	0.028	(0.532, 0.645)	60
			0.650	0.620	0.036	(0.548, 0.691)	54
			0.750	0.664	0.022	(0.620, 0.708)	117
			0.850	0.718	0.023	(0.672, 0.764)	125
			0.945	0.735	0.022	(0.692, 0.779)	135
			1.000	0.841	0.019	(0.803, 0.879)	144
		Incorrect	0.500	0.508	0.039	(0.429, 0.587)	36
			0.550	0.552	0.033	(0.486, 0.619)	39
			0.650	0.616	0.052	(0.510, 0.721)	29
			0.750	0.452	0.056	(0.339, 0.565)	36
			0.850	0.635	0.064	(0.502, 0.768)	23
			0.945	0.624	0.074	(0.466, 0.782)	15
			1.000	0.807	0.097	(0.391, 1.223)	3
	Democrat	Correct	0.500	0.586	0.028	(0.527, 0.644)	29
			0.550	0.547	0.033	(0.480, 0.614)	36
			0.650	0.586	0.051	(0.481, 0.690)	30
			0.750	0.667	0.031	(0.605, 0.730)	65
			0.850	0.753	0.028	(0.696, 0.810)	62
			0.945	0.779	0.028	(0.724, 0.835)	69
			1.000	0.826	0.026	(0.774, 0.877)	92
		Incorrect	0.500	0.478	0.079	(0.305, 0.652)	13
			0.550	0.623	0.036	(0.548, 0.697)	19
			0.650	0.628	0.070	(0.472, 0.784)	11
			0.750	0.436	0.093	(0.240, 0.633)	17
			0.850	0.570	0.083	(0.384, 0.756)	11
			0.945	0.531	0.121	(0.234, 0.829)	7
			1.000	0.855	0.145	(-0.987, 2.697)	2

Table C.2: Estimates plotted in Figures 5.3 and 5.4 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
	Republican	Correct	0.500	0.502	0.180	(-0.071, 1.076)	4
			0.550	0.701	0.046	(0.605, 0.798)	17
			0.650	0.645	0.058	(0.524, 0.767)	19
			0.750	0.615	0.035	(0.544, 0.686)	40
			0.850	0.668	0.041	(0.586, 0.750)	53
			0.945	0.672	0.037	(0.597, 0.747)	56
			1.000	0.882	0.029	(0.822, 0.941)	36
		Incorrect	0.500	0.595	0.053	(0.476, 0.714)	10
			0.550	0.503	0.065	(0.363, 0.642)	15
			0.650	0.596	0.080	(0.423, 0.769)	15
			0.750	0.504	0.073	(0.347, 0.661)	15
			0.850	0.731	0.098	(0.513, 0.949)	11
			0.945	0.705	0.085	(0.504, 0.906)	8
			1.000	0.710	0.000	(NA, NA)	1
GDP growth below 4%	All respondents	Correct	0.500	0.532	0.026	(0.479, 0.584)	57
			0.550	0.565	0.023	(0.518, 0.611)	67
			0.650	0.586	0.032	(0.522, 0.651)	55
			0.750	0.629	0.028	(0.574, 0.684)	73
			0.850	0.677	0.038	(0.600, 0.753)	48
			0.945	0.718	0.033	(0.653, 0.784)	54
			1.000	0.888	0.023	(0.843, 0.933)	59
		Incorrect	0.500	0.500	0.037	(0.426, 0.575)	41
			0.550	0.609	0.024	(0.561, 0.657)	83
			0.650	0.637	0.028	(0.582, 0.693)	66
			0.750	0.644	0.024	(0.596, 0.692)	102
			0.850	0.641	0.028	(0.586, 0.696)	85
			0.945	0.685	0.043	(0.597, 0.772)	48
			1.000	0.774	0.051	(0.668, 0.880)	22
	Democrat	Correct	0.500	0.477	0.028	(0.420, 0.534)	34
			0.550	0.573	0.027	(0.519, 0.628)	53
			0.650	0.612	0.040	(0.531, 0.693)	36
			0.750	0.649	0.033	(0.583, 0.715)	46
			0.850	0.719	0.048	(0.618, 0.820)	22
			0.945	0.725	0.044	(0.635, 0.814)	31
			1.000	0.912	0.022	(0.867, 0.956)	39
		Incorrect	0.500	0.472	0.051	(0.364, 0.580)	20
			0.550	0.580	0.034	(0.512, 0.648)	47
			0.650	0.632	0.045	(0.540, 0.724)	31
			0.750	0.665	0.026	(0.612, 0.718)	50
			0.850	0.640	0.042	(0.556, 0.724)	38
			0.945	0.643	0.073	(0.483, 0.803)	13
			1.000	0.842	0.066	(0.686, 0.999)	8
	Republican	Correct	0.500	0.690	0.089	(0.480, 0.900)	8
			0.550	0.512	0.046	(0.409, 0.615)	10
			0.650	0.547	0.062	(0.415, 0.680)	15
			0.750	0.579	0.067	(0.437, 0.722)	18
			0.850	0.641	0.067	(0.502, 0.780)	21
			0.945	0.710	0.049	(0.608, 0.812)	23
			1.000	0.873	0.046	(0.775, 0.971)	17
		Incorrect	0.500	0.583	0.077	(0.408, 0.758)	10
			0.550	0.654	0.043	(0.565, 0.743)	24
			0.650	0.622	0.037	(0.545, 0.699)	28
			0.750	0.583	0.052	(0.476, 0.689)	36
			0.850	0.652	0.038	(0.575, 0.730)	45
			0.945	0.682	0.055	(0.570, 0.793)	33
			1.000	0.715	0.073	(0.556, 0.873)	13
Inflation below average	All respondents	Correct	0.500	0.578	0.025	(0.528, 0.628)	71
			0.550	0.606	0.020	(0.566, 0.646)	83
			0.650	0.617	0.036	(0.546, 0.689)	59
			0.750	0.648	0.027	(0.594, 0.702)	89
			0.850	0.700	0.031	(0.639, 0.762)	69
			0.945	0.726	0.050	(0.624, 0.828)	32

Table C.2: Estimates plotted in Figures 5.3 and 5.4 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			1.000	0.858	0.029	(0.801, 0.916)	51
		Incorrect	0.500	0.555	0.022	(0.510, 0.600)	63
			0.550	0.657	0.019	(0.619, 0.695)	76
			0.650	0.648	0.028	(0.592, 0.704)	53
			0.750	0.608	0.028	(0.552, 0.665)	86
			0.850	0.585	0.028	(0.529, 0.641)	81
			0.945	0.597	0.040	(0.517, 0.677)	43
			1.000	0.715	0.087	(0.528, 0.903)	15
	Democrat	Correct	0.500	0.560	0.032	(0.495, 0.625)	33
			0.550	0.615	0.023	(0.568, 0.662)	44
			0.650	0.608	0.061	(0.482, 0.734)	25
			0.750	0.690	0.038	(0.612, 0.767)	38
			0.850	0.661	0.047	(0.565, 0.757)	29
			0.945	0.813	0.051	(0.703, 0.922)	14
			1.000	0.896	0.032	(0.830, 0.962)	27
		Incorrect	0.500	0.559	0.025	(0.508, 0.610)	44
			0.550	0.648	0.021	(0.606, 0.689)	53
			0.650	0.648	0.030	(0.587, 0.709)	40
			0.750	0.614	0.042	(0.529, 0.699)	48
			0.850	0.619	0.038	(0.543, 0.695)	43
			0.945	0.588	0.073	(0.432, 0.744)	15
			1.000	0.728	0.101	(0.506, 0.949)	12
	Republican	Correct	0.500	0.578	0.049	(0.477, 0.679)	23
			0.550	0.591	0.045	(0.497, 0.684)	27
			0.650	0.654	0.046	(0.559, 0.749)	29
			0.750	0.627	0.041	(0.544, 0.710)	42
			0.850	0.727	0.046	(0.634, 0.821)	34
			0.945	0.621	0.083	(0.445, 0.798)	16
			1.000	0.839	0.047	(0.740, 0.938)	20
		Incorrect	0.500	0.513	0.066	(0.367, 0.659)	11
			0.550	0.657	0.062	(0.522, 0.792)	14
			0.650	0.650	0.088	(0.451, 0.849)	10
			0.750	0.581	0.045	(0.489, 0.673)	28
			0.850	0.525	0.046	(0.431, 0.619)	33
			0.945	0.583	0.053	(0.473, 0.694)	24
			1.000	0.710	0.000	(NA, NA)	1
Obama DAPA reversal	All respondents	Correct	0.500	0.560	0.028	(0.504, 0.617)	43
			0.550	0.575	0.024	(0.526, 0.623)	71
			0.650	0.575	0.038	(0.498, 0.651)	51
			0.750	0.592	0.032	(0.528, 0.656)	72
			0.850	0.585	0.034	(0.518, 0.652)	67
			0.945	0.688	0.039	(0.610, 0.766)	47
			1.000	0.770	0.060	(0.645, 0.894)	21
		Incorrect	0.500	0.573	0.018	(0.536, 0.609)	101
			0.550	0.604	0.014	(0.576, 0.632)	122
			0.650	0.602	0.025	(0.551, 0.653)	69
			0.750	0.639	0.027	(0.586, 0.692)	94
			0.850	0.603	0.030	(0.542, 0.663)	71
			0.945	0.704	0.043	(0.617, 0.792)	33
			1.000	0.734	0.062	(0.604, 0.864)	19
	Democrat	Correct	0.500	0.572	0.028	(0.513, 0.631)	24
			0.550	0.572	0.027	(0.518, 0.626)	40
			0.650	0.547	0.053	(0.436, 0.657)	24
			0.750	0.603	0.045	(0.511, 0.695)	42
			0.850	0.590	0.042	(0.504, 0.676)	26
			0.945	0.687	0.079	(0.511, 0.863)	11
			1.000	0.790	0.072	(0.623, 0.957)	9
		Incorrect	0.500	0.569	0.024	(0.522, 0.617)	54
			0.550	0.588	0.019	(0.551, 0.625)	69
			0.650	0.614	0.031	(0.551, 0.677)	42
			0.750	0.665	0.029	(0.607, 0.724)	57

Table C.2: Estimates plotted in Figures 5.3 and 5.4 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			0.850	0.647	0.039	(0.569, 0.725)	45
			0.945	0.658	0.057	(0.538, 0.777)	18
			1.000	0.818	0.050	(0.710, 0.925)	16
	Republican	Correct	0.500	0.582	0.067	(0.428, 0.736)	9
			0.550	0.601	0.048	(0.501, 0.701)	24
			0.650	0.553	0.055	(0.438, 0.669)	20
			0.750	0.589	0.059	(0.467, 0.711)	22
			0.850	0.582	0.048	(0.484, 0.680)	41
			0.945	0.678	0.047	(0.582, 0.774)	34
			1.000	0.749	0.101	(0.525, 0.973)	11
		Incorrect	0.500	0.555	0.045	(0.461, 0.649)	23
			0.550	0.621	0.024	(0.571, 0.670)	38
			0.650	0.577	0.055	(0.461, 0.692)	21
			0.750	0.587	0.060	(0.464, 0.710)	30
			0.850	0.502	0.055	(0.388, 0.617)	21
			0.945	0.762	0.070	(0.610, 0.914)	13
			1.000	0.290	0.000	(0.290, 0.290)	2
Political awareness	All respondents	Correct	0.500	0.591	0.016	(0.558, 0.623)	181
			0.550	0.631	0.016	(0.599, 0.663)	195
			0.650	0.639	0.018	(0.603, 0.675)	145
			0.750	0.665	0.019	(0.627, 0.703)	211
			0.850	0.740	0.018	(0.705, 0.775)	246
			0.945	0.804	0.017	(0.771, 0.837)	274
			1.000	0.942	0.006	(0.930, 0.954)	979
		Incorrect	0.500	0.519	0.015	(0.490, 0.549)	197
			0.550	0.565	0.016	(0.533, 0.597)	208
			0.650	0.592	0.020	(0.553, 0.632)	159
			0.750	0.637	0.017	(0.603, 0.671)	214
			0.850	0.629	0.019	(0.591, 0.667)	203
			0.945	0.656	0.028	(0.599, 0.712)	124
			1.000	0.752	0.030	(0.693, 0.812)	81
	Democrat	Correct	0.500	0.585	0.021	(0.544, 0.626)	102
			0.550	0.638	0.024	(0.591, 0.686)	105
			0.650	0.638	0.023	(0.593, 0.683)	99
			0.750	0.685	0.024	(0.637, 0.733)	121
			0.850	0.744	0.025	(0.694, 0.794)	128
			0.945	0.855	0.020	(0.814, 0.895)	144
			1.000	0.942	0.008	(0.926, 0.957)	596
		Incorrect	0.500	0.530	0.019	(0.492, 0.568)	96
			0.550	0.571	0.024	(0.523, 0.618)	96
			0.650	0.615	0.024	(0.567, 0.663)	90
			0.750	0.655	0.024	(0.607, 0.703)	98
			0.850	0.665	0.030	(0.605, 0.726)	86
			0.945	0.678	0.040	(0.596, 0.759)	49
			1.000	0.759	0.033	(0.691, 0.827)	40
	Republican	Correct	0.500	0.592	0.032	(0.527, 0.657)	54
			0.550	0.642	0.025	(0.592, 0.692)	67
			0.650	0.636	0.037	(0.561, 0.711)	34
			0.750	0.633	0.039	(0.555, 0.711)	70
			0.850	0.706	0.029	(0.649, 0.764)	93
			0.945	0.726	0.028	(0.670, 0.783)	114
			1.000	0.936	0.012	(0.912, 0.961)	291
		Incorrect	0.500	0.541	0.027	(0.487, 0.595)	48
			0.550	0.561	0.024	(0.513, 0.609)	88
			0.650	0.568	0.038	(0.491, 0.645)	54
			0.750	0.620	0.029	(0.563, 0.677)	88
			0.850	0.588	0.026	(0.536, 0.640)	103
			0.945	0.641	0.042	(0.557, 0.724)	69
			1.000	0.759	0.056	(0.642, 0.875)	32
Trump-Russia collusion	All respondents	Correct	0.500	0.534	0.041	(0.451, 0.617)	38
			0.550	0.619	0.033	(0.552, 0.687)	37

Table C.2: Estimates plotted in Figures 5.3 and 5.4 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			0.650	0.654	0.037	(0.578, 0.729)	44
			0.750	0.732	0.026	(0.681, 0.784)	85
			0.850	0.730	0.030	(0.669, 0.790)	94
			0.945	0.813	0.025	(0.764, 0.863)	83
			1.000	0.900	0.014	(0.873, 0.928)	181
		Incorrect	0.500	0.561	0.060	(0.433, 0.688)	16
			0.550	0.609	0.033	(0.542, 0.676)	40
			0.650	0.712	0.031	(0.649, 0.775)	47
			0.750	0.620	0.033	(0.555, 0.685)	65
			0.850	0.625	0.034	(0.556, 0.694)	65
			0.945	0.727	0.036	(0.656, 0.799)	55
			1.000	0.729	0.046	(0.633, 0.825)	25
	Democrat	Correct	0.500	0.567	0.054	(0.448, 0.686)	13
			0.550	0.649	0.040	(0.567, 0.731)	23
			0.650	0.680	0.038	(0.601, 0.758)	25
			0.750	0.718	0.034	(0.650, 0.787)	53
			0.850	0.788	0.033	(0.720, 0.855)	49
			0.945	0.837	0.030	(0.777, 0.897)	41
			1.000	0.882	0.025	(0.832, 0.931)	84
		Incorrect	0.500	0.557	0.095	(0.313, 0.800)	6
			0.550	0.620	0.039	(0.540, 0.700)	26
			0.650	0.697	0.039	(0.617, 0.777)	34
			0.750	0.635	0.045	(0.545, 0.726)	36
			0.850	0.593	0.045	(0.502, 0.685)	43
			0.945	0.755	0.045	(0.661, 0.849)	24
			1.000	0.744	0.051	(0.635, 0.853)	16
	Republican	Correct	0.500	0.538	0.108	(0.288, 0.788)	9
			0.550	0.506	0.091	(0.282, 0.729)	7
			0.650	0.643	0.088	(0.455, 0.832)	15
			0.750	0.790	0.035	(0.718, 0.862)	27
			0.850	0.640	0.058	(0.523, 0.757)	37
			0.945	0.764	0.048	(0.667, 0.861)	33
			1.000	0.911	0.017	(0.877, 0.944)	80
		Incorrect	0.500	0.450	0.096	(0.203, 0.697)	6
			0.550	0.584	0.081	(0.401, 0.767)	10
			0.650	0.745	0.063	(0.602, 0.888)	10
			0.750	0.597	0.055	(0.483, 0.711)	25
			0.850	0.664	0.059	(0.539, 0.788)	17
			0.945	0.675	0.056	(0.560, 0.790)	28
			1.000	0.775	0.168	(0.241, 1.309)	4
Trump Article II claim	All respondents	Correct	0.500	0.572	0.034	(0.503, 0.640)	32
			0.550	0.644	0.023	(0.597, 0.691)	51
			0.650	0.666	0.024	(0.618, 0.713)	76
			0.750	0.691	0.022	(0.647, 0.735)	94
			0.850	0.684	0.026	(0.632, 0.736)	100
			0.945	0.701	0.027	(0.646, 0.756)	84
			1.000	0.859	0.022	(0.814, 0.903)	105
		Incorrect	0.500	0.588	0.034	(0.518, 0.657)	38
			0.550	0.554	0.035	(0.484, 0.623)	60
			0.650	0.603	0.037	(0.528, 0.677)	45
			0.750	0.596	0.040	(0.515, 0.677)	61
			0.850	0.676	0.035	(0.605, 0.747)	51
			0.945	0.754	0.041	(0.671, 0.836)	34
			1.000	0.785	0.045	(0.695, 0.875)	43
	Democrat	Correct	0.500	0.549	0.043	(0.458, 0.640)	16
			0.550	0.643	0.031	(0.580, 0.706)	31
			0.650	0.657	0.027	(0.602, 0.712)	43
			0.750	0.700	0.023	(0.653, 0.747)	58
			0.850	0.700	0.031	(0.639, 0.762)	65
			0.945	0.748	0.034	(0.679, 0.816)	50
			1.000	0.878	0.022	(0.835, 0.922)	86

Table C.2: Estimates plotted in Figures 5.3 and 5.4 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
		Incorrect	0.500	0.521	0.046	(0.423, 0.620)	15
			0.550	0.496	0.061	(0.372, 0.620)	29
			0.650	0.598	0.051	(0.491, 0.705)	23
			0.750	0.562	0.087	(0.381, 0.744)	20
			0.850	0.710	0.056	(0.592, 0.828)	18
			0.945	0.781	0.080	(0.605, 0.956)	12
			1.000	0.506	0.158	(0.067, 0.945)	5
	Republican	Correct	0.500	0.540	0.065	(0.385, 0.695)	8
			0.550	0.660	0.038	(0.578, 0.742)	15
			0.650	0.663	0.049	(0.563, 0.763)	27
			0.750	0.638	0.057	(0.520, 0.755)	26
			0.850	0.641	0.057	(0.524, 0.759)	28
			0.945	0.624	0.046	(0.530, 0.719)	32
			1.000	0.757	0.090	(0.561, 0.953)	13
		Incorrect	0.500	0.674	0.078	(0.497, 0.851)	10
			0.550	0.594	0.043	(0.505, 0.683)	21
			0.650	0.623	0.052	(0.513, 0.733)	19
			0.750	0.600	0.046	(0.506, 0.694)	37
			0.850	0.664	0.059	(0.543, 0.785)	25
			0.945	0.752	0.046	(0.655, 0.850)	18
			1.000	0.810	0.048	(0.713, 0.906)	34
Trump said 'grab them'	All respondents	Correct	0.500	0.692	0.086	(0.470, 0.913)	6
			0.550	0.699	0.056	(0.583, 0.816)	20
			0.650	0.615	0.059	(0.492, 0.737)	24
			0.750	0.701	0.041	(0.619, 0.783)	54
			0.850	0.688	0.028	(0.632, 0.745)	90
			0.945	0.791	0.025	(0.743, 0.840)	114
			1.000	0.928	0.008	(0.912, 0.944)	337
		Incorrect	0.500	0.621	0.041	(0.534, 0.707)	19
			0.550	0.652	0.043	(0.562, 0.742)	22
			0.650	0.540	0.054	(0.430, 0.651)	26
			0.750	0.620	0.043	(0.534, 0.706)	41
			0.850	0.686	0.038	(0.609, 0.764)	47
			0.945	0.764	0.040	(0.683, 0.845)	39
			1.000	0.759	0.042	(0.674, 0.845)	51
	Democrat	Correct	0.500	1.000	0.000	(NA, NA)	1
			0.550	0.696	0.075	(0.525, 0.867)	10
			0.650	0.614	0.087	(0.425, 0.803)	13
			0.750	0.711	0.053	(0.602, 0.820)	29
			0.850	0.697	0.036	(0.624, 0.769)	47
			0.945	0.790	0.034	(0.723, 0.858)	62
			1.000	0.927	0.010	(0.908, 0.945)	240
		Incorrect	0.500	0.605	0.044	(0.501, 0.709)	8
			0.550	0.656	0.079	(0.470, 0.842)	8
			0.650	0.608	0.082	(0.425, 0.791)	11
			0.750	0.726	0.061	(0.595, 0.857)	14
			0.850	0.675	0.063	(0.540, 0.810)	16
			0.945	0.820	0.069	(0.667, 0.973)	12
			1.000	0.739	0.098	(0.527, 0.950)	15
	Republican	Correct	0.500	0.660	0.150	(-1.246, 2.566)	2
			0.550	0.803	0.078	(0.602, 1.005)	6
			0.650	0.592	0.112	(0.327, 0.858)	8
			0.750	0.648	0.095	(0.447, 0.850)	16
			0.850	0.633	0.051	(0.530, 0.737)	35
			0.945	0.776	0.040	(0.697, 0.856)	47
			1.000	0.917	0.022	(0.873, 0.960)	64
		Incorrect	0.500	0.588	0.099	(0.335, 0.842)	6
			0.550	0.652	0.064	(0.505, 0.799)	9
			0.650	0.505	0.080	(0.330, 0.679)	13
			0.750	0.548	0.063	(0.416, 0.680)	22
			0.850	0.694	0.054	(0.583, 0.804)	28
			0.945	0.749	0.055	(0.635, 0.863)	23

Table C.2: Estimates plotted in Figures 5.3 and 5.4 (continued)

Question	Party	Answer	Certainty	Estimate	SE	CI	N
			1.000	0.764	0.050	(0.662, 0.867)	30
Unemployment declined	All respondents	Correct	0.500	0.631	0.047	(0.533, 0.729)	25
			0.550	0.661	0.028	(0.606, 0.716)	57
			0.650	0.695	0.024	(0.647, 0.743)	79
			0.750	0.702	0.022	(0.658, 0.746)	120
			0.850	0.769	0.019	(0.732, 0.806)	146
			0.945	0.802	0.020	(0.763, 0.842)	121
			1.000	0.887	0.015	(0.858, 0.916)	144
		Incorrect	0.500	0.480	0.060	(0.351, 0.609)	16
			0.550	0.542	0.035	(0.469, 0.616)	22
			0.650	0.602	0.049	(0.499, 0.705)	21
			0.750	0.590	0.047	(0.495, 0.685)	37
			0.850	0.564	0.043	(0.476, 0.652)	41
			0.945	0.644	0.035	(0.572, 0.716)	32
			1.000	0.567	0.140	(0.235, 0.900)	8
	Democrat	Correct	0.500	0.587	0.058	(0.460, 0.714)	13
			0.550	0.664	0.036	(0.591, 0.737)	37
			0.650	0.690	0.030	(0.631, 0.750)	53
			0.750	0.712	0.026	(0.659, 0.765)	76
			0.850	0.744	0.028	(0.687, 0.800)	76
			0.945	0.815	0.029	(0.758, 0.872)	59
			1.000	0.883	0.029	(0.825, 0.940)	54
		Incorrect	0.500	0.411	0.083	(0.216, 0.607)	8
			0.550	0.574	0.037	(0.495, 0.653)	18
			0.650	0.607	0.071	(0.453, 0.761)	13
			0.750	0.589	0.061	(0.463, 0.715)	22
			0.850	0.526	0.071	(0.378, 0.674)	20
			0.945	0.674	0.051	(0.561, 0.787)	12
			1.000	0.590	0.140	(0.229, 0.951)	6
	Republican	Correct	0.500	0.684	0.092	(0.429, 0.939)	5
			0.550	0.641	0.053	(0.526, 0.756)	14
			0.650	0.684	0.061	(0.553, 0.815)	15
			0.750	0.682	0.049	(0.582, 0.781)	33
			0.850	0.789	0.028	(0.733, 0.844)	54
			0.945	0.782	0.030	(0.723, 0.841)	56
			1.000	0.882	0.019	(0.845, 0.919)	74
		Incorrect	0.500	0.650	0.160	(-1.383, 2.683)	2
			0.550	0.290	0.000	(NA, NA)	1
			0.650	0.608	0.080	(0.403, 0.814)	6
			0.750	0.595	0.083	(0.410, 0.779)	11
			0.850	0.595	0.054	(0.481, 0.709)	20
			0.945	0.632	0.050	(0.526, 0.738)	19
			1.000	0.500	0.500	(-5.853, 6.853)	2

Appendix D

Appendix to Chapter 6

D.1 Analytic Evidence on the Properties of Certainty Bias

This section uses algebra to prove some facts about certainty bias that are described in the main text and used as a basis for decomposing certainty bias in parts of the results.

Rewriting the average belief.

Recall that in the main text, the average belief for group j was defined as $\frac{1}{N_j}\sum p_i$. Further, recall that p_i can be re-written as $a_i c_i + (1-a_i)(1-c_i)$, where a_i is an indicator for a correct answer and c_i is certainty.

This allows the average belief to be rewritten as follows:

$$\begin{aligned}\frac{1}{N_j}\sum p_i &= \frac{1}{N_j}\Big(\sum a_i c_i + \sum (1-a_i)(1-c_i)\Big)\\ &= \frac{1}{N_j}\Big(N_j\,\mathbb{E}[a_i c_i] + N_j\,\mathbb{E}[(1-a_i)(1-c_i)]\Big)\\ &= \mathbb{E}[a_i c_i] + \mathbb{E}[(1-a_i)(1-c_i)]\end{aligned}$$

where $\mathbb{E}[\cdot]$, the expectation operator, simply takes the average. By the law of total probability, this can be further rewritten as

$$\begin{aligned}&= \mathbb{E}[a_i c_i|a=1]Pr(a=1) + \mathbb{E}[a_i c_i|a=0]Pr(a=0)\\ &\quad + \mathbb{E}[(1-a_i)(1-c_i)|a=1]Pr(a=1) + \mathbb{E}[(1-a_i)(1-c_i)|a=0]Pr(a=0).\end{aligned}$$

Because the middle two terms equal zero, this reduces to

$$= \mathbb{E}[a_i c_i|a=1]Pr(a=1) + \mathbb{E}[(1-a_i)(1-c_i)|a=0]Pr(a=0),$$

and because a_i is a constant in both terms, this can be rewritten as

$$= \mathbb{E}[c_i|a=1]Pr(a=1) + \mathbb{E}[(1-c_i)|a=0](1-Pr(a=1)). \tag{D.1}$$

Verbally, expression (D.1) says that the average belief can be written as a weighted sum of the belief among respondents who answer correctly and incorrectly.

To express this rewritten version of party j's average belief more compactly, define $\bar{p}_j$ as the average belief, $\bar{a}_j$ as the probability that a respondent who preferred party j answers correctly, $\bar{c}_{1j}$ as average certainty among respondents who answer correctly, and $\bar{c}_{0j}$ as average certainty among respondents who answer incorrectly. This allows (D.1) to be expressed as

$$\bar{p}_j = \bar{a}_j \bar{c}_{1j} + (1 - \bar{a}_j)\bar{c}_{0j}. \tag{D.2}$$

Decomposing the partisan difference.

In the main text, the partisan belief difference was expressed as

$$\Delta = \frac{1}{N_D} \sum_{i \in \mathbb{D}} p_i - \frac{1}{N_R} \sum_{i \in \mathbb{R}} p_i. \tag{D.3}$$

Based on the above, this can be rewritten as

$$\begin{aligned}
\Delta &= \Big(\bar{a}_D \bar{c}_{1D} + (1 - \bar{a}_D)(1 - \bar{c}_{0D}\Big) - \Big(\bar{a}_R \bar{c}_{1R} + (1 - \bar{a}_R)(1 - \bar{c}_{0R})\Big) \\
&= \Big(\bar{a}_D \bar{c}_{1D} - \bar{a}_R \bar{c}_{1R}\Big) + \Big((1 - \bar{a}_D)(1 - \bar{c}_{0D}) - (1 - \bar{a}_R)(1 - \bar{c}_{0R})\Big) \\
&= \bar{c}_{1D}(\bar{a}_D - \bar{a}_R) + (\bar{c}_{1D} - \bar{c}_{1R})\bar{a}_R + (1 - \bar{c}_{0D})(\bar{a}_R - \bar{a}_D) + (\bar{c}_{0R} - \bar{c}_{0D})(1 - \bar{a}_R) \\
&= \underbrace{(\bar{c}_{1D} + \bar{c}_{0D} - 1)(\bar{a}_D - \bar{a}_R)}_{\text{Term 1}} + \underbrace{(\bar{c}_{1D} - \bar{c}_{1R})\bar{a}_R}_{\text{Term 2}} + \underbrace{(\bar{c}_{0R} - \bar{c}_{0D})(1 - \bar{a}_R)}_{\text{Term 3}}.
\end{aligned} \tag{D.4}$$

In expression (D.4), the first term is the partisan belief difference that would realize if both parties were equally certain of their correct and incorrect answers; that is, if $\bar{c}_{1R} = \bar{c}_{1D}$ and $\bar{c}_{0R} = \bar{c}_{0D}$. To see this, note that if both of these conditions hold, the second and third terms equal zero. Given this, the second and third terms can be thought of as the contribution that partisan differences in certainty make to the partisan belief difference.

To capture this, define the first term as Δ_{EC}, the belief difference under the equal-certainty counterfactual:

$$\Delta_{EC} \equiv (\bar{c}_{1D} + \bar{c}_{0D} - 1)(\bar{a}_D - \bar{a}_R).$$

In turn, this implies that the sum of the second and third terms can be written as $\Delta - \Delta_{EC}$. Define this as the effect of the certainty gap on the partisan belief difference:

$$\Delta - \Delta_{EC} = (\bar{c}_{1D} - \bar{c}_{1R})\bar{a}_R + (\bar{c}_{0R} - \bar{c}_{0D})(1 - \bar{a}_R). \tag{D.5}$$

Before using these expressions to examine the nature of certainty bias, it is worth making observing that when $\bar{c}_{1R} = \bar{c}_{1D} = 1$, expression (D.4) simplifies to $\bar{a}_D - \bar{a}_R$, which is the partisan difference in the probability of answering correctly. In the main text, this quantity was defined as Δ_{BG}, the partisan difference in best guesses. This means that the partisan difference in best guesses is a special case of the equal-certainty counterfactual.

Effect of generalized uncertainty.

In the main text, certainty bias was defined as $\Delta - \Delta_{BG}$. Based on the above observations, this can be rewritten as the sum of two components:

$$\Delta_{BG} - \Delta = \underbrace{\Delta_{BG} - \Delta_{EC}}_{\text{Effect of general uncert.}} + \underbrace{\Delta_{EC} - \Delta}_{\text{Effect of cert. gap}} \tag{D.6}$$

Just as in the main text, these components are defined as the effect of generalized uncertainty and the effect of the partisan certainty gap. Note that the effect of the certainty gap on certainty bias is the inverse of its effect on the partisan belief difference as defined in equation (D.5).

To understand how generalized uncertainty contributes to certainty bias, express it using the terms above:

$$\begin{aligned}\Delta_{BG} - \Delta_{EC} &= (\bar{a}_D - \bar{a}_R) - (\bar{c}_{1D} + \bar{c}_{0D} - 1)(\bar{a}_D - \bar{a}_R) \\ &= (2 - \bar{c}_{1D} - \bar{c}_{0D})(\bar{a}_D - \bar{a}_R) \end{aligned} \tag{D.7}$$

Note that whenever all respondents are completely certain of their answers, $\bar{c}_{1j} = \bar{c}_{1j} = 1$, and expression (D.7) is equal to zero. Whenever any respondent expresses uncertainty, at least one of $\bar{c}_{1j}$ or $\bar{c}_{0j}$ is less than 1, and (D.7) takes the same sign as the partisan difference in best guesses. If all respondents were completely uncertain, $\bar{c}_{1j} = \bar{c}_{1j} = 1/2$, and expression (D.7) is equal to the partisan difference in best guesses, $\bar{a}_D - \bar{a}_R$.

This gives rise to the following:

> **Theorem 1.** The effect of generalized uncertainty is bounded on the interval $[0, \bar{a}_D - \bar{a}_R]$. Whenever at least one respondent is uncertain, the interval is $(0, \bar{a}_D - \bar{a}_R]$.

> **Lemma 1.** The effect of generalized uncertainty on certainty bias always takes the same sign as the partisan difference in best guesses.

Lemma 1 can be thought of as the reinforcement effect of generalized uncertainty on measured partisan belief differences.

Effect of partisan differences in certainty.

Before examining the role of partisan certainty gaps, it is worth being clear about the relationship between the effect of such gaps on the partisan belief difference and the effect of such gaps on certainty bias. In the definition of certainty bias ($\Delta_{BG} - \Delta$), the partisan belief difference (Δ) enters as a negative. Anything that makes Δ more negative will make certainty bias more positive; anything that makes Δ more positive will make certainty bias more negative.

In practice, it is easier to first examine how partisan certainty gaps affect partisan belief differences, then use this "inverse property" to draw out the implications for the effect on certainty bias.

Consider the first half of expression (D.5), $(\bar{c}_{1D} - \bar{c}_{1R})\bar{a}_R$, which is the effect of certainty gaps among respondents who answered correctly on the partisan belief difference. The first derivative with respect to Democrats' greater certainty about correct answers, $\frac{\partial}{\partial(\bar{c}_{1D} - \bar{c}_{1R})}$, is simply $\bar{a}_R$, which is weakly positive (and always positive when at least one Republican answers correctly). This means that when the partisan difference is defined as it was in expression (D.3) — Democrat minus Republican — any case in which Democrats are more certain of correct answers than Republicans will have a positive effect on the partisan belief difference. Conversely, any case in which Republicans are more certain of correct answers than Democrats will have a negative effect.

Now consider the second half of (D.5), $(\bar{c}_{0R} - \bar{c}_{0D})(1 - \bar{a}_R)$, which is the effect of certainty gaps among respondents who answered *in*correctly on the partisan belief difference. The first derivative with respect to Republicans' greater certainty about incorrect answers, $\frac{\partial}{\partial(\bar{c}_{1D} - \bar{c}_{1R})}$, is simply $(1 - \bar{a}_R)$, which is weakly positive (and always positive when at least one Republican answers incorrectly). This means that when the partisan difference is defined as it was in expression (D.3) — Democrat minus Republican — any case in which Republicans are more certain of incorrect answers than Democrats will have a positive effect on the partisan belief difference. Conversely, any case in which

Democrats are more certain of incorrect answers than Republicans will have a negative effect.

This means that the effect of partisan certainty gaps has predictable effects on partisan belief differences. Whether those effects reinforce or counteract the partisan belief difference depends on the sign of the difference. When Democrats answer correctly more often than Republicans, the partisan belief difference is reinforced (counteracted) whenever $(\bar{c}_{1D} - \bar{c}_{1R})$ and $(\bar{c}_{0R} - \bar{c}_{0D})$ are both positive (negative). When Republicans answer correctly more often than Democrats, the partisan belief difference is reinforced (counteracted) whenever $(\bar{c}_{1D} - \bar{c}_{1R})$ and $(\bar{c}_{0R} - \bar{c}_{0D})$ are both negative (positive).

Now, note again that the "inverse property" means that any effect on the partisan belief difference has the opposite effect on certainty bias.

This gives rise to the following:

> **Theorem 2.** Partisan certainty gaps reinforce the partisan belief difference, and counteract certainty bias, whenever two conditions are met: the party that answers correctly more often is also (a) more certain of correct answers and (b) less certain of incorrect answers.

> **Theorem 3.** Partisan certainty gaps counteract the partisan belief difference, and reinforce certainty bias, whenever two conditions are met: the party that answers correctly more often is also (a) less certain of correct answers and (b) more certain of incorrect answers.

As discussed in the text, theories of partisan bias predict that partisans know more about facts that benefit their party, and are less susceptible to believing false things that would be inconvenient for their party if true. This gives rise to the expectation that the conditions in Theorem 2 will be far more likely to hold than the conditions in Theorem 3, i.e., that partisan certainty gaps will usually reinforce partisan belief differences and counteract certainty bias.

D.2 Simulation Study

To examine the degree to which threshold-based measurement techniques can successfully approximate the partisan belief difference, I conducted a simulation study using the same data analyzed in the main text. To each question, I applied the two techniques described there: the "drop" strategy that calculates the belief difference only among those whose certainty level exceeds the threshold (i.e., those who do not say DK), and the "treat-as-uncertain" strategy that scores all respondents whose certainty level is beneath the threshold as an 0.5, as if they were completely uncertain. I calculated these estimates for each possible choice of thresholds between 0.5 and 1, rounded to the nearest hundredth.

Figures D.1 and D.2 display the results. In each panel, the x-axis is the threshold and the y-axis is the estimate of the belief difference. The dashed line in each panel is the probabilistic belief difference. The threshold-based estimates are represented by the dotted line (drop) and the solid line (score as 0.5). At the leftmost point on the x-axis, 0.5, these two lines always touch because a threshold of 0.5 is equivalent to the best guess strategy.

The results suggest that the "drop" strategy is generally prone to over-estimating the partisan belief difference, while the "score as 0.5" strategy will usually get things right at least some threshold. For the score as 0.5 strategy, the point of intersection usually falls at the middle of the scale, suggesting that the low- and high-threshold strategies used in most research would over- or under-estimate the difference. The drop strategy is wildly inconsistent, generally inflating partisan belief differences but doing so by drastically different magnitudes.

There is substantial variation across questions. For some questions, any choice of strategy would be about right, while for others, the choice of strategy matters a lot. Occasionally, the drop strategy reduces partisan differenceswhile the score as 0.5 strategy increases them.

Together, these results suggest that the score as 0.5 strategy could sometimes be a reasonable approximation, given an informed choice of measurement technologies. The drop strategy is highly distortive and should probably not be used.

Figure D.1: Simulation study results, controversy questions.

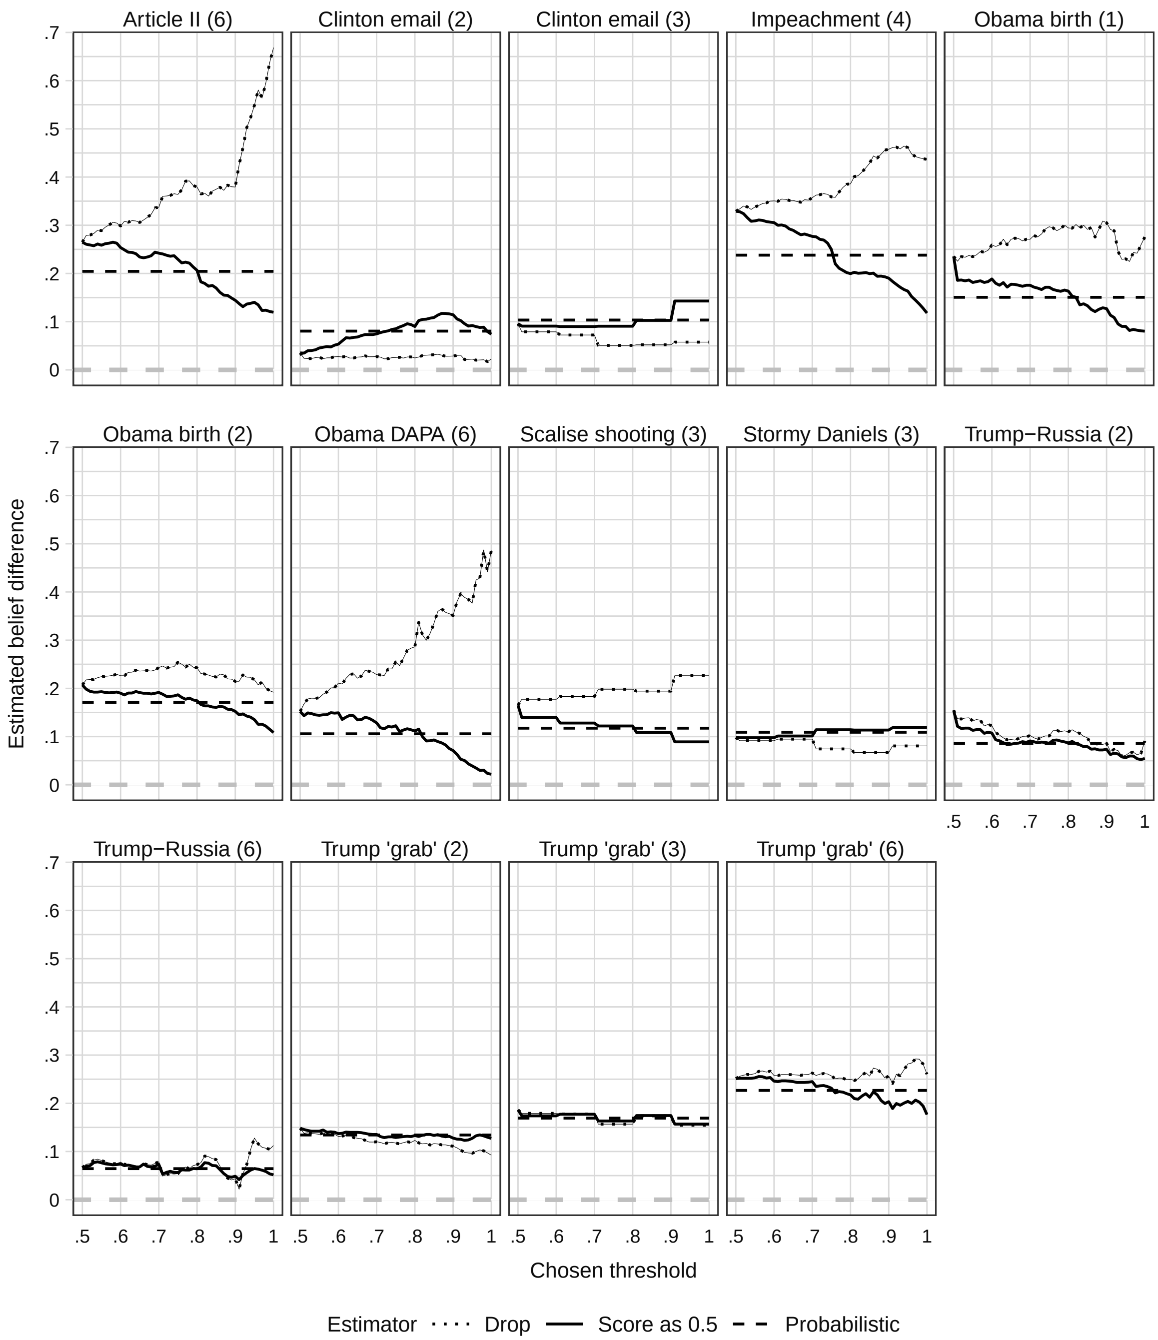

Figure D.2: Simulation study results, economic questions.

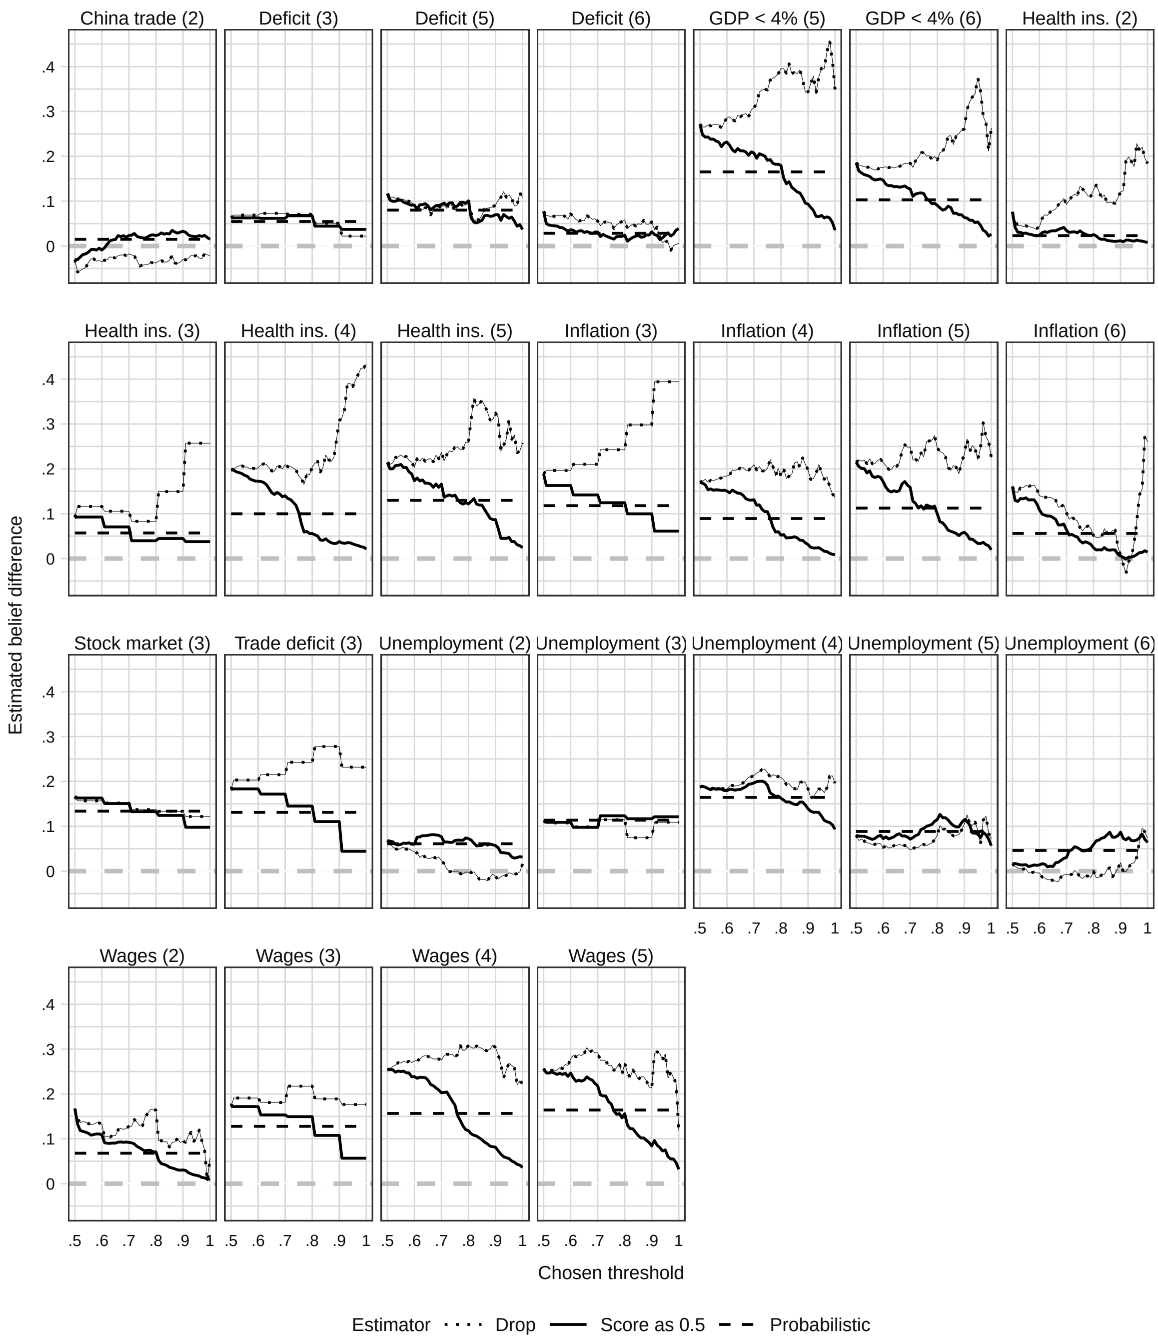

D.3 Supporting Tables

Table D.1: Estimates plotted in Figure 6.2

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
Budget deficit increased	Democrat	3	Probabilistic	(Intercept)	0.763	0.013	0.000	(0.738, 0.789)
				Republican	-0.055	0.020	0.005	(-0.093, -0.016)
			Best guess	(Intercept)	0.824	0.018	0.000	(0.789, 0.859)
				Republican	-0.062	0.027	0.022	(-0.116, -0.009)
China trade deficit	Democrat	2	Probabilistic	(Intercept)	0.657	0.007	0.000	(0.642, 0.671)
				Republican	0.015	0.012	0.201	(-0.008, 0.038)
			Best guess	(Intercept)	0.867	0.013	0.000	(0.841, 0.894)
				Republican	-0.030	0.020	0.135	(-0.070, 0.009)
Clinton email	Republican	2	Probabilistic	(Intercept)	0.800	0.009	0.000	(0.782, 0.818)
				Republican	0.080	0.012	0.000	(0.056, 0.105)
			Best guess	(Intercept)	0.930	0.010	0.000	(0.911, 0.949)
				Republican	0.035	0.013	0.007	(0.010, 0.061)
		3	Probabilistic	(Intercept)	0.815	0.011	0.000	(0.795, 0.836)
				Republican	0.103	0.013	0.000	(0.077, 0.129)
			Best guess	(Intercept)	0.873	0.014	0.000	(0.846, 0.900)
				Republican	0.096	0.016	0.000	(0.065, 0.127)
Health ins. decreased	Democrat	2	Probabilistic	(Intercept)	0.507	0.008	0.000	(0.490, 0.524)
				Republican	-0.023	0.012	0.056	(-0.047, 0.001)
			Best guess	(Intercept)	0.535	0.020	0.000	(0.496, 0.574)
				Republican	-0.076	0.028	0.008	(-0.132, -0.020)
		3	Probabilistic	(Intercept)	0.515	0.013	0.000	(0.490, 0.540)
				Republican	-0.057	0.018	0.002	(-0.093, -0.021)
			Best guess	(Intercept)	0.543	0.021	0.000	(0.503, 0.583)
				Republican	-0.092	0.031	0.003	(-0.153, -0.031)
		4	Probabilistic	(Intercept)	0.514	0.011	0.000	(0.494, 0.535)
				Republican	-0.100	0.015	0.000	(-0.130, -0.070)
			Best guess	(Intercept)	0.532	0.018	0.000	(0.496, 0.567)
				Republican	-0.198	0.027	0.000	(-0.251, -0.145)
Impeachment witnesses	Democrat	4	Probabilistic	(Intercept)	0.551	0.013	0.000	(0.527, 0.576)
				Republican	-0.238	0.018	0.000	(-0.273, -0.203)
			Best guess	(Intercept)	0.611	0.018	0.000	(0.575, 0.646)

Table D.1: Estimates plotted in Figure 6.2 (continued)

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
				Republican	-0.329	0.026	0.000	(-0.380, -0.278)
Inflation below avg.	Republican	3	Probabilistic	(Intercept)	0.382	0.011	0.000	(0.359, 0.404)
				Republican	0.118	0.019	0.000	(0.082, 0.154)
			Best guess	(Intercept)	0.317	0.019	0.000	(0.279, 0.355)
				Republican	0.189	0.030	0.000	(0.129, 0.248)
		4	Probabilistic	(Intercept)	0.425	0.009	0.000	(0.407, 0.443)
				Republican	0.089	0.015	0.000	(0.059, 0.119)
			Best guess	(Intercept)	0.326	0.017	0.000	(0.293, 0.360)
				Republican	0.170	0.027	0.000	(0.118, 0.223)
Obama birth certificate	Democrat	1	Probabilistic	(Intercept)	0.690	0.016	0.000	(0.659, 0.721)
				Republican	-0.151	0.025	0.000	(-0.199, -0.102)
			Best guess	(Intercept)	0.787	0.024	0.000	(0.740, 0.834)
				Republican	-0.235	0.040	0.000	(-0.315, -0.156)
		2	Probabilistic	(Intercept)	0.725	0.013	0.000	(0.700, 0.750)
				Republican	-0.171	0.019	0.000	(-0.208, -0.134)
			Best guess	(Intercept)	0.770	0.017	0.000	(0.738, 0.803)
				Republican	-0.207	0.026	0.000	(-0.258, -0.156)
Softball shooting	Republican	3	Probabilistic	(Intercept)	0.561	0.013	0.000	(0.536, 0.586)
				Republican	0.117	0.019	0.000	(0.081, 0.154)
			Best guess	(Intercept)	0.507	0.021	0.000	(0.466, 0.547)
				Republican	0.165	0.030	0.000	(0.106, 0.224)
Stock market increased	Republican	3	Probabilistic	(Intercept)	0.678	0.015	0.000	(0.648, 0.709)
				Republican	0.134	0.019	0.000	(0.096, 0.172)
			Best guess	(Intercept)	0.716	0.022	0.000	(0.672, 0.759)
				Republican	0.166	0.027	0.000	(0.113, 0.219)
Stormy Daniels payment	Democrat	3	Probabilistic	(Intercept)	0.854	0.009	0.000	(0.836, 0.871)
				Republican	-0.109	0.016	0.000	(-0.140, -0.078)
			Best guess	(Intercept)	0.924	0.011	0.000	(0.902, 0.945)
				Republican	-0.098	0.021	0.000	(-0.139, -0.057)
Trade deficit increased	Democrat	3	Probabilistic	(Intercept)	0.637	0.015	0.000	(0.608, 0.667)
				Republican	-0.131	0.022	0.000	(-0.175, -0.087)
			Best guess	(Intercept)	0.687	0.023	0.000	(0.641, 0.733)
				Republican	-0.182	0.035	0.000	(-0.250, -0.114)

Table D.1: Estimates plotted in Figure 6.2 (continued)

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
Trump collusion	Republican	2	Probabilistic	(Intercept)	0.622	0.013	0.000	(0.596, 0.647)
				Republican	0.086	0.019	0.000	(0.049, 0.122)
			Best guess	(Intercept)	0.607	0.019	0.000	(0.568, 0.645)
				Republican	0.155	0.026	0.000	(0.103, 0.206)
Trump said 'grab them'	Democrat	2	Probabilistic	(Intercept)	0.863	0.008	0.000	(0.846, 0.879)
				Republican	-0.134	0.015	0.000	(-0.164, -0.105)
			Best guess	(Intercept)	0.928	0.010	0.000	(0.909, 0.948)
				Republican	-0.148	0.020	0.000	(-0.187, -0.110)
		3	Probabilistic	(Intercept)	0.857	0.010	0.000	(0.837, 0.876)
				Republican	-0.169	0.019	0.000	(-0.206, -0.132)
			Best guess	(Intercept)	0.897	0.013	0.000	(0.872, 0.921)
				Republican	-0.186	0.025	0.000	(-0.235, -0.138)
Unemp. decreased	Republican	2	Probabilistic	(Intercept)	0.672	0.009	0.000	(0.655, 0.689)
				Republican	0.061	0.013	0.000	(0.035, 0.087)
			Best guess	(Intercept)	0.829	0.015	0.000	(0.800, 0.857)
				Republican	0.069	0.020	0.001	(0.030, 0.107)
		3	Probabilistic	(Intercept)	0.643	0.013	0.000	(0.618, 0.667)
				Republican	0.114	0.019	0.000	(0.076, 0.151)
			Best guess	(Intercept)	0.692	0.019	0.000	(0.654, 0.729)
				Republican	0.112	0.027	0.000	(0.060, 0.164)
		4	Probabilistic	(Intercept)	0.553	0.011	0.000	(0.532, 0.575)
				Republican	0.164	0.017	0.000	(0.131, 0.198)
			Best guess	(Intercept)	0.571	0.018	0.000	(0.537, 0.606)
				Republican	0.188	0.025	0.000	(0.139, 0.237)
Wages increased	Republican	2	Probabilistic	(Intercept)	0.569	0.008	0.000	(0.555, 0.584)
				Republican	0.068	0.011	0.000	(0.046, 0.090)
			Best guess	(Intercept)	0.687	0.018	0.000	(0.651, 0.723)
				Republican	0.167	0.023	0.000	(0.121, 0.213)
		3	Probabilistic	(Intercept)	0.591	0.011	0.000	(0.569, 0.612)
				Republican	0.128	0.016	0.000	(0.096, 0.159)
			Best guess	(Intercept)	0.673	0.019	0.000	(0.636, 0.711)
				Republican	0.172	0.026	0.000	(0.122, 0.223)

Table D.1: Estimates plotted in Figure 6.2 (continued)

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
		4	Probabilistic	(Intercept)	0.540	0.009	0.000	(0.522, 0.559)
				Republican	0.157	0.013	0.000	(0.130, 0.183)
			Best guess	(Intercept)	0.587	0.018	0.000	(0.552, 0.621)
				Republican	0.254	0.023	0.000	(0.208, 0.299)

Table D.2: Estimates plotted in Figure 6.3

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
Budget deficit increased	Democrat	5	Best guess, stated	(Intercept)	0.874	0.023	0.000	(0.830, 0.919)
				Republican	-0.117	0.038	0.002	(-0.192, -0.041)
			Probabilistic, stated	(Intercept)	0.776	0.016	0.000	(0.744, 0.807)
				Republican	-0.082	0.025	0.001	(-0.132, -0.032)
			Best guess, revealed	(Intercept)	0.800	0.027	0.000	(0.746, 0.854)
				Republican	-0.084	0.042	0.050	(-0.167, -0.000)
			Probabilistic, revealed	(Intercept)	0.691	0.019	0.000	(0.654, 0.728)
				Republican	-0.043	0.029	0.143	(-0.101, 0.015)
Economic growth below 4%	Democrat	5	Best guess, stated	(Intercept)	0.624	0.033	0.000	(0.559, 0.690)
				Republican	-0.272	0.048	0.000	(-0.366, -0.177)
			Probabilistic, stated	(Intercept)	0.556	0.019	0.000	(0.520, 0.593)
				Republican	-0.157	0.028	0.000	(-0.212, -0.103)
			Best guess, revealed	(Intercept)	0.545	0.034	0.000	(0.477, 0.612)
				Republican	-0.181	0.049	0.000	(-0.278, -0.085)
			Probabilistic, revealed	(Intercept)	0.548	0.022	0.000	(0.505, 0.590)
				Republican	-0.106	0.032	0.001	(-0.168, -0.043)
Inflation below avg.	Republican	5	Best guess, stated	(Intercept)	0.403	0.033	0.000	(0.337, 0.469)
				Republican	0.215	0.048	0.000	(0.120, 0.309)
			Probabilistic, stated	(Intercept)	0.452	0.019	0.000	(0.415, 0.489)
				Republican	0.107	0.027	0.000	(0.054, 0.161)
			Best guess, revealed	(Intercept)	0.463	0.034	0.000	(0.396, 0.530)
				Republican	0.093	0.049	0.059	(-0.004, 0.190)
			Probabilistic, revealed	(Intercept)	0.513	0.021	0.000	(0.472, 0.553)
				Republican	0.043	0.030	0.147	(-0.015, 0.102)

Table D.2: Estimates plotted in Figure 6.3 (continued)

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
Unemployment decreased	Republican	5	Best guess, stated	(Intercept)	0.727	0.030	0.000	(0.667, 0.787)
				Republican	0.080	0.041	0.053	(-0.001, 0.161)
			Probabilistic, stated	(Intercept)	0.637	0.020	0.000	(0.598, 0.676)
				Republican	0.089	0.030	0.003	(0.030, 0.148)
			Best guess, revealed	(Intercept)	0.681	0.032	0.000	(0.618, 0.743)
				Republican	0.107	0.043	0.013	(0.022, 0.191)
			Probabilistic, revealed	(Intercept)	0.632	0.021	0.000	(0.592, 0.673)
				Republican	0.073	0.030	0.014	(0.014, 0.131)
Budget deficit increased	Democrat	6	Best guess, stated	(Intercept)	0.827	0.018	0.000	(0.793, 0.862)
				Republican	-0.077	0.031	0.012	(-0.137, -0.017)
			Probabilistic, stated	(Intercept)	0.727	0.011	0.000	(0.705, 0.749)
				Republican	-0.033	0.019	0.089	(-0.070, 0.005)
			Best guess, revealed	(Intercept)	0.778	0.019	0.000	(0.740, 0.816)
				Republican	-0.108	0.033	0.001	(-0.173, -0.042)
			Probabilistic, revealed	(Intercept)	0.669	0.013	0.000	(0.644, 0.695)
				Republican	-0.049	0.021	0.023	(-0.091, -0.007)
Economic growth below 4%	Democrat	6	Best guess, stated	(Intercept)	0.558	0.023	0.000	(0.513, 0.603)
				Republican	-0.186	0.036	0.000	(-0.257, -0.115)
			Probabilistic, stated	(Intercept)	0.537	0.012	0.000	(0.513, 0.561)
				Republican	-0.093	0.022	0.000	(-0.136, -0.051)
			Best guess, revealed	(Intercept)	0.515	0.023	0.000	(0.470, 0.560)
				Republican	-0.120	0.036	0.001	(-0.191, -0.048)
			Probabilistic, revealed	(Intercept)	0.535	0.013	0.000	(0.510, 0.560)
				Republican	-0.064	0.022	0.003	(-0.106, -0.021)
Inflation below avg.	Republican	6	Best guess, stated	(Intercept)	0.452	0.023	0.000	(0.406, 0.497)
				Republican	0.161	0.036	0.000	(0.090, 0.231)
			Probabilistic, stated	(Intercept)	0.496	0.012	0.000	(0.473, 0.520)
				Republican	0.047	0.020	0.018	(0.008, 0.086)
			Best guess, revealed	(Intercept)	0.488	0.023	0.000	(0.443, 0.534)
				Republican	0.095	0.036	0.009	(0.024, 0.166)
			Probabilistic, revealed	(Intercept)	0.512	0.013	0.000	(0.487, 0.537)
				Republican	0.056	0.020	0.006	(0.016, 0.096)
Unemployment decreased	Republican	6	Best guess, stated	(Intercept)	0.788	0.019	0.000	(0.751, 0.825)
				Republican	0.016	0.029	0.575	(-0.041, 0.074)
			Probabilistic, stated	(Intercept)	0.677	0.012	0.000	(0.653, 0.701)
				Republican	0.051	0.021	0.014	(0.010, 0.092)
			Best guess, revealed	(Intercept)	0.741	0.020	0.000	(0.701, 0.781)

Table D.2: Estimates plotted in Figure 6.3 (continued)

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
				Republican	0.022	0.032	0.487	(-0.040, 0.084)
			Probabilistic, revealed	(Intercept)	0.680	0.013	0.000	(0.655, 0.705)
				Republican	0.029	0.020	0.150	(-0.010, 0.068)
Health ins. decreased	Democrat	5	Best guess, stated	(Intercept)	0.609	0.033	0.000	(0.544, 0.674)
				Republican	-0.214	0.048	0.000	(-0.309, -0.120)
			Probabilistic, stated	(Intercept)	0.554	0.020	0.000	(0.515, 0.593)
				Republican	-0.119	0.028	0.000	(-0.175, -0.063)
			Best guess, revealed	(Intercept)	0.600	0.033	0.000	(0.535, 0.665)
				Republican	-0.190	0.048	0.000	(-0.285, -0.095)
			Probabilistic, revealed	(Intercept)	0.560	0.022	0.000	(0.517, 0.603)
				Republican	-0.082	0.032	0.011	(-0.145, -0.019)
Wages increased	Republican	5	Best guess, stated	(Intercept)	0.545	0.034	0.000	(0.477, 0.612)
				Republican	0.256	0.045	0.000	(0.169, 0.344)
			Probabilistic, stated	(Intercept)	0.534	0.018	0.000	(0.498, 0.571)
				Republican	0.155	0.026	0.000	(0.104, 0.205)
			Best guess, revealed	(Intercept)	0.573	0.034	0.000	(0.506, 0.640)
				Republican	0.228	0.044	0.000	(0.141, 0.316)
			Probabilistic, revealed	(Intercept)	0.558	0.021	0.000	(0.517, 0.599)
				Republican	0.107	0.029	0.000	(0.051, 0.163)
Obama DAPA statement	Republican	6	Best guess, stated	(Intercept)	0.369	0.022	0.000	(0.326, 0.412)
				Republican	0.152	0.036	0.000	(0.081, 0.223)
			Probabilistic, stated	(Intercept)	0.456	0.011	0.000	(0.435, 0.477)
				Republican	0.099	0.019	0.000	(0.063, 0.136)
			Best guess, revealed	(Intercept)	0.413	0.023	0.000	(0.369, 0.457)
				Republican	0.131	0.036	0.000	(0.060, 0.202)
			Probabilistic, revealed	(Intercept)	0.455	0.011	0.000	(0.433, 0.477)
				Republican	0.062	0.020	0.002	(0.024, 0.101)
Trump Article II claim	Democrat	6	Best guess, stated	(Intercept)	0.741	0.020	0.000	(0.701, 0.781)
				Republican	-0.265	0.035	0.000	(-0.333, -0.197)
			Probabilistic, stated	(Intercept)	0.682	0.012	0.000	(0.657, 0.706)
				Republican	-0.196	0.022	0.000	(-0.239, -0.154)
			Best guess, revealed	(Intercept)	0.747	0.020	0.000	(0.708, 0.787)
				Republican	-0.268	0.035	0.000	(-0.336, -0.200)
			Probabilistic, revealed	(Intercept)	0.650	0.013	0.000	(0.624, 0.676)
				Republican	-0.172	0.022	0.000	(-0.215, -0.129)
Trump collusion	Republican	6	Best guess, stated	(Intercept)	0.609	0.022	0.000	(0.565, 0.653)

Table D.2: Estimates plotted in Figure 6.3 (continued)

Question	Valence	Survey	Estimator	Term	Estimate	SE	p	CI
				Republican	0.066	0.035	0.057	(-0.002, 0.135)
			Probabilistic, stated	(Intercept)	0.595	0.015	0.000	(0.566, 0.625)
				Republican	0.063	0.024	0.009	(0.016, 0.110)
			Best guess, revealed	(Intercept)	0.605	0.023	0.000	(0.560, 0.649)
				Republican	0.058	0.035	0.101	(-0.011, 0.127)
			Probabilistic, revealed	(Intercept)	0.608	0.015	0.000	(0.579, 0.638)
				Republican	0.031	0.024	0.201	(-0.017, 0.079)
Trump said 'grab them'	Democrat	6	Best guess, stated	(Intercept)	0.827	0.017	0.000	(0.793, 0.861)
				Republican	-0.251	0.033	0.000	(-0.316, -0.186)
			Probabilistic, stated	(Intercept)	0.808	0.013	0.000	(0.782, 0.834)
				Republican	-0.222	0.025	0.000	(-0.272, -0.173)
			Best guess, revealed	(Intercept)	0.811	0.018	0.000	(0.776, 0.846)
				Republican	-0.209	0.033	0.000	(-0.274, -0.144)
			Probabilistic, revealed	(Intercept)	0.753	0.014	0.000	(0.725, 0.781)
				Republican	-0.164	0.025	0.000	(-0.213, -0.115)

Table D.3: Estimates plotted in Figure 6.4

Category	Answer	Valence	Estimate	SE	p	CI
All questions	Correct	Congenial	0.824	0.002	0.000	(0.821, 0.828)
		Uncongenial	0.787	0.002	0.000	(0.783, 0.791)
		Difference	0.037	0.002	0.000	(0.033, 0.041)
	Incorrect	Congenial	0.737	0.003	0.000	(0.731, 0.742)
		Uncongenial	0.738	0.002	0.000	(0.734, 0.742)
		Difference	-0.001	0.003	0.621	(-0.007, 0.004)
Controversies	Correct	Congenial	0.876	0.002	0.000	(0.872, 0.881)
		Uncongenial	0.832	0.003	0.000	(0.826, 0.838)
		Difference	0.044	0.003	0.000	(0.038, 0.050)
	Incorrect	Congenial	0.739	0.005	0.000	(0.729, 0.749)
		Uncongenial	0.754	0.004	0.000	(0.746, 0.762)
		Difference	-0.015	0.006	0.008	(-0.027, -0.004)
Economic	Correct	Congenial	0.786	0.002	0.000	(0.782, 0.791)
		Uncongenial	0.758	0.002	0.000	(0.753, 0.762)
		Difference	0.029	0.003	0.000	(0.024, 0.034)
	Incorrect	Congenial	0.736	0.003	0.000	(0.730, 0.742)
		Uncongenial	0.730	0.002	0.000	(0.725, 0.735)
		Difference	0.006	0.003	0.080	(-0.001, 0.013)

Table D.4: Estimates plotted in Figure 6.5

Category	Estimator	Estimate	SE	CI
All questions	Best guess (Δ_{BG})	0.152	0.005	(0.142, 0.161)
	Equal certainty (Δ_{EC})	0.085	0.003	(0.080, 0.091)
	Probabilistic (Δ)	0.108	0.003	(0.102, 0.114)
	Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.067	0.002	(-0.071, -0.062)
	Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.023	0.002	(0.019, 0.026)
	Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.044	0.003	(-0.049, -0.038)
Controversies	Best guess (Δ_{BG})	0.170	0.007	(0.156, 0.183)
	Equal certainty (Δ_{EC})	0.104	0.004	(0.096, 0.113)
	Probabilistic (Δ)	0.138	0.005	(0.128, 0.147)
	Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.065	0.003	(-0.071, -0.060)
	Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.034	0.003	(0.027, 0.039)
	Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.032	0.004	(-0.039, -0.024)
Economic	Best guess (Δ_{BG})	0.138	0.006	(0.126, 0.151)
	Equal certainty (Δ_{EC})	0.072	0.003	(0.066, 0.079)
	Probabilistic (Δ)	0.086	0.004	(0.079, 0.094)
	Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.066	0.003	(-0.072, -0.060)
	Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.014	0.002	(0.010, 0.018)
	Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.052	0.004	(-0.059, -0.044)

Table D.5: Estimates plotted in Figure 6.6

Question	Valence	Estimator	Estimate	SE	CI
Article II (6)	D	Best guess (Δ_{BG})	0.265	0.036	(0.194, 0.333)
		Equal certainty (Δ_{EC})	0.141	0.020	(0.102, 0.180)
		Probabilistic (Δ)	0.204	0.024	(0.156, 0.253)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.124	0.017	(-0.157, -0.091)
		Difference ($\Delta - \Delta_{EC}$)	0.064	0.013	(0.039, 0.089)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.060	0.019	(-0.096, -0.022)
China trade (2)	D	Best guess (Δ_{BG})	0.030	0.020	(-0.008, 0.072)
		Equal certainty (Δ_{EC})	0.009	0.006	(-0.002, 0.020)
		Probabilistic (Δ)	-0.015	0.012	(-0.037, 0.010)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.022	0.015	(-0.052, 0.006)
		Difference ($\Delta - \Delta_{EC}$)	-0.023	0.010	(-0.044, -0.003)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.045	0.017	(-0.080, -0.013)
Deficit (3)	D	Best guess (Δ_{BG})	0.062	0.027	(0.008, 0.115)
		Equal certainty (Δ_{EC})	0.038	0.016	(0.005, 0.069)
		Probabilistic (Δ)	0.055	0.019	(0.015, 0.092)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.025	0.011	(-0.045, -0.003)
		Difference ($\Delta - \Delta_{EC}$)	0.017	0.010	(-0.004, 0.037)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.008	0.015	(-0.037, 0.022)
Deficit (5)	D	Best guess (Δ_{BG})	0.117	0.038	(0.044, 0.197)
		Equal certainty (Δ_{EC})	0.066	0.022	(0.024, 0.110)
		Probabilistic (Δ)	0.080	0.027	(0.028, 0.130)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.051	0.017	(-0.088, -0.019)
		Difference ($\Delta - \Delta_{EC}$)	0.015	0.016	(-0.017, 0.046)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.037	0.023	(-0.084, 0.010)
Deficit (6)	D	Best guess (Δ_{BG})	0.077	0.030	(0.018, 0.135)
		Equal certainty (Δ_{EC})	0.040	0.016	(0.009, 0.071)
		Probabilistic (Δ)	0.028	0.020	(-0.011, 0.067)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.037	0.014	(-0.064, -0.009)
		Difference ($\Delta - \Delta_{EC}$)	-0.012	0.011	(-0.035, 0.011)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.049	0.018	(-0.084, -0.013)
GDP $<$ 4% (5)	D	Best guess (Δ_{BG})	0.272	0.047	(0.183, 0.366)
		Equal certainty (Δ_{EC})	0.133	0.024	(0.089, 0.181)
		Probabilistic (Δ)	0.165	0.029	(0.106, 0.221)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.139	0.024	(-0.186, -0.092)
		Difference ($\Delta - \Delta_{EC}$)	0.032	0.016	(-0.001, 0.064)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.106	0.028	(-0.162, -0.055)
GDP $<$ 4% (6)	D	Best guess (Δ_{BG})	0.186	0.036	(0.114, 0.257)
		Equal certainty (Δ_{EC})	0.084	0.017	(0.051, 0.116)
		Probabilistic (Δ)	0.103	0.023	(0.059, 0.150)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.101	0.020	(-0.141, -0.062)
		Difference ($\Delta - \Delta_{EC}$)	0.019	0.012	(-0.004, 0.044)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.083	0.021	(-0.124, -0.045)
Health ins. (2)	D	Best guess (Δ_{BG})	0.076	0.028	(0.018, 0.131)
		Equal certainty (Δ_{EC})	0.024	0.009	(0.006, 0.040)
		Probabilistic (Δ)	0.023	0.012	(0.001, 0.047)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.053	0.019	(-0.090, -0.013)
		Difference ($\Delta - \Delta_{EC}$)	-0.000	0.008	(-0.017, 0.017)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.053	0.021	(-0.097, -0.009)
Health ins. (3)	D	Best guess (Δ_{BG})	0.092	0.032	(0.029, 0.155)
		Equal certainty (Δ_{EC})	0.046	0.016	(0.015, 0.077)
		Probabilistic (Δ)	0.057	0.019	(0.019, 0.094)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.046	0.016	(-0.076, -0.015)
		Difference ($\Delta - \Delta_{EC}$)	0.011	0.011	(-0.011, 0.032)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.035	0.020	(-0.075, 0.003)
Health ins. (4)	D	Best guess (Δ_{BG})	0.198	0.028	(0.142, 0.248)
		Equal certainty (Δ_{EC})	0.102	0.014	(0.073, 0.129)
		Probabilistic (Δ)	0.100	0.016	(0.070, 0.129)
		Difference ($\Delta_{EC} - \Delta_{BG}$)	-0.096	0.014	(-0.121, -0.069)
		Difference ($\Delta - \Delta_{EC}$)	-0.002	0.008	(-0.016, 0.014)
		Certainty bias ($\Delta - \Delta_{BG}$)	-0.098	0.016	(-0.128, -0.066)

Table D.5: Estimates plotted in Figure 6.6 (continued)

Category	Estimator	Estimate	SE	CI	
Health ins. (5)	D	Best guess (Δ_{BG})	0.214	0.049	(0.116, 0.313)
		Equal certainty (Δ_{EC})	0.117	0.027	(0.064, 0.172)
		Probabilistic (Δ)	0.130	0.030	(0.074, 0.188)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.097	0.022	(-0.143, -0.052)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.013	0.015	(-0.016, 0.043)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.084	0.027	(-0.134, -0.026)
Impeachment (4)	D	Best guess (Δ_{BG})	0.329	0.025	(0.280, 0.381)
		Equal certainty (Δ_{EC})	0.204	0.016	(0.173, 0.237)
		Probabilistic (Δ)	0.238	0.018	(0.204, 0.273)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.125	0.010	(-0.147, -0.106)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.034	0.010	(0.015, 0.052)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.091	0.013	(-0.118, -0.065)
Obama birth (1)	D	Best guess (Δ_{BG})	0.235	0.041	(0.156, 0.313)
		Equal certainty (Δ_{EC})	0.108	0.020	(0.068, 0.147)
		Probabilistic (Δ)	0.151	0.025	(0.100, 0.196)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.128	0.023	(-0.173, -0.084)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.043	0.018	(0.007, 0.080)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.085	0.029	(-0.141, -0.025)
Obama birth (2)	D	Best guess (Δ_{BG})	0.207	0.026	(0.157, 0.257)
		Equal certainty (Δ_{EC})	0.130	0.017	(0.097, 0.162)
		Probabilistic (Δ)	0.171	0.019	(0.132, 0.206)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.077	0.010	(-0.096, -0.057)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.041	0.012	(0.017, 0.064)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.036	0.015	(-0.066, -0.006)
Stormy Daniels (3)	D	Best guess (Δ_{BG})	0.098	0.021	(0.059, 0.143)
		Equal certainty (Δ_{EC})	0.056	0.012	(0.034, 0.082)
		Probabilistic (Δ)	0.109	0.016	(0.079, 0.140)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.042	0.010	(-0.063, -0.025)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.053	0.011	(0.032, 0.073)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	0.011	0.013	(-0.016, 0.037)
Trade deficit (3)	D	Best guess (Δ_{BG})	0.182	0.036	(0.110, 0.251)
		Equal certainty (Δ_{EC})	0.100	0.020	(0.060, 0.139)
		Probabilistic (Δ)	0.131	0.023	(0.086, 0.172)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.081	0.016	(-0.114, -0.050)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.031	0.011	(0.007, 0.051)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.051	0.019	(-0.089, -0.015)
Trump 'grab' (2)	D	Best guess (Δ_{BG})	0.148	0.020	(0.109, 0.187)
		Equal certainty (Δ_{EC})	0.080	0.011	(0.058, 0.101)
		Probabilistic (Δ)	0.134	0.015	(0.103, 0.163)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.069	0.010	(-0.088, -0.050)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.055	0.010	(0.035, 0.075)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.014	0.013	(-0.039, 0.010)
Trump 'grab' (3)	D	Best guess (Δ_{BG})	0.186	0.024	(0.137, 0.235)
		Equal certainty (Δ_{EC})	0.117	0.016	(0.086, 0.148)
		Probabilistic (Δ)	0.169	0.018	(0.132, 0.204)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.069	0.010	(-0.091, -0.050)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.052	0.012	(0.028, 0.074)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.017	0.013	(-0.044, 0.008)
Trump 'grab' (6)	D	Best guess (Δ_{BG})	0.251	0.033	(0.188, 0.315)
		Equal certainty (Δ_{EC})	0.183	0.024	(0.136, 0.231)
		Probabilistic (Δ)	0.227	0.026	(0.176, 0.275)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.068	0.010	(-0.089, -0.049)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.044	0.011	(0.022, 0.066)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.024	0.012	(-0.049, 0.000)
Clinton email (2)	R	Best guess (Δ_{BG})	0.035	0.013	(0.011, 0.060)
		Equal certainty (Δ_{EC})	0.022	0.008	(0.007, 0.037)
		Probabilistic (Δ)	0.080	0.012	(0.059, 0.104)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.013	0.005	(-0.025, -0.004)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.059	0.010	(0.041, 0.078)

Table D.5: Estimates plotted in Figure 6.6 (continued)

Category	Estimator	Estimate	SE	CI	
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	0.045	0.011	(0.024, 0.068)
Clinton email (3)	R	Best guess (Δ_{BG})	0.096	0.016	(0.066, 0.129)
		Equal certainty (Δ_{EC})	0.064	0.012	(0.044, 0.089)
		Probabilistic (Δ)	0.103	0.013	(0.077, 0.131)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.032	0.007	(-0.048, -0.020)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.039	0.010	(0.019, 0.060)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	0.007	0.010	(-0.013, 0.026)
Inflation (3)	R	Best guess (Δ_{BG})	0.189	0.031	(0.132, 0.252)
		Equal certainty (Δ_{EC})	0.100	0.017	(0.070, 0.133)
		Probabilistic (Δ)	0.118	0.018	(0.084, 0.155)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.089	0.015	(-0.120, -0.061)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.018	0.011	(-0.004, 0.038)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.071	0.019	(-0.107, -0.033)
Inflation (4)	R	Best guess (Δ_{BG})	0.170	0.027	(0.117, 0.224)
		Equal certainty (Δ_{EC})	0.088	0.014	(0.061, 0.116)
		Probabilistic (Δ)	0.089	0.015	(0.061, 0.119)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.082	0.013	(-0.109, -0.056)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.001	0.008	(-0.015, 0.018)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.081	0.016	(-0.113, -0.051)
Inflation (5)	R	Best guess (Δ_{BG})	0.215	0.045	(0.119, 0.307)
		Equal certainty (Δ_{EC})	0.108	0.023	(0.062, 0.156)
		Probabilistic (Δ)	0.112	0.027	(0.060, 0.162)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.107	0.023	(-0.151, -0.060)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.004	0.016	(-0.027, 0.036)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.102	0.027	(-0.159, -0.049)
Inflation (6)	R	Best guess (Δ_{BG})	0.161	0.036	(0.092, 0.230)
		Equal certainty (Δ_{EC})	0.083	0.019	(0.046, 0.119)
		Probabilistic (Δ)	0.056	0.020	(0.016, 0.096)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.077	0.017	(-0.113, -0.044)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	-0.027	0.012	(-0.052, -0.002)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.105	0.022	(-0.152, -0.062)
Obama DAPA (6)	R	Best guess (Δ_{BG})	0.152	0.036	(0.084, 0.227)
		Equal certainty (Δ_{EC})	0.073	0.017	(0.039, 0.109)
		Probabilistic (Δ)	0.106	0.020	(0.066, 0.144)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.079	0.019	(-0.120, -0.043)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.033	0.012	(0.010, 0.057)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.046	0.023	(-0.092, -0.004)
Scalise shooting (3)	R	Best guess (Δ_{BG})	0.165	0.029	(0.109, 0.225)
		Equal certainty (Δ_{EC})	0.081	0.015	(0.054, 0.110)
		Probabilistic (Δ)	0.117	0.019	(0.081, 0.156)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.084	0.015	(-0.114, -0.055)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.036	0.012	(0.013, 0.061)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.048	0.019	(-0.082, -0.010)
Stock market (3)	R	Best guess (Δ_{BG})	0.166	0.028	(0.113, 0.223)
		Equal certainty (Δ_{EC})	0.105	0.018	(0.072, 0.141)
		Probabilistic (Δ)	0.134	0.020	(0.095, 0.172)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.061	0.011	(-0.082, -0.041)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.028	0.010	(0.009, 0.049)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.033	0.014	(-0.060, -0.006)
Trump-Russia (2)	R	Best guess (Δ_{BG})	0.155	0.026	(0.105, 0.205)
		Equal certainty (Δ_{EC})	0.091	0.016	(0.061, 0.121)
		Probabilistic (Δ)	0.086	0.019	(0.048, 0.121)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.064	0.011	(-0.085, -0.043)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	-0.005	0.012	(-0.028, 0.019)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.069	0.016	(-0.100, -0.038)
Trump-Russia (6)	R	Best guess (Δ_{BG})	0.066	0.035	(-0.002, 0.136)
		Equal certainty (Δ_{EC})	0.045	0.023	(-0.001, 0.090)
		Probabilistic (Δ)	0.064	0.026	(0.014, 0.113)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.022	0.012	(-0.046, 0.001)

Table D.5: Estimates plotted in Figure 6.6 (continued)

Category	Estimator	Estimate	SE	CI	
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.020	0.011	(-0.003, 0.041)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.002	0.016	(-0.036, 0.028)
Unemployment (2)	R	Best guess (Δ_{BG})	0.069	0.019	(0.034, 0.105)
		Equal certainty (Δ_{EC})	0.034	0.010	(0.016, 0.054)
		Probabilistic (Δ)	0.061	0.013	(0.037, 0.087)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.034	0.010	(-0.054, -0.017)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.027	0.010	(0.008, 0.048)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.008	0.014	(-0.035, 0.021)
Unemployment (3)	R	Best guess (Δ_{BG})	0.112	0.026	(0.059, 0.164)
		Equal certainty (Δ_{EC})	0.076	0.018	(0.040, 0.112)
		Probabilistic (Δ)	0.114	0.019	(0.074, 0.151)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.036	0.009	(-0.054, -0.019)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.038	0.011	(0.018, 0.060)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	0.002	0.014	(-0.026, 0.029)
Unemployment (4)	R	Best guess (Δ_{BG})	0.188	0.025	(0.141, 0.234)
		Equal certainty (Δ_{EC})	0.126	0.017	(0.094, 0.159)
		Probabilistic (Δ)	0.164	0.017	(0.133, 0.196)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.061	0.009	(-0.079, -0.045)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.038	0.009	(0.021, 0.055)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.023	0.013	(-0.047, 0.001)
Unemployment (5)	R	Best guess (Δ_{BG})	0.080	0.040	(0.003, 0.160)
		Equal certainty (Δ_{EC})	0.061	0.031	(0.003, 0.124)
		Probabilistic (Δ)	0.088	0.030	(0.035, 0.149)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.019	0.010	(-0.040, -0.001)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.028	0.015	(-0.002, 0.055)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	0.008	0.019	(-0.030, 0.046)
Unemployment (6)	R	Best guess (Δ_{BG})	0.016	0.029	(-0.038, 0.075)
		Equal certainty (Δ_{EC})	0.012	0.020	(-0.027, 0.051)
		Probabilistic (Δ)	0.046	0.021	(0.007, 0.089)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.005	0.008	(-0.022, 0.011)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.034	0.010	(0.014, 0.054)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	0.030	0.014	(0.000, 0.058)
Wages (2)	R	Best guess (Δ_{BG})	0.167	0.024	(0.122, 0.213)
		Equal certainty (Δ_{EC})	0.056	0.009	(0.040, 0.073)
		Probabilistic (Δ)	0.068	0.011	(0.045, 0.090)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.111	0.016	(-0.142, -0.080)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.012	0.009	(-0.005, 0.029)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.100	0.019	(-0.136, -0.061)
Wages (3)	R	Best guess (Δ_{BG})	0.172	0.026	(0.120, 0.225)
		Equal certainty (Δ_{EC})	0.087	0.014	(0.060, 0.114)
		Probabilistic (Δ)	0.128	0.016	(0.096, 0.158)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.085	0.013	(-0.110, -0.059)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.041	0.011	(0.019, 0.062)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.044	0.017	(-0.077, -0.009)
Wages (4)	R	Best guess (Δ_{BG})	0.254	0.023	(0.209, 0.297)
		Equal certainty (Δ_{EC})	0.125	0.012	(0.102, 0.149)
		Probabilistic (Δ)	0.157	0.013	(0.130, 0.182)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.129	0.012	(-0.153, -0.106)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.032	0.008	(0.016, 0.047)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.097	0.014	(-0.125, -0.068)
Wages (5)	R	Best guess (Δ_{BG})	0.256	0.045	(0.166, 0.340)
		Equal certainty (Δ_{EC})	0.138	0.025	(0.090, 0.188)
		Probabilistic (Δ)	0.164	0.027	(0.111, 0.214)
		Difference ($\Delta_{\mathrm{EC}} - \Delta_{\mathrm{BG}}$)	-0.118	0.022	(-0.159, -0.075)
		Difference ($\Delta - \Delta_{\mathrm{EC}}$)	0.026	0.017	(-0.008, 0.056)
		Certainty bias ($\Delta - \Delta_{\mathrm{BG}}$)	-0.092	0.026	(-0.146, -0.042)

Appendix E

Survey Information

E.1 Certainty Scales

This section provides visual support for the descriptions of the certainty scales in the main text. Respondents were first asked to choose which of two response options is more likely to be true:

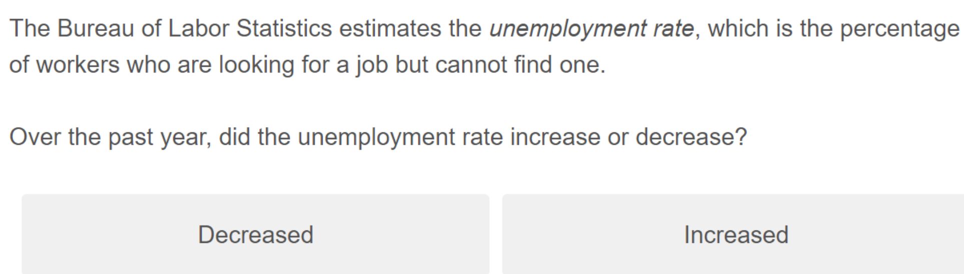

Immediately after the respondent selected their choice, a certainty scale appeared just below the question:

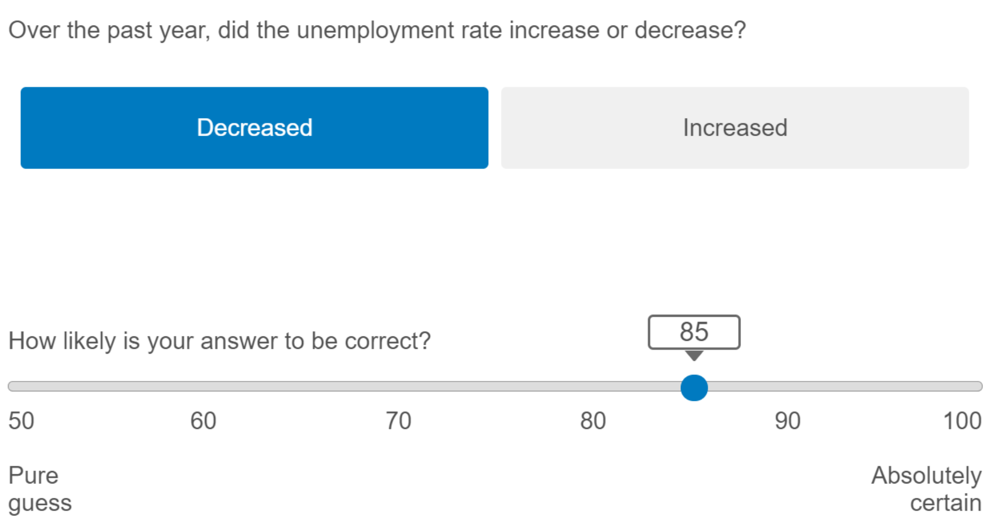

E.2 Revealing Beliefs through Costly Choices

As discussed in Chapter 4, several of the surveys included a measure that reveals respondents' beliefs through a series of costly discrete choices. This section describes how this procedure was implemented in the surveys.

Training

When subjects first reach the point in the survey containing the revealed belief measure, they are told that they will be making a series of choices that could result in them winning a bonus:

> BONUS OPPORTUNITY:
>
> At the end of the survey, one person will be chosen to be eligible to win a **$100 bonus**.
>
> In this part of the survey, you'll answer some questions about what you think gives you the best chance to win the bonus. If you are selected, the computer will use one of your choices to decide if you win.
>
> To show you how it works, let's start with two practice questions.
>
> →

The first practice question asks subjects to choose which side they think a coin will land on: heads or tails. Immediately after making this choice, a second choice appears below it: would you rather win the bonus if the coin lands on the side you chose, or would you prefer an 8 in 10 chance of winning the bonus? After making this choice, an additional line of text appears, explaining why an 8 in 10 chance is the better choice.

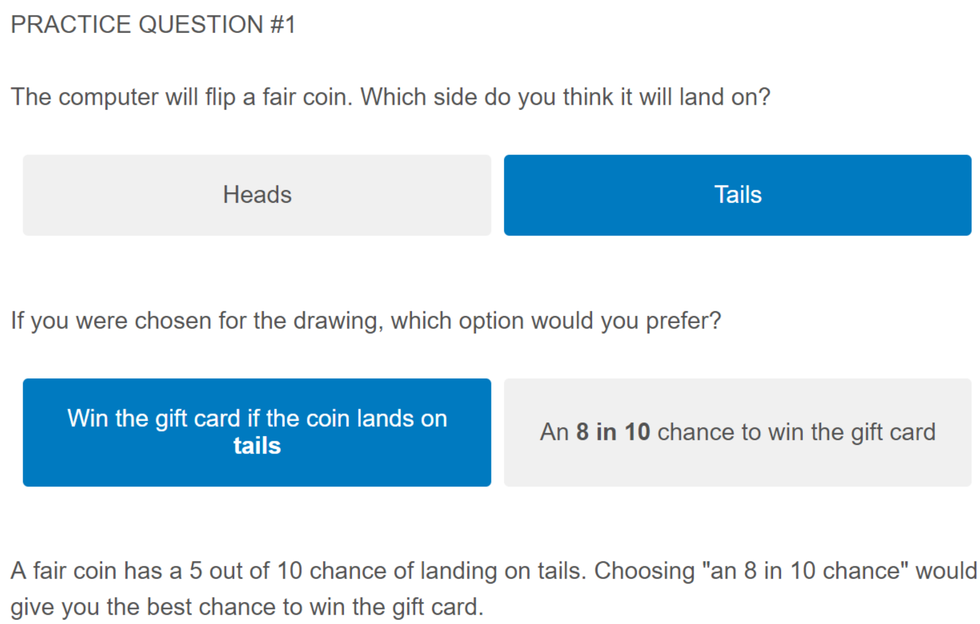

Next, subjects complete this same procedure for a choice between whether a shape pictured is a square or a triangle. This time, the intuition is that they are sure they answered correctly, they should choose payment for a correct answer over an 8 in 10 chance to win.

Having been introduced to the idea that they might sometimes prefer the lottery and sometimes prefer payment for a correct answer, subjects are next introduced to the idea that they may make multiple such choices on a single page.

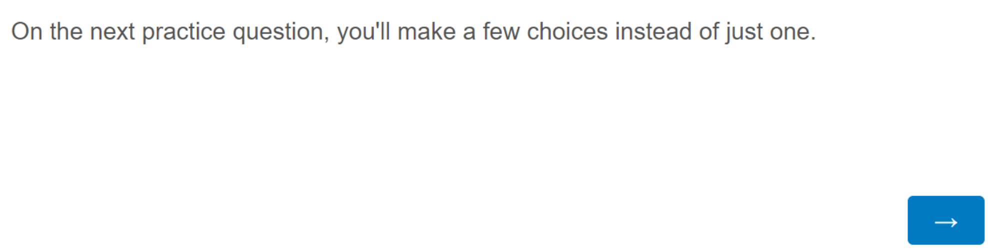

The next screen begins with a question about which a reasonable person might be uncertain: is Earth more or less than 50 million miles from the sun?

Consider this question:

Is Earth more than 50 million miles away from the sun?

Which option do you prefer?

Win if Earth is **less than** 50 million miles from the sun.	Win if Earth is **more than** 50 million miles from the sun.

→

After making this choice, five additional choices appear, in the same format they saw before. Subjects are explicitly told that they are making five separate choices and are encouraged to choose whatever gives them the best chance to win.

PRACTICE QUESTION

Consider this question:

Is Earth more than 50 million miles away from the sun?

Which option do you prefer?

Win if Earth is **less than** 50 million miles from the sun.	Win if Earth is **more than** 50 million miles from the sun.

Below is a list of five choices.

For each choice, which option gives you the best chance to win?

Choice 1	Win if the sun is **less than** 50 million miles away	**6 in 10** chance to win
Choice 2	Win if the sun is **less than** 50 million miles away	**7 in 10** chance to win
Choice 3	Win if the sun is **less than** 50 million miles away	**8 in 10** chance to win
Choice 4	Win if the sun is **less than** 50 million miles away	**9 in 10** chance to win
Choice 5	Win if the sun is **less than** 50 million miles away	**99 in 100** chance to win

Once subjects have tried this procedure for themselves, the next screen displays an explanation of how to fill out the screen.

Please read the following explanation carefully.

Most people aren't completely sure how far away the sun is.

Suppose you thought there was a 75 in 100 chance that the sun is more than 50 million miles away. Here's what you should choose.

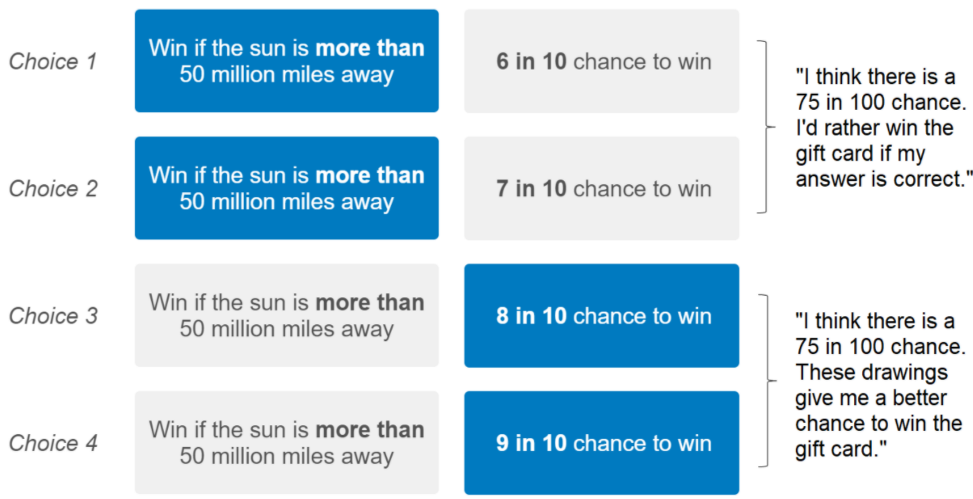

As you move down the list, **you should only "cross over" from the left column to the right column — never from right to left**. After all, if you'd rather have an 8 in 10 chance than get paid if your answer is correct, you should also prefer a 9 in 10 chance.

Example

After completing the training, subjects see another transition slide informing them that it is time to make the real choices that will be used to award the bonus payment.

For each belief, the choices begin by asking respondents which option gives them the better chance to win: being paid if one answer choice is true, or the other one is true.

> Consider this question:
>
> > What is John Roberts' job or political office?
>
> Which would you prefer?
>
> | Win $100 if Roberts is **Secretary of Defense**. | Win $100 if Roberts is **Chief Justice**. |
> |---|---|

After making this initial choice, five additional choices appear, just as in the final training slide.

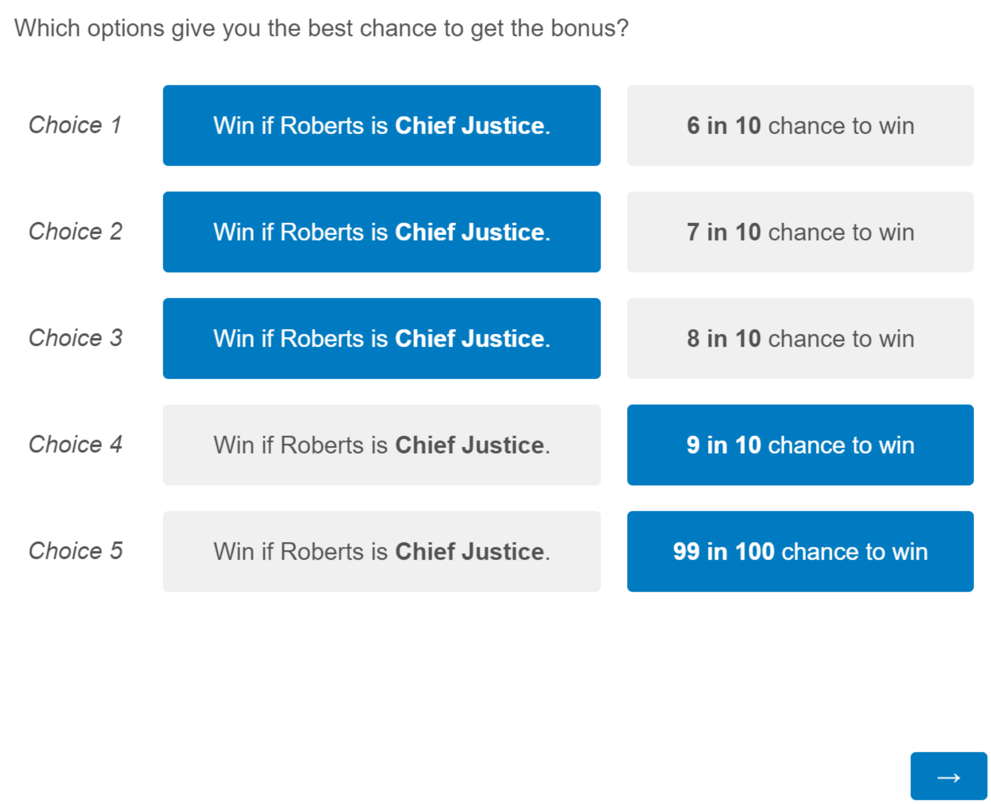

Whenever respondents stray from the instruction never to cross from the right to the left, a warning immediately appears above the question, and detailed instructions for fixing the error appear below the question.

Warning: You just "crossed over" from right to left. This hurts your chance to win the bonus. See below for details.

Which options give you the best chance to get the bonus?

Choice 1	Win if growth was **4% or more**.	**6 in 10** chance to win
Choice 2	Win if growth was **4% or more**.	**7 in 10** chance to win
Choice 3	Win if growth was **4% or more**.	**8 in 10** chance to win
Choice 4	Win if growth was **4% or more**.	**9 in 10** chance to win
Choice 5	Win if growth was **4% or more**.	**99 in 100** chance to win

Detailed warning:

You said you prefer a 6 in 10 chance, but not a 7 in 10 chance.

This will hurt your chances to win the $100 bonus.

To have the best chance to win, you should only cross from the left to the right as you move down the list — never from the right to the left. Please change this before you continue.

E.3 Detecting Information Search

Implementation

To detect whether respondents were looking up the answers to the questions, the four most recently-conducted original surveys — the October 2019 MTurk survey, December 2019 Lucid survey, February 2020 Lucid survey, and March 2020 Lucid survey — contained snippets of Javascript that detected whether the survey remained visible on the respondent's screen. Similar code is pervasive in contemporary web design, e.g. to stop or start a video when a respondent looks away from or at a web page. The method was implemented by embedding the following Javascript:

```
Qualtrics.SurveyEngine.addOnload(function(){

  function recordLeave() {
    var prefix = "itemName_";
    if (document.hidden) {
      var theStart = new Date();
      var embedName = prefix + "everLeft";
      Qualtrics.SurveyEngine.setEmbeddedData(embedName, "1");
      var embedName = prefix + "timeLeft";
      Qualtrics.SurveyEngine.setEmbeddedData(embedName, theStart.getTime()/1000);
    }
  };
  document.addEventListener('visibilitychange', recordLeave, false);

  function recordArrive() {
    var prefix = "itemName_";
      if (document.hidden){}else {
        var theEnd = new Date();
        var embedName = prefix + "timeReturn";
        Qualtrics.SurveyEngine.setEmbeddedData(embedName, theEnd.getTime()/1000);
      }
    };
  document.addEventListener('visibilitychange', recordArrive, false);

  $('NextButton').onclick = function (event) {
    document.removeEventListener('visibilitychange', recordLeave);
    document.removeEventListener('visibilitychange', recordArrive);
    Qualtrics.SurveyEngine.navClick(event, 'NextButton')
  }

});
```

The only difference across questions was the variable `prefix`, which added a question code to the variables suffixed `_everLeft`, `_timeLeft`, and `_timeReturn`.

This code adds three variables to the dataset: whether the respondent ever lost page visibility, the time at which the respondent lost page visibility, and the time at which the respondent returned to the page. In the results, all losses of page visibility lasting 5 seconds or more were dropped. In surveys 5 or 6, a loss of page visibility on either measure (stated or revealed) resulted in all observations for these questions being dropped. This ensures that differences between the stated and revealed measure are not affected by differences in information search behavior across the two measures.

An Estimator for Undetected Information Search

This strategy for detecting information search is affected by measurement error. Not all respondents who look up the information are flagged as losing page visibility, and not all respondents who are flagged as losing page visibility actually looked up the information. For the results presented in the main text, the former aspect of measurement error is the greatest threat: if a substantial amount of information search goes undetected, the method cannot be viewed as a successful strategy for eliminating such responses from the data.

To evaluate the extent to which undetected cases of information search are likely to be present in the data, three pieces of information were used: (1) the percentage of respondents flagged as losing page visibility on each question, (2) an estimate of how likely respondents are to be flagged when they actually look up the information, and (3) an estimate of how likely respondents are to be flagged when they do not look up the information. The first quantity varies at the question level, while the latter two are static features of the measurement technology. Consequently, this strategy allows one to estimate how the prevalence of information search varies across questions while also accounting for measurement error.

To see how this strategy works, begin by the law of total probability to rewrite the probability of being flagged, $P(\text{flag})$:

$$P(\text{flag}) = P(\text{look})P(\text{flag}|\text{look}) + (1 - P(\text{look}))P(\text{flag}|\neg\text{look}), \tag{E.1}$$

then re-arrange terms to isolate the probability of looking up the information, $P(\text{look})$:

$$P(\text{look}) = \frac{P(\text{flag}) - P(\text{flag}|\neg\text{look})}{P(\text{flag}|\text{look}) - P(\text{flag}|\neg\text{look})}. \tag{E.2}$$

This expresses the probability of looking up the answer in terms of the three quantities listed above: (1) $P(\text{flag})$, (2) $P(\text{flag}|\text{look})$, and (3) $P(\text{flag}|\neg\text{look})$.

The ultimate quantity of interest is the probability of looking up the information *given that one was not flagged as looking it up*. Mathematically, this quantity is $P(\text{look}|\neg\text{flag})$. To express this quantity in terms of the above, first use Bayes' rule to rewrite it as

$$P(\text{look}|\neg\text{flag}) = \frac{P(\neg\text{flag}|\text{look})P(\text{look})}{P(\neg\text{flag})}, \tag{E.3}$$

then rewrite terms from (E.3), and substitute (E.2) into the second term in the denominator, to yield

$$P(\text{look}|\neg\text{flag}) = \frac{\Big(1 - P(\text{flag}|\text{look})\Big)\Big(\frac{P(\text{flag}) - P(\text{flag}|\neg\text{look})}{P(\text{flag}|\text{look}) - P(\text{flag}|\neg\text{look})}\Big)}{1 - P(\text{flag})}. \tag{E.4}$$

Here, the left-hand side is the quantity of interest, $P(\text{look}|\neg\text{flag})$. The right-hand side expresses this quantity in terms of the three observable quantities listed above. Thus, the empirical version of the right-hand side is an estimator for the prevalence of undetected information search.

Estimating Undetected Information Search

The estimator for undetected information search requires the percentage of flagged responses to be combined with two other quantities: the probability of being flagged when one looked up the information, $P(\text{flag}|\text{look})$, and the probability of being flagged when one did not look up the information, $P(\text{flag}|\neg\text{look})$.

Table E.1: Estimates of undetected information search.

Category	Survey	Question	P(flag)	P(look)	P(look|¬flag)
Controversies	6	Article II	0.025	0.001	0.000
		Obama DAPA	0.025	0.001	0.000
		Trump-Russia	0.036	0.016	0.005
		Trump 'grab'	0.027	0.003	0.001
Economic	5	Deficit	0.031	0.008	0.003
		GDP $< 4\%$	0.031	0.008	0.003
		Health ins.	0.017	-0.012	-0.004
		Inflation	0.031	0.008	0.003
		Unemployment	0.024	-0.002	-0.001
		Wages	0.031	0.008	0.003
	6	Deficit	0.039	0.020	0.006
		GDP $< 4\%$	0.047	0.033	0.010
		Inflation	0.032	0.010	0.003
		Unemployment	0.026	0.002	0.001
Political awareness	5	Chief justice	0.050	0.037	0.012
		Fed chair	0.081	0.083	0.027
		Foreign aid	0.029	0.006	0.002
		House control	0.029	0.006	0.002
		Reconciliation	0.034	0.013	0.004
	6	Chief justice	0.066	0.061	0.020
		Fed chair	0.073	0.071	0.023
		Foreign aid	0.026	0.002	0.001
		House control	0.040	0.022	0.007

To estimate the first quantity, respondents to surveys 5 and 6 were offered a financial incentive to report a quantity that can quickly and easily be found using an Internet search, the population of Nicaragua. About 85 percent of respondents reported a correct answer. About 70 percent of respondents who answered correctly were flagged, including 89 percent of respondents who indicated that they had used the same device that they are using to take the survey to look up the answer; about two-thirds of respondents with unflagged correct answers indicated that they had used a different device. Based on these results, the estimates assume that $P(\text{flag}|\text{look}) = 0.7$.

To estimate the second quantity, respondents in an unrelated survey on Lucid were asked to state their best guess as to the population of Nicaragua, with strong discouragement to look it up. Just 3 percent of respondents reported a correct answer. Of those who reported an incorrect answer, about 2.5 percent were flagged as having looked it up. Based on these results, the estimates assume that $P(\text{flag}|\neg\text{look}) = 0.025$.

With these estimates in hand, the percentage of respondents who were flagged on each question can be plugged into (E.4) to produce an estimate of the prevalence of undetected information search. Table E.1 presents these estimates for all of the questions in the February 2020 Lucid and March 2020 MTurk surveys.

On the controversy and economic questions, the estimated prevalence of undetected information search never exceeds 1 percent, and rarely exceeds 0.5 percent. Undetected information search is a bit more common on the questions about the job or political office held by John Roberts (Chief Justice of the Supreme Court) and Jerome Powell (Chairman of the Federal Reserve), reaching 1.2 to 2.7 percent of the data.

E.4 Survey Text

This section provides information about each original survey analyzed in this volume, including the platform, date, sample size, question text. In most of the volume, surveys are referred to using the vendor and the date. In Chapter 6, they are numbered for display in the figures. These numbers appear in parentheses after the name of the survey.

September 2017 Lucid Survey

Platform: Lucid.

Date: September 2017.

Number of subjects: 1729.

Anti-cheating measures: Pledge.

Description: The only subjective scale analyzed in Chapter 4.

Full text of questions analyzed:

1. What is the name of the first ten amendments to the U.S. Constitution?

   [The Bill of Rights, The Preamble, The Articles of Confederation]

2. Which country does Recep Erdogan lead?

   [Turkey, Egypt, Ukraine]

3. What position in government does John Roberts currently hold?

   [Chief Justice of the Supreme Court, U.S. Senator, Secretary of Defense]

4. What leadership position does Angela Merkel hold?

   [Chancellor of Germany, Prime Minister of the United Kingdom, Secretary-General of the United Nations]

5. In the U.S. Senate, what procedure allows changes to the budget to be passed by a simple majority? Reconciliation, Filibuster, Prerogative]

6. Who is the Senate Minority Leader?

   [Chuck Schumer, Nancy Pelosi, Mark Warner]

7. How long is a U.S. Senator's term in office?

   [6 years, 4 years, 2 years]

8. How many votes in Congress are needed to override a presidential veto?

   [2/3 majority, 3/4 majority, Simple majority]

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were asked, "How certain are you that your answer is correct?" and provided a five-item scale labelled:

- Don't know (pure guess).
- Not too certain.

- Somewhat certain (good guess).
- Very certain.
- Absolutely certain (definitely correct).

May 2018 MTurk Survey (Survey 1)

Platform: Amazon Mechanical Turk.

Date: May 2018.

Number of subjects: 519.

Anti-cheating measures: Pledge, catch question.

Description: Included in Chapter 6.

Full text of questions analyzed:

1. Some people say that former president Barack Obama was not born in the United States. They have asked Obama to release his long-form birth certificate.

   Did Obama release his long-form birth certificate?

   [Yes, No]

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were asked, "How likely is your answer to be correct?" and provided a quasi-continuous 50 to 100 scale.

June 2019 MTurk Survey (Survey 2)

Platform: Amazon Mechanical Turk.

Date: June 2019 (wave 1), June 2020 (wave 2).

Number of subjects: 1,244.

Anti-cheating measures: Pledge, catch question.

Description: Included in Chapters 5 and 6.

Full text of questions analyzed:

1. The Bureau of Labor Statistics estimates the *unemployment rate*, which is the percentage of workers who are looking for a job but cannot find one.

   Between April 2018 and April 2019, did the unemployment rate increase or decrease?

   [Unemployment went up, Unemployment went down]

2. The amount of money people earn at their jobs is often measured using the *median real wage*. "Median" means the person right in the middle and "real" means adjusted for inflation.

   Between Spring 2018 and Spring 2019, did the median real wage in the U.S. go up or down?

   [Wages went up, Wages went down]

3. Over the past year, has the percentage of Americans who have health insurance gone up or down?

[Higher percentage has insurance now, Lower percentage has insurance now]

4. When the U.S. buys more products from a country than it sells to the country, the U.S. has a *trade deficit* with that country.

   Is the following statement true or false?

   In 2018, the U.S. trade deficit with China reached a new record high.

   [True, False]

5. Is the following statement true or false?

   Before becoming president, Donald Trump was tape recorded saying that he kisses women and grabs them between the legs without their consent.

   [True, False]

6. Is the following statement true or false?

   While she was Secretary of State, Hillary Clinton used a private email server to send and receive classified information.

   [True, False]

7. Robert Mueller was in charge of the special counsel investigation into possible Russian interference in the 2016 election.

   Is the following statement true or false?

   Robert Mueller's final report stated that there is "undeniable proof" that President Trump personally conspired with Russian agents to influence the 2016 election.

   [True, False]

8. Is the following statement true or false?

   Barack Obama has never released his birth certificate.

   [True, False]

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were asked, "How many chances in 100 does your answer have to be correct?" and presented with a quasi-continuous 50 to 100 scale.

July 2019 Lucid Survey (Survey 3)

Platform: Lucid.

Date: July 2019.

Number of subjects: 1,217 assigned to form A, 1,223 assigned to form B.

Anti-cheating measures: Pledge.

Description: Included in Chapters 3, 4, and 6.

Full text of questions analyzed: See the Appendix A, which is the appendix to Chapter 3.

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were asked, "How sure are you about that?" and presented with a six-item scale with the options 50, 60, 70, 80, 90, 100 and labels at 50 and 100.

October 2019 MTurk Survey

Platform: Amazon Mechanical Turk.

Date: October 2019 (wave 1), November 2019 (wave 2).

Number of subjects: 1,244 (wave 1), 631 (wave 2).

Anti-cheating measures: Pledge, cheating detection Javascript (see section E.3).

Description: Included in Chapter 4.

Full text of questions analyzed:

1. Which party currently has the most members in the U.S. Senate?

   [Democratic, Republican]

2. On which of the following does the U.S. federal government currently spend the **least**?

   [Medicare, Foreign aid]

3. Which country does Angela Merkel lead?

   [Austria, Germany]

4. Which country does Viktor Orban lead?

   [Hungary, Turkey]

Format of certainty scale: See the "application" section of Chapter 4

December 2019 Lucid Survey (Survey 4)

Platform: Lucid.

Date: December 2019.

Number of subjects: 4,125 respondents answered a random two of the five total questions. About 1,650 respondents answered each question.

Anti-cheating measures: Pledge, cheating detection Javascript.

Description: Included in Chapter 6.

Full text of questions analyzed:

1. The Bureau of Labor Statistics estimates the *unemployment rate*, which is the percentage of workers who are looking for a job but cannot find one.

   Over the past year, did the unemployment rate increase or decrease?

   [Unemployment increased, Unemployment decreased]

2. The amount of money people earn at their jobs is often measured using the *median real wage*. "Median" means the person right in the middle and "real" means adjusted for inflation.

   Over the past year, did the median real wage in the U.S. go up or down?

   [Wages went up, Wages went down]

3. The rate of *inflation* measures how quickly prices are rising.

   Over the past year, has inflation been higher or lower than the historical average (since 1945)?

[Higher than average, Lower than average]

4. Over the past year, has the percentage of Americans who have health insurance gone up or down?

   [Higher percentage has insurance now, Lower percentage has insurance now]

5. As you may know, the U.S. House of Representatives recently concluded an impeachment investigation into President Trump.

   Is this statement true or false?

   During the impeachment investigation, congressional Democrats did not allow Republicans to select any of the witnesses.

   [True, False]

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were randomly assigned to be asked, "How likely is your answer to be correct?" or "How sure are you about that?" and provided a quasi-continuous 50 to 100 scale with labels at 50, 75, and 100. No systematic differences between the scale wordings was found.

February 2020 Lucid Survey (Survey 5)

Platform: Lucid.

Date: February 2020.

Number of subjects: 532.

Anti-cheating measures: Pledge, cheating detection Javascript.

Description: Included in Chapters 4 and 6.

Full text of questions analyzed:

1. The Bureau of Labor Statistics estimates the *unemployment rate*, which is the percentage of workers who are looking for a job but cannot find one.

   Over the past year, did the unemployment rate increase or decrease?

   [Decreased, Increased]

2. The amount of money people earn at their jobs is often measured using the *median real wage*. "Median" means the person right in the middle and "real" means adjusted for inflation.

   Over the past year, did the median real wage in the U.S. increase or decrease?

   [Decreased, Increased]

3. The rate of *inflation* measures how quickly prices are rising. Since World War II, the average inflation rate has been about 4 percent.

   Over the past year, has inflation been higher or lower than the historical average?

   [Above average, Below average]

4. Over the past year, has the percentage of Americans who have health insurance gone up or down?

   [Higher percentage has insurance now, Lower percentage has insurance now]

5. The size of the U.S. economy is usually measured using gross domestic product (GDP). *Economic growth* is the annual rate of change in GDP.

   Over the past year, what was the rate of economic growth in the United States?

   [Less than 4%, 4% or more]

6. Most years, the U.S. national government spends more than it collects in taxes. In these years, the government has an annual *budget deficit*.

   Compared with the 2017 fiscal year, was 2019's budget deficit higher or lower?

   [Higher, Lower]

7. What job or political office does John Roberts hold?

   [Secretary of Defense, Chief Justice of the Supreme Court]

8. What job or political office does Jerome Powell hold?

   [Treasury Secretary, Chairman of the Federal Reserve]

9. On which of the following does the U.S. federal government currently spend the least?

   [Medicare, Foreign Aid]

10. Which party currently has the most members in the U.S. House of Representatives?

    [Democrats, Republicans]

11. What Senate procedure allows budget changes with a simple majority vote?

    [Filibuster, Reconciliation]

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were randomly assigned to be asked, "How likely is your answer to be correct?" or "How sure are you about that?" and provided a quasi-continuous 50 to 100 scale with labels at 50 and 100. No systematic differences between the scales were found.

March 2020 MTurk Survey (Survey 6)

Platform: Amazon Mechanical Turk.

Date: March 2020 (wave 1), August 2020 (wave 2).

Number of subjects: 939 (wave 1), 420 (wave 2).

Anti-cheating measures: Pledge, cheating detection Javascript.

Description: Included in Chapters 4, 5, and 6.

Full text of questions analyzed:

1. The Bureau of Labor Statistics estimates the *unemployment rate*, which is the percentage of workers who are looking for a job but cannot find one.

   Over the past year, did the unemployment rate increase or decrease?

   [Decreased, Increased]

2. The rate of *inflation* measures how quickly prices are rising. Since World War II, the average inflation rate has been about 4 percent.

Over the past year, has inflation been higher or lower than the historical average?

[Above average, Below average]

3. The size of the U.S. economy is usually measured using gross domestic product (GDP). *Economic growth* is the annual rate of change in GDP.

   Over the past year, what was the rate of economic growth in the United States?

   [Less than 4%, 4% or more]

4. Most years, the U.S. national government spends more than it collects in taxes. In these years, the government has an annual *budget deficit*.

   Compared with the 2017 fiscal year, was 2019's budget deficit higher or lower?

   [Higher, Lower]

5. Is the following statement true or false?

   Before becoming president, Donald Trump was tape recorded saying that he kisses women and grabs them between the legs without their consent.

   [True, False]

6. Robert Mueller was in charge of the special counsel investigation into possible Russian interference in the 2016 election.

   Is this statement true or false? *Robert Mueller's report stated that President Trump personally conspired with Russia to influence the 2016 election.*

   [True, False]

7. Article II of the U.S. Constitution describes the president's powers.

   Is this statement true or false? *President Trump has said that Article II gives him the power to do whatever he wants.*

   [True, False]

8. In 2014, former President Barack Obama issued an order that would stop most deportations of unauthorized immigrants who have U.S. citizen children.

   Is this statement true or false? *About a year earlier, Obama said that he would be ignoring the law if he issued such an order.*

   [True, False]

9. What job or political office does John Roberts hold?

   [Secretary of Defense, Chief Justice of the Supreme Court]

10. What job or political office does Jerome Powell hold?

    [Treasury Secretary, Chairman of the Federal Reserve]

11. Which party currently has the most members in the U.S. House of Representatives?

    [Democrats, Republicans]

12. On which of the following does the U.S. federal government currently spend the least?

    [Medicare, Foreign Aid]

Format of certainty scale: The certainty scale appeared immediately after each respondent chose their answer. Respondents were randomly assigned to be asked, "How likely is your answer to be correct?" or "How sure are you about that?" and provided a quasi-continuous 50 to 100 scale with labels at 50 and 100. No systematic differences between the scales were found.