What Do We Actually Know? Correcting for Differential Item Functioning with a Survey Experiment

By

Jason Means

Submitted to

Central European University

Department of Political Science

In partial fulfilment of the requirements for the degree of Master of Arts

Supervisor: Dr. Gabor Tóka

Abstract

Despite being one of the best-documented features of contemporary U.S. politics, scholars are still in disagreement on how to address the extent to which voters are uninformed about politics. A large portion of the electorate is comprised of low-information voters—a fact that is dissonant to most characterizations of democratic citizenship. This study seeks to discern whether the electorate, specifically African Americans, are actually uninformed about politics or whether their general level of political knowledge is misrepresented due to knowledge batteries biased in favor of majority (White) Americans. Before testing for differential item functioning, a methodological problem causing respondents to answer survey questions in vastly different ways from one another, a theoretical argument against disenfranchising low-information voters is presented, and is followed by an overview of the literature on political knowledge, and its measurement. This study's findings remain inconclusive. Though its treatment has an effect, the extent of that effect is uncertain. Multiple regression analysis on factors influencing respondents' level of political knowledge shows a statistically significant relationship between race and knowledge, however, only among control-group respondents. Further research is needed to assess whether item bias may be reduced through altering the construction of surveys measuring political knowledge.

Acknowledgements

First and foremost, I am extremely grateful to the Department of Political Science, Dr. Gabor Toka and Dr. Levente Littvay for all of their support and guidance throughout the process of writing this thesis. Additionally, this research would not have been possible without a generous grant from the Central European University Foundation of Budapest—thank you.

Karl, you're an absolute gem. Our many conversations about home, family, and race constantly reaffirmed my dedication to finishing this beast of a document.

Gretchen, Regina, Karen, and Cady: Thank you for keeping me sane throughout this entire process. Your friendship means the world to me.

Liam, this thesis is dedicated to you. It's long, so you don't have to read it all. But I truly wish you all the best that life has to offer; keep learning, keep growing, and always remember that you can do anything you set your mind to. Come visit, eh?

Table of Contents

Abstracti
Acknowledgementsii
List of Figuresv
List of Tables vii
Introduction1
Chapter 1: Epistocracy2
1.1 The Competence Principle
1.2 The Principle of Equal Consideration of Interests5
1.3 Disenfranchisement and Authority9
1.4 Voter Competence Tests10
1.5 Concluding Remarks12
Chapter 2: Political Knowledge 14
2.1 The Condorcet Theorem16
2.2 The Heuristic Category17
2.3 The Heuristic-Skeptic Category18
2.4 The Measurement of Political Knowledge19
2.4.1 Exacerbating Gaps?21
2.5 Differential Item Functioning and the Blackbox Transpose22
2.5.1 DIF Along Racial Lines
Chapter 3: An Experiment Testing for DIF and Substantive Gaps

3.1 Introduction	26
3.2 Hypotheses	26
3.3. Experimental Design	27
3.4 Sample Size Estimates	29
3.5 Pilot Study	
3.6 Results	
3.6.1 Results (Regression Analysis)	47
3.7. Discussion	50
3.8. Limitations	51
3.9. Conclusion	52
References	
Appendix 1: MTurk Survey	

List of Figures

Figure 1: Base Knowledge (Control)
Figure 2: Base Knowledge (Treatment)
Figure 3: Name Recognition (Control)
Figure 4: Name Recognition (Treatment)
Figure 5: Base Knowledge (Control)
Figure 6: Base Knowledge (Treatment)
Figure 7: Name Recognition (Control)
Figure 8: Name Recognition (Treatment)
Figure 9: Incorrect (0) vs. "Don't Know" (99) Responses, Court Item (Control)
Figure 10: Incorrect (0) vs. "Don't Know" (99) Responses, Secretary Item (Control) 39
Figure 11: Incorrect (0) vs. "Don't Know" Responses, Court Item (Treatment) 40
Figure 12: Incorrect (0) vs. "Don't Know" Responses, Secretary Item (Treatment) 40
Figure 13: White Respondents' Placement of Democratic Party on 11-Point Scale 41
Figure 14: Black Respondents' Placement of Democratic Party on 11-Point Scale 42
Figure 15: Blackbox Transpose Corrected Estimates (White Respondents)

Figure 16: Blackbox Transpose Corrected Estimates (Black Respondents) 46

List of Tables

Table 1: Respondents' Placements on 11-Point L-R Scale (in Percent)
Table 2: Multiple Regression Analysis of Factors Influencing Control-Group Respondents' Level
of Political Knowledge
Table 3: Multiple Regression Analysis of Factors Influencing Treatment-Group Respondents'
Level of Political Knowledge

Introduction

Political knowledge and its true utility in elections has been a widely contested topic within UScentric literature. Voter ignorance is considered unsurprising (Delli Carpini and Keeter 1993; Lupia 1994; Bartels 1996) and there are considerable inequalities among the level of knowledge held by different groups of citizens (Verba, Schlozman and Brady; Mondak and Anderson 2004; Abrajano 2014; Pérez 2015). This is not a new phenomenon: Galston (2001) finds that today's college graduates do not know more than high school graduates in the 1950s, which isn't that much. The lack of a well-informed citizenry is troubling for democracy, considering that a commonly held democratic belief is that citizens should be competent enough to engage in their political system. Subsequently, Platonic-esque arguments supporting the disenfranchisement of low-information voters have gained more prominence. Should we strip away the voting rights of the poorly informed? How much do citizens actually not know? Recent literature on the subject suggests that knowledge gaps between particular groups are not as large as scholars have previously claimed, specifically with regard to knowledge gaps across racial groups. This paper begins with a theoretical critique of one of the many individuals supporting voter disenfranchisement, Jason Brennan (2011). Afterwards, the literature on political knowledge, knowledge "gaps", and the measurement of political knowledge will be explored in order to lay the foundation for an experiment testing the gap in political knowledge between White and Black Americans, drawing upon an approach inspired by both Abrajano (2014) and Pérez (2015).

Chapter 1: Epistocracy

Jason Brennan (2011) asserts that contemporary democracies would benefit from moving to restricted suffrage from universal suffrage. This move, he argues, would be a "moral improvement" (701). Restricted suffrage, or epistocracy, whereby those deemed incompetent would have their voting rights stripped away, is a non-ideal theory—a theory of the second-best (Brennan, 712). Nonetheless, Brennan claims that though restricted suffrage is unjust, it is better than universal suffrage in that it does not violate the competence principle. Though the case for voter disenfranchisement seems convincing, there are systematic flaws in Brennan's argumentation. In this chapter, the basic premises of Brennan's call-to-arms for restricted suffrage will be evaluated; this chapter concludes, with the help of Christiano (2004) and Rawls (1958), that voter disenfranchisement is not a moral improvement. Perhaps even more striking, this chapter also concludes that though many individuals and institutions promote and attempt the construction and implementation of tests measuring political knowledge, they are inherently problematic and difficult to employ in an unbiased manner without knowing the interests of individuals and groups. Therein lies the existential reason for this chapter-the theoretical arguments presented in this chapter paired with Chapter 3 present what I believe is a strong theoretical and practical case against the use of political knowledge tests to disenfranchise the electorate. Though seemingly sensible—low information voters (depending on the political preferences of the reader) have certainly had negative effects on electoral outcomes-disenfranchisement is antithetical to democracy.

This chapter is structured as follows: The first section will provide the foundation for my argument, where I will detail Brennan's support of what he calls the competence principle. The second section

will discuss the principles of equal consideration of interests and the fair equality of opportunity, and the difference between judgements and interests; key assumptions within the competence principle neglect to address this distinction. The third section will discuss authority—how the violation of the equal consideration of interests negates the obligation of those who would find themselves disenfranchised by the competence principle to adhere to the authority of the state. The fourth and final section will argue against Brennan's desired competence tests as a means to decide who can and cannot maintain the right to vote.

1.1 The Competence Principle

"Many of my fellow citizens are incompetent, ignorant, irrational, and morally unreasonable about politics. Despite that, they hold political power over me" (Brennan, 700).

In true Platonic fashion, Brennan argues on behalf of suffrage restricted to citizens who can demonstrate political competence. The foundation of his argument stems from the logic that as an innocent person, he should not be subject to his incompetent fellow citizens exercising political power over himself or other competent individuals. As a matter of justice, he believes that these citizens should be excluded from holding the right to vote. This argument can be broken down to the following premises:

- 1) There are many incompetent and morally unreasonable people.
- 2) Universal suffrage allows these people to vote.
- 3) It is wrong for an individual to be forced to submit to the decisions of these voters.
- Restricting voting to individuals who are not incompetent or morally unreasonable would be a moral improvement.

5) Hence, restricted suffrage is morally superior to universal suffrage.

With regard to formal logic, premises 1 and 2 are both valid and true; that citizens use little information and may be irrational in casting their vote is one of the best-documented findings in political science.¹ Premises 3-5 may be logically valid, but are they true? Premise 3 is based on Brennan's competence principle, whereby:

"It is unjust to deprive citizens of life, liberty, or property, or to alter their life prospects significantly, by force and threats of force as a result of decisions made by an incompetent or morally unreasonably deliberative body, or as a result of decisions made in an incompetent and morally unreasonable way." (704)

The application of this principle to support Brennan's third premise is perhaps partially true. It assumes that voters within a system of universal suffrage are not always equal, which would certainly violate the democratic principle of equality. However, is my 1/nth share of decision-making power not the same as Brennan's? Electoral power and political influence are two separate things: For example, Brennan claims that Obama has more *electoral* decision-making power than him. While Obama may have much more political influence, that does not mean that his actual vote is any more than 1/nth of the share of electoral decision-making power. This argument is also misleading in that it does not support Brennan's fourth or fifth premises: Obama and other influential political actors would maintain their superior influence regardless of whether restricted

¹ For a rather deterministic example of this see Todorov et al. 2005, whose findings suggest that inferences of competence based solely on facial appearance predicted the outcomes of U.S. congressional elections better than chance.

suffrage were adopted. Therefore, the focus should shift from those "superior" individuals who would pass Brennan's competence test to those who are deemed to be incompetent and morally unreasonable. Premises 4 and 5 will be discussed in the following sections.

1.2 The Principle of Equal Consideration of Interests

"The poorest he that is in England has a life to live as the greatest he"

(Rainsborough, in Christiano 2004a, 43).

In the previous section I introduced the competence principle, and argued that Brennan's third premise falsely assumes that his share of decision-making power is any less than mine. Politically resourceful individuals will maintain their superior influence whether or not restricted suffrage is adopted, therefore I argued that Brennan's focus should stay on the incompetent and morally unreasonable electorate. This section will discuss Brennan's fourth premise, that restricted suffrage is morally superior to universal suffrage. This argument is rooted in the debate on whether incompetent and morally unreasonable voters are making false judgements (voting decisions that are misaligned with their interests), which would determine whether their voting rights should be revoked.

Christiano (2004a) points out that judgements may be incorrect or correct, though interests are neither; "they are simply attributes of a person" (43). Therefore, it is difficult to determine whether a judgement is incorrect or correct without knowing the interests of the individual(s) making the judgement. I could argue that the interests of low-income workers should be low inflation due to its subsequent effects on economic growth and job creation. I could also argue that these low-income workers choosing to vote in favor of costly welfare programs that reduced economic

growth and job creation would be making an incorrect judgement. However, utilizing these arguments to propose that low information voters making "incorrect" judgements should be stripped of their voting rights carries a misleading assumption (of which Brennan is also guilty of making):

 Those with a higher level of political competence are more likely to know the interests of citizens with lower levels of competence.

This assumption is misleading in that higher-competence individuals may make *more informed* decisions in line with political realities, however, they are just as likely as individuals with less knowledge to make judgements based on their intrinsic interests. The aforementioned workers might prioritize healthcare over economic growth. If so, a vote for welfare programs would be a correct judgement—not grounds for disenfranchisement.

Christiano aptly notes that the "[e]qual consideration of interests means that advancing the interests of one person is as important as advancing the interests of any other person" (2004a, 43). Just as the above quote by Thomas Rainsborough notes, those individuals who are worst off in society have as much as right to a good life as those who are better off. The democratic principle of the equal consideration of interests is truly egalitarian in that each individual has the same ability to influence the results of an election as their fellow citizens, whether they are "competent" or not. As noted above, it would be difficult to assess the true interests of the electorate, therefore we cannot definitively say whether or not their judgements are incorrect. If Brennan were advocating autocracy, which he is not, his call for restricted suffrage may be sensible—however, his arguments support an allegedly morally superior approach to decision-making within the realm of

democracy, and as such, his violation of the principle of equal consideration of interests is decidedly undemocratic and morally reprehensible.

It should be noted, however, that Brennan does not claim restricted suffrage to be the *most* morally superior choice for democracy; instead, it is morally superior to universal suffrage. For this assertion to hold, his advocacy for partial disenfranchisement must stand up to the Rawlsian view of fair equality of opportunity (1958), a principle that finds partial loss of freedom for the sake of the greater good to be unjust. Injustice, Rawls argues, is only acceptable when it is necessary to avoid greater injustice (4). Therefore, for Brennan's claim to hold we must see the adoption of restricted suffrage as the prevention of the greater injustices perpetuated by universal suffrage. But aside from the injustice of incompetent voters holding "political" power over us (which I have already questioned above—electoral power is different from political influence; all citizens only have 1/nth share of the vote), which injustices stem from universal suffrage? Despite further inquiry into Brennan's text I could not find a single one, though he discusses an analogy of juries and competence at length to support his disenfranchisement argument, which the following paragraph will address.

Brennan rightly notes that individuals of bad epistemic and moral character should not serve on juries. As such, jury decisions that are known to be incompetent (even if made by individuals who usually show a certain degree of competence and morally reasonable behavior) should not be enforced, and the affected defendants have no duty to submit (704). This is already corrected for in jury interviews, which is perhaps why Brennan thought this to be an apt analogy for restricted suffrage. However, it is fundamentally inappropriate to compare juries to political voters. Firstly, a jury is not meant to make decisions based on their individual considerations of what is right—

rather, they are told what the law is and are subsequently asked to reach a collective judgement based on the facts of a case. Secondly, a jury vote has a *direct* effect on the individual on trial, whereas voting has *indirect* effects; considering that a political representative is the individual who makes binding and semi-binding decisions on behalf of the electorate (excluding referenda, of course), we cannot say that incompetent voting exerts political power over the competent, but rather, that elected officials do. Brennan counters this argument by asserting that if citizens demand competence from their political representatives then, *prima facie*, incompetent decision-making or decisions made by incompetent individuals is unjust (707). My counter-counter argument is the following: Citizens do not have the inherent right to demand competence from their political representatives; rather, citizens have a right to hold their democratically elected leaders accountable. Incompetent decision-making is a negative externality of democratic governance and is therefore not inherently unjust; should their interests dictate, citizens may correct for incompetent decision-making at the ballot box.

Brennan's fourth premise presupposes that voters with a higher level of competence are likely to know the interests of less-competent voters. As interests and judgements are not the same—the former is a precondition for the determining the evaluation of the latter—and considering the difficulty discerning another individual's actual interests, this presupposition was deemed to be untrue. Furthermore, Brennan's competence principle violates not only the equal consideration of interests principle, but he fails to provide compelling evidence to justify its violation of the fair equality of opportunity principle. To be disenfranchised is to lose the equal share of power that comes with fully enfranchised voting, rendering the lives of disenfranchised voters worse off than those who meet a particular knowledge threshold. We therefore cannot instantiate the disenfranchisement of low-information voters into any societal prescription without violating

major precepts of democratic theory. Though valid, Brennan's fourth premise is untrue: Restricted suffrage is not morally superior to universal suffrage.

1.3 Disenfranchisement and Authority

The previous sections have discussed my objections to Brennan's arguments regarding the role of the citizen in society regarding the right to a good life. This section will take a theoretical, albeit brief, look at what role low-competence citizens would play in a society that has adopted restricted suffrage with regard to the obligation to respect the authority of decisions made by the enfranchised electorate.

Brennan's argument is instrumentalist in that it approaches the question of what democracy ought to be as a matter of what produces the "best" outcomes. Christiano (2004b) notes that instrumentalists argue that all citizens have the same duties to comply with the authority of the state, but that this depends on showing that each citizen has "duties that are the same as everyone else's" (2). In a partially enfranchised society, all citizens would be required to comply with the decisions enacted by voters and democratically elected officials, but how does this conception of democratic authority fit the instrumental justification of this mandatory compliance? Before providing an answer, let us look at the premises outlined by Christiano that are associated with instrumentalism and authority:

- 1) Democratic decisions ought to be made in a way that will produce the best outcomes.
- 2) These outcomes are just because the best outcomes promote the common good.
- 3) Each citizen has the same political duties as everyone else.
- 4) Therefore, each citizen has the obligation to comply with democratic decisions.

The third premise is where we begin to question the nature of Brennan's instrumental argument regarding the "inherent" obligation to comply with authority. If low-competence voters had their voting rights revoked they would no longer have the same duties as their enfranchised peers, thus eliminating Christiano's fourth premise. Consequentially, Brennan's disenfranchised electorate would no longer have the obligation to comply with decisions that they no longer had a role in. The authority possessed by a decision-maker is entirely piecemeal on the justification of the fourth premise; under partial enfranchisement, a large percentage of citizens would no longer be morally required to comply with democratic decision-making—an obvious problem for the functionality of a society.

1.4 Voter Competence Tests

Now that I have argued against premises 3-5 of Brennan's argument for partial enfranchisement, it is paramount to discuss the problematic nature of testing for competence. As I have previously mentioned, Brennan has ignored a substantial body of literature on voting behavior and testing for political knowledge. Though Brennan presents a normative argument, he also claims that his societal prescriptions would be ideal for democracies *today*, therefore it is perfectly acceptable to utilize contemporary social science literature to show how unrealistic the creation of a non-biased test would be. First, I will argue against low-competence voters as being inherently detrimental to democracy. Second, I will present an outline of why testing for competence is unjust.

Though the literature on low-information voters is the subject of great debate, there are findings that are worth pointing out in defense of a fully enfranchised electorate. Lupia's (1994) findings suggest that low-information voters can emulate highly informed behavior, and that political campaigning creates incentives for elite actors to provide and widely spread cues that allow for this emulation. Lee's (2012) research identifies low-information voters as being better informed than previous research had realized, noting that those who know little about federal politics may know more about local-level politics. These texts are but two of many that detail the ability of lowinformation voters to either vote in a more informed way through the use of heuristics—also known as information short cuts—and show that these voters might actually be more informed than they seem upon first glance. Brennan asserts that we cannot easily test for this kind of competence (715), but it has nevertheless been done. Judgements vs. interests aside, Brennan's incompetent electorate has the potential to be more competent than he assumes.

Though the aforementioned tests of competence have been done, they are by no means a flawless tool to be used to disenfranchise voters.² One primary reason why voter competence tests should not be created for the purpose of disenfranchisement is that tests for political knowledge may further disadvantage already-disadvantaged members of society. Brennan acknowledges this consequence of partial enfranchisement, though he argues that voter tests would only "*reflect* and *result from* an underlying injustice, but that does not mean that [they] would [themselves] be unjust" (720). By this logic any competence test is not responsible for disadvantaged people not meeting a required knowledge threshold. I strongly oppose this view. Any such institutional arrangement would reflect, result from, and *perpetuate* underlying injustice in that it would further exacerbate the division between the well- and worst-off members of society. Consider the low-income worker I mentioned earlier: If the socioeconomic status of this worker were low enough to the point that they were unable to provide for their family, you might imagine that they would be

² The measurement of political knowledge will be discussed in greater detail in Chapter 2.

unable to afford an exceptionally good education for their children. It is perfectly reasonable to assume that the lack of a good education will prevent these children from passing a competence test, subsequently preventing them from participating in elections that could affect whether social welfare or mobility programs are enacted to improve their socioeconomic status. This disenfranchisement will then be passed on to their children, their children's children, and so on, all things being equal.

1.5 Concluding Remarks

In this chapter I have argued that the use of the competence principle as a justification for restricted suffrage violates the principles of equal consideration of interests and the fair equality of opportunity. I have also argued that if incompetent voters were disenfranchised, they would have no moral obligation to comply with the authority of democratic decision-making by enfranchised voters. Lastly, through examples drawn from contemporary social science research I showed that incompetent voters may not be as incompetent as one might think, and that the creation of a voter competence test is impractical in that it would most likely further disadvantage those members of society who are already disadvantaged. If one were to create a test measuring political knowledge, as this paper sets out to do, it would be important to attempt to assess individual interests in addition to knowledge in order to determine whether these individuals are making the "right" decisions when voting. Brennan's epistocracy may be compelling at first glance, but a deeper look at his premises shows that restricted suffrage is morally inferior to universal suffrage. Universal suffrage is not perfect—there are very real criticisms to be made; however, it is currently the best option for democratic society.

Though seemingly tangential to the following chapters, this theoretical argument presents a cautionary tale of the arbitrary nature of disenfranchising voters based on low levels of political knowledge. This thesis provides a direct example of the difficulties associated with measuring political knowledge—if we cannot definitively identify the extent to which voters are informed about politics—whether due to biased measurement, simply asking the wrong questions, or the inability to grasp the underlying interests of the individuals being surveyed—we cannot begin to consider rescinding their democratic rights.

The following chapter explores what political knowledge actually is, the arguments in support of and against whether citizens should actually be politically informed, how political knowledge is measured, and discusses knowledge "gaps" across racial boundaries.

Chapter 2: Political Knowledge

What is political knowledge? Simply put, it refers to factual knowledge about politics that can be recalled from the conscious memory of citizens (Delli Carpini and Keeter 1996). To provide a practical example, it is objectively true that the President of the United States can legally serve two terms in office. The same objective truth cannot be attributed to the idea that taxes should be higher or lower—the former is a characteristic of political knowledge whereas the latter is a political attitude.

Despite this clear definition, there are a number of similar-yet-different terms used by public opinion scholars that obfuscate knowledge's conceptual formation and measurement:

"[V]ariables purporting to measure 'political awareness,' 'political expertise,' 'political sophistication,' 'cognitive sophistication,' 'political information,' 'political involvement,' 'media exposure,' and 'political interest' appear regularly in the public opinion literature and are used (along with education) more or less interchangeably to explain the same family of dependent variables." (Zaller 1992, 126)

Though this grouping of dependent variables is certainly comprised of correlates, the interchangeable use of terms has the potential to create conceptual disorder when referring to factual knowledge about politics. For example, if a scholar were to use a survey item asking respondents to identify an incumbent or candidate, I could code a correct response as being an indicator of *political awareness*, *political sophistication*, or both. What does this mean for the discipline's state of the art? As Kulinski and Quirk (2001) assert, "conceptual elements need explicit definition" (306). Knowledge, therefore, requires specific conceptualization so as to not render findings useless and/or ambiguous.

This paper refers to political knowledge with the intention of conveying Deli Carpini and Keeter (1996)'s above definition without being mired in other terminology from the conceptual family; it is not measuring whether individuals are politically sophisticated, aware, or informed, but rather the extent to which surveys can or cannot accurately measure the factual information about politics that individuals and groups know—in essence, this paper is operating off of a rather narrow definition. There are certainly consequences to being misinformed or poorly informed about politics, however, this chapter does not make any empirical claims regarding the extent to which an individual's "good" or "bad" political decision-making is a result of their level of knowledge; it simply presents the divergent schools of thought surrounding the importance (or lack thereof) of political knowledge as a foundation towards further understanding its measurement.

That being said, what is knowledge for? What is its argued utility? Democratic theorists generally agree that political knowledge is an essential component of democratic citizenship (Keeter 2011, 2)—citizens are obligated to maintain a certain degree of knowledge through which they can contribute to the democratic process. Hence, the utility of knowledge is that citizens can effectively participate in the democratic process. However ideal, this is not the reality, particularly in a U.S.-context. Dissonant to most characterizations of democratic citizenship, the subject literature consistently finds that there are large amounts of variance in the levels of political knowledge across population groups (Keeter 2011, 3) and scholars generally agree that overall, the amount of knowledge held by citizens is usually too low to effectively participate in politics (Ibid.). Yet, despite being "one of the best-documented features of contemporary politics" (Bartels 1996, 94), scholars are still in disagreement on how to address the electorate's general lack of information about politics. This has led to diverging "schools" of thought regarding political knowledge and

voting³—1) The Condorcet Theorem, whereby individual voting error cancels out in the aggregate (Condorcet 1972); 2) The Heuristic Category, whereby low-information voters are posited to efficiently use cues, shortcuts, and/or heuristics in their given environments to make better-informed choices (Berelson, Lazarfeld and Mcphee 1954; McKelvey and Ordeshook 1985; Neuman 1986; Page and Shapiro 1992); and 3) The Heuristic-Skeptic Category, whereby the extent to which heuristic-use is efficient is questioned—this view posits that while voters do indeed use shortcuts, the extent to which they are used efficiently depends on a voter's level of political knowledge or cognitive ability (Lau and Redlawsk 2001; Todorov et al. 2005).

2.1 The Condorcet Theorem

The Condorcet Jury Theorem, first established in 1785, sets a theoretical basis for democracy (Ladha 1992, p. 617). In essence, it establishes that the majority of a group with limited access to information about two alternate decisions has a higher probability to choose the "better" decision than any one member of the group (Ladha 1992, p. 619). This theorem, however, has been identified to have two major assumptions that limit its viability: 1) It assumes that individuals vote independently, and 2) that voters have a shared goal (Ladha 1992, p. 617). While there are some theoretical merits to this theorem when these two assumptions are relaxed (Miller 1982; Ladha 1992), in praxis, it has generally been declared to be unrealistic—a prerequisite for individual error canceling out at the aggregate level is that this error must be random; however, several findings indicate that nonrandom error exists within individual-level voting (Althaus 1998; Graber 1984;

³ These three "schools" have been categorized by this paper; I am not operating off of previous literature that does so, but am rather grouping similar scholarly arguments together in order to display trends within the literature on political knowledge, voting, and the democratic obligation of citizens to be informed about politics.

Lau and Redlawsk 2001). Althaus (1998), for example, shows that the introduction of information effects (nonrandom error) biases collective preferences to represent the views of an elite minority. Graber (2004) suggests that journalists do not see themselves as educators and therefore do not report on everything that one would deem newsworthy; additionally, citizens selectively choose which sources they pay attention to and which ones they ignore, subsequently introducing more nonrandom error (550-551). Thus, we begin to see that while theoretically interesting, the Condorcet Theorem is perhaps not as applicable to the practical conundrum of low-information voters as once thought.

2.2 The Heuristic Category

If the Condorcet Theorem doesn't help to clearly parse the problem of low-information voting, then perhaps we can look to the Heuristic Category. This category posits that voters can (and do) compensate for their lack of knowledge through the use of heuristics as a mechanism to make more informed decisions. McKelvey and Ordershook's (1985) findings, for example, suggest that "two-candidate democratic systems can work…even if voters have no direct knowledge of candidate positions and candidates have no direct knowledge of voter preferences" (508) due to the use of endorsements as a heuristic. As mentioned in the previous chapter, Lupia (1994), in a study of voting behavior during a referendum on insurance reform in California, found that low-information voters could emulate highly informed behavior, and that campaigns created incentives for elite actors to provide and widely spread cues that allowed for such emulation (though these findings are unable to be generalized).

This category presents an ideal response to the problem of low-information voting: If individuals with less knowledge have the capacity to make better informed choices by either emulating their

fellow informed citizens or using endorsements and other heuristics then, in essence, political knowledge does not matter. This category's solution lies in the elimination of the problem altogether. However, the following category provides compelling evidence to refute this assertion, or at the very least it suggests that we should take this category's findings with a grain of salt.

2.3 The Heuristic-Skeptic Category

Though the previous group of scholars presents evidence suggesting that individuals do use heuristics to make better-informed decisions, more recent literature within experimental political science and cognitive psychology indicates that this heuristic use is not always efficient. Todorov et al. (2005), for example, present findings suggesting that voters use a heuristic as arbitrary as facial appearance (within a 1-second exposure, nonetheless) to determine candidate competence. Though facial cues are indeed important for successful social interaction, Todorov et al. point out that campaigns spend millions of dollars on disseminating information about parties and candidates in order to convince voters to choose them (2005, 1623). If voters rely on facial appearance as a strong indicator of candidate competence, then one could assume that campaigns are essentially wasting their money. This also implies that voters are behaving irrationally—an individual may vote for a candidate whose policy platform aligns with the voters' interests, however, Todorov et al.'s study suggests that this is unlikely to be the case.

Lau and Redlawsk (2001) use a series of experiments that show individuals with high levels of political knowledge to be more likely to effectively use heuristics, however, low-information individuals had a *lower* probability of effectively using heuristics.

These results, and the findings from the other literature categories within studies on political knowledge and its effects on voting behavior, suggest that political knowledge is a complex topic of which we still have much to explore. Nevertheless, in many situations it appears that we aren't the voters we think we are, sometimes subconsciously using cues, shortcuts, or heuristics to make the process of voting more manageable. However, one topic of particular concern when discussing political knowledge is its measurement. Have scholars truly found a way to effectively measure the level of knowledge an individual has?

2.4 The Measurement of Political Knowledge

After defining a concept, it is important to accurately measure the concept's indicators. Now that we have defined political knowledge (factual knowledge about politics that can be recalled from the conscious memory of citizens) we must decide which indicators adequately measure the level of knowledge individuals hold. This section will provide an overview of how political knowledge batteries—series of knowledge questions—are created, along with the costs and benefits of particular measurement choices (i.e. open vs. closed questioning, coding, and survey design that leads to cheating).

Given our definition of political knowledge, it seems sensible to create a series of questions, also known as *items*, that ask individuals factual questions about national institutions, processes, and political actors. However, this leads us to the following questions: Which questions should be asked? How many? Which kind? Should there be a "don't know" option? In using educational testing for guidance, scholars generally agree that tests of knowledge "should include many questions, cover most or all important subdomains of knowledge, and have varying levels of difficulty" (Keeter 2011, 3).

Regarding the format of knowledge items, Delli Carpini and Keeter (1993) mention that multiple choice items have many proponents, but suggest that they are awkward in telephone surveys and should therefore be avoided unless the item response count is limited to three options per question (1183). Items may also be true/false, or open-ended (also known as "free response"), however true/false items encourage guessing (Delli Carpini and Keeter 1993, 1183) and open-ended questions are subject to coding error (DeBell 2013).

The quintessential case of open-ended questions leading to coding error can be found when looking at one of the most prominent data centers in the world, the American National Election Studies (ANES). This is documented in DeBell (2013): Before 2008, none of the optimal practices for open-ended coding were followed, leading to the 2004 ANES coder-instruction manual claiming that "The Prime Minister of the United Kingdom" was an incorrect response to the question "What job or political office does Tony Blair now hold" (394). Decades of ANES' coding instructions remained vague or didn't exist at all until 2008, causing a data quality crisis (DeBell 394). Though ANES has worked towards remedying this problem, we can infer from this crisis that it may be methodologically safer to rely on closed questioning in order to reduce coding error.

Cheating can also harm the validity of responses obtained from political knowledge batteries, especially when these batteries are administered online (Clifford and Jerit 2016). When survey respondents were directly asked whether they cheated in political knowledge surveys as much as 22 percent of the sample admitted that they had (Jensen and Thomson 2014), showing that cheating is a problem that requires further investigation: Clifford and Jerit (2016)'s findings suggest that cheating is reduced through the inclusion of a "commitment item" that directly asks respondents whether they are willing to answer knowledge items without aid (873). Other techniques that

Clifford and Jerit suggest are less effective are directly asking respondents to not cheat (867), and including a timer (872). Mechanical Turk (MTurk)—the platform that this paper used to conduct research—has an inherent financial incentive for respondents to complete surveys as fast as possible, preventing most instances of cheating (Clifford and Jerit 2016, 873).

2.4.1 Exacerbating Gaps?

Despite these potential shortcomings in the construction of knowledge items, what is perhaps more difficult to assess and solve than coding errors and the tendency to cheat is the prevalence of differential item functioning (DIF), a methodological problem stemming from survey respondents answering the same question in "vastly different ways" (King et al. 2004, 191; Brady 1985). DIF has the potential to artificially inflate gaps in the levels of political knowledge various groups hold. For example, Mondak and Anderson's (2004) findings suggest that men are more likely to guess on survey items whereas women are more likely to not answer or simply select don't know (usually over-inflating men's knowledge by an average of 50%).

Regarding race, Abrajano's (2015) research suggests that the reported gap in knowledge between White, Latino, and Black Americans suffers from DIF. Through differing group identity construction and shared experiences, Abrajano argues that minority groups possess a "perceptual 'bias' that leads them to interpret [survey] questions in a way that differs from the way the majority of Americans interpret those same questions" (44). To correct for DIF occurring in alreadyadministered surveys, she employs a scaling procedure known as the Blackbox Transpose (Aldrich and McKelvey 1977; Poole 1998; Abrajano 2014) on ANES 2008 and 2012 raw data. After this procedure, DIF was found to diminish significantly, indicating that the political knowledge gap between racial groups is not as stark as previous research has indicated. Pérez's (2015) research was inspired by similar instances of respondents employing different interpretations of survey questions, however, instead of applying a statistical correction, he updated survey questions to better match Latino respondents through the inclusion of Latino incumbents such as Marco Rubio, for example. His findings suggest that creating Latino-themed items lessens the political knowledge gap between White and Latino Americans without reducing the level of knowledge reportedly held by White respondents. Therefore, the following study presents a first-attempt at checking the extent to which Pérez's findings hold; namely, an A/B survey experiment was conducted that utilized a battery of political knowledge questions, including a treatment with updated ANES questions that are more relevant to Black Americans. If Pérez's findings do generalize to other racial minorities, we will have further evidence to support the claim that surveys utilizing inclusive items are more likely to better represent the actual levels of political knowledge of all Americans, not just the majority. But first, the following sections will dive deeper in exploring the technical aspects of DIF and why it may occur among African American respondents.

2.5 Differential Item Functioning and the Blackbox Transpose

The scaling method Abrajano (2015) uses to correct for differential item functioning, the Blackbox Transpose, is a generalized scaling procedure meant to account for the existence of perceptual biases in ordered issue scales (Abrajano 48-49; Poole 1998). This procedure addresses DIF through the assumption that stimuli, in Abrajano's case two major presidential candidates and parties in the 2008 election, possess "true" positions on a left-right scale, Y_j , yet the way respondents perceive these positions are somehow systematically distorted (Abrajano 2015, 48-49). The

Transpose model assumes that *J* stimuli—prominent political actors—are to be located on a scale by *N* respondents. The "true" or latent position of these political figures is θ_j ($1 \le j \le J$). Each survey respondent, *i*, is assumed to also possess a latent perception of the *j*-th political actor ($1, \le i \le N$). These latent perceptions are defined as the "true" positions with errors that adhere to Gauss-Markov assumptions (homoscedasticity, 0 mean, and independence [Aldrich and McKelvey 1977, 113])⁴; however, researchers are not able to capture true positions, and are instead limited to capturing observed perceptions on the left-right scale of political figure *j* by respondent *i*—*Y*_{*ij*}. The Blackbox Transpose, then, accounts for these perceptual biases (DIF) through the inclusion of stretch parameters α_1 and β_1 . These parameters distort respondent *i*'s placement of the *j*'th political figure so that:

$$\theta_i + \epsilon_{ij} = \theta_{ij} = \alpha_1 + \beta_1 Y_{ij}$$

Based on these assumptions, the Blackbox Transpose estimates stretch parameters α_1 and β_1 and the positions of the political figures⁵ θ_j by reducing the sum of squared residuals of perceived positions regressed on Y_j for all respondents and political figures⁶:

 $\Sigma_{ij} = \Sigma_{\alpha_i} + \beta_i Y_{ij} - \theta_j, \forall_{ij} \forall_{ij}$

CEU eTD Collection

⁴ If the rigid assumption (no perceptual bias) regarding individual candidates holds, then the error term may be Gauss-Markov. However, if it is Gauss-Markov this does not necessarily mean that the assumptions about the nature of perceptual bias hold.

⁵ In Chapter 3 we will see that this includes the positions of political parties, as well.

⁶ This model is from Abrajano (2015), inspired by Poole (1998).

Substantively speaking, the above equations suggest that there is only random variation across respondents in how they understand the substantive meaning of the issues that the scales refer to (i.e. Democrats, Republicans, Elizabeth Warren, and Marco Rubio, in the case of the following experiment). The model parameters (α_1 and β_1) only vary across respondents, but not within respondents and across candidates. Subsequently, these parameters are not capturing the tendency to misperceive individual candidates or to see specific candidates in more extreme or centrist than others regardless of scale and true position. Rather, each individual has their own interpretation and/or understanding of the issue scale. That is, respondents are assumed to have a different understanding of the endpoints in terms of where their mean is on the underlying issue dimension relative to where the respondent's or the candidate's position is – hence the alpha parameter – and also in terms of how far the endpoints are from each other, relative to how far the respondent's or the candidate's positions are from each other – hence the beta parameter.

Essentially, the above model is applied to issue scales in an attempt to reduce DIF, or "interpersonal comparability", through accounting for preferences and affective orientations that bias respondent evaluations of the political world (Armstrong et al. 2014, 40).

2.5.1 DIF Along Racial Lines

The previous section presents a technical introduction to DIF; this section will present a sociological one. While we can display a statistical correction of respondents' answers on a spatial plane, this does not tell us the root cause(s) of the original methodological error.

Pérez (2009) finds that English and Spanish survey items are not created equally—Spanish items are usually direct translations of their English-language counterparts, not accounting for semantic

differences. In many Latin American countries, for example, the term "Liberal" refers to Rightwing parties and candidates. If a Spanish-speaking survey respondent were to position the Democratic party on the Left-Right spectrum as being closer to the Right, we would mark their response as incorrect though the question itself fails to take semantic differences into account (Abrajano 2015, 46). The African American community, similarly, has the potential to misinterpret political knowledge items, though this is based on cultural differences as opposed to general linguistic heritage. Dawson (1994) suggests that African Americans are socialized through particular informal institutions to perceive their fate as being "linked" to their fellow African Americans, thus influencing their political behavior in a way that supports their racial group subsequently, African American political behavior-how they inform themselves about politics, and how they utilize this political information, is collectively (instead of individually) based. If Dawson's theory of linked fate is correct, we can assume that African American respondents will be more knowledgeable about candidates and policies that are relevant to the community as a whole (i.e. candidates who may be perceived as acting on behalf of racial-group interests [African American representatives] or policies that directly affect the community).

Chapter 3: An Experiment Testing for DIF and Substantive Gaps

3.1 Introduction

Chapter 1 provided a theoretical argument against the notion that we should partially disenfranchise the electorate based on political knowledge testing, especially considering that tests for political knowledge have the tendency to be at least partially methodologically flawed. Chapter 2 presented an introduction to political knowledge, and provided the foundation for my experiment in this chapter. Now, based on the aforementioned literature covering knowledge measurement, the inflation of gaps in political knowledge (particularly across racial groups), and differential item functioning, this chapter presents a survey experiment that was conducted with the intention of checking whether the combination of Abrajano and Pérez's findings hold. Does the inclusion of African American-specific items reduce substantive gaps in political knowledge across racial lines? Do these new items reduce the amount of differential item functioning found in issue scales? This experiment is a first-attempt at discerning whether this is the case.

3.2 Hypotheses

As mentioned above, this experiment seeks to discern whether African American-relevant survey questions will reduce the reported racial gap in political knowledge. If it does, this will also likely reduce the level of differential item functioning (DIF) found within the survey. Based on the previous literature, particularly Dawson (1994), Abrajano (2015), and Pérez (2015), my hypotheses are as follows:

H₁: The inclusion of African American-specific items will reduce the substantive gap in political knowledge between African American and White respondents.

H₂: The inclusion of African American-specific items will show a visible reduction of differential item functioning when plotting average item responses against DIF-corrected estimates.

H₃: The inclusion of African American-specific items will not diminish the average reported knowledge of White respondents.

H₃ is perhaps the most important of the three hypotheses; though this experiment attempts to provide a better measure of political knowledge, this should not be done at the expense of other racial groups. If H₃ proves false, it is not the intention of this paper to tout this approach as being an effective way to correct for bias. As Pérez aptly notes:

"The goal is not to replace knowledge items that are 'stacked against' one group answering correctly for items that are 'stacked in favor' of another. Rather, it is to design items that adhere to a common formulation of political knowledge *and* perform equally across groups. If this standard is met, then members of each group will have equal odds of answering a knowledge question correctly, provided they have the same knowledge level." (2015, 938)

3.3. Experimental Design

This experiment utilizes A/B-style testing whereby two surveys are created and administered to a randomized control and treatment group of both White and African Americans. The control group was administered a survey containing a battery of standard political knowledge questions utilized by the American National Election Study (ANES—see Appendix 1). The treatment group received a survey containing the same battery of political knowledge questions with the inclusion of two

questions relevant to African Americans; the control group received an updated version of the standard ANES office-recognition questions "What office is held by Stephen Breyer," and "What office is held by Rick Perry" while the treatment group received versions with African American politicians in similar positions—namely, "What office is held by Clarence Thomas" and "What office is held by Ben Carson."

In line with the political science literature on designating "anchors" to measure DIF (Aldrich and McKelvey 1977; Pzerworksi and Teune 1966-67; Thissen, Steinberg and Wainer 1993; King et al. 2004; King and Wand 2007), both control and treatment groups will first receive the aforementioned battery of political knowledge questions designed to assess base levels of political knowledge. Afterwards, respondents will be asked to position political candidates on a left-right ideology scale. This will enable the raw results to be analyzed for DIF in that it checks whether respondents with the same base levels of political knowledge position candidates differently. If there is significant variance, this is a sign that DIF is indeed present, meaning that the initial base knowledge questions may have biased the results.

Though ANES uses open-ended questioning, this experiment alters questions to be closed and contain a "Don't Know" (DK) option. As mentioned in Chapter 2, closed questioning provides less coding problems (Krosnick et al. 2008) and a "Don't Know" option increases the ability of the researcher to check whether a reported gap in political knowledge is artificially inflated (Mondak and Anderson 2004). As also mentioned in Chapter 2, this experiment did not utilize any specific language to prevent cheating as the use of Mechanical Turk (see below) inherently prevents cheating due to the financial incentive to complete surveys as efficiently as possible.
This experiment's sample was drawn from Mechanical Turk (MTurk), a digital platform through Amazon that allows social scientists to conduct experimental and survey-based research online. Though scholars have questioned the external validity of inferences made from MTurk samples, Berinsky et al. (2012) have shown that respondents recruited via MTurk are usually more representative of the US population than in-person convenience samples. Though less representative than internet-based panels or national probability samples (Huff and Tingley 2015), MTurk will suffice for initial research on racially-inclusive measures of political knowledge in surveys.

3.4 Sample Size Estimates

As funding was limited for this project, correctly estimating the appropriate sample size was key.⁷ In testing various measures of DIF on CCES data, Armstrong et al. (2014) found that Blackbox Transpose results for a sample size of 100 respondents obtained a sum of all errors (between the estimated and true stimuli positions within their sample) of 0.128—a 0.115 difference between the sum of all errors obtained from a sample size of 1,500 respondents. Subsequently, to reduce the overall error within this study's results, I estimated that a sample size of as close to 1,500 is necessary.

Previous research has shown that respondents are most likely to participate in MTurk studies when the incentive is larger than the minimum. Berinsky et al. (2012) were able to obtain more than 200

⁷ Estimating this experiments sample size reduced the risk of not obtaining a large enough sample of African Americans—if Black respondents were too few, as the reader will see in my pilot study, it would have been impossible to test for differential item functioning.

subjects per day when paying as little as \$0.25 per survey (however, this was on a survey that lasted 2-4 minutes) (353). As my own experiment is likely to last at least 10-15 minutes, I deemed it necessary to pay my respondents a minimum of \$0.50 (an effective hourly rate of \$6.00), which is still significantly lower than the traditional costs of undergraduate and nonstudent campus samples, and temporary agency subjects (Berinsky et al. 2012, 353).⁸

The last thing considered when approaching the subject of sample size was the desired attributes of respondents. This experiment specifically requires White and African American respondents, so I deemed it necessary to administer a pre-survey asking for a variety of demographic or opinion-based information which included self-reported racial identity. Those that identified as White or Black received the longer randomized control or treatment surveys, and the other participants finished after taking this initial pre-survey (made possible through the use of Qualtrics). This increased my sample size that that the recruitment of African American respondents was inferred to be at a slower rate than White Americans (6%, or 180 of Huff and Tingley's [2015] sample of 2,706 respondents were Black Americans). I paid respondents who only took the pre-survey a minimum of \$0.1.

In considering the above information, I determined that, at the *very least*, 833 respondents would need to be recruited through MTurk in order to obtain a sample of at least 50 African Americans. Roughly 649 of these respondents were likely to be White Americans (78% of Huff and Tingley's [2015] sample were White). The remaining 134 were likely to be of a different racial group.

⁸ This was perhaps the largest miscalculation on my end—the average respondent time for my survey was about 5 minutes, though the sample was not obtained as quickly as Berinsky et al. due to the experiment's pre-survey, something that is further explored in section 3.9.

However, it is of great importance that I obtained at least 50 Black Americans within my sample in order to provide any statistically sound measure of DIF.

3.5 Pilot Study

A preliminary sample size of n = 300 was obtained. Of this 300, 125 men (41%) and 175 women (59%) were surveyed. The average age of respondents was 39 with a median completed education at the undergraduate (BA or equivalent) level. Less than 10% of respondents identified themselves to be African American (n = 28), therefore any results from this pilot should be interpreted with a grain of salt. Nevertheless, the data present intriguing findings.

Figures 1 and 2 show the results from the designated anchors designed to assess the base knowledge of respondents. Paired-sample t tests show no statistically significant difference between each base knowledge item; in both control and treatment groups a plurality of respondents answered at least three questions correctly.

The remaining two items were designed to test whether racially-inclusive questions better measure the level of political knowledge of minority groups. Paired-sample t tests conducted on each item (between control and treatment groups) indicate an effect on the first name recognition question (regarding Supreme Court Justices; t = 2.09, df = 144, p = 0.03), whereas the second name recognition question is only statistically significant at the p < 0.10 level (Cabinet-level officials; t= 1.77, df = 144, p = 0.07). The statistically significant result of the first name recognition question does suggest that making a question racially-inclusive, at the very least regarding Supreme Court Justices, may better capture the actual level of knowledge of respondents. However, this is

Figure 1: Base Knowledge (Control)



Figure 2: Base Knowledge (Treatment)



Race

insufficient evidence to completely reject the null hypothesis of no difference. For example, looking at the large number of '0' responses in figures 3 and 4 (meaning neither name recognition questions were answered correctly by respondents) we cannot definitively say whether individuals generally do not recognize the listed officials or if this is due to sample size. In contradiction to this study's hypothesis, the results of the treatment actually show substantively negative effects on the reported level of knowledge of individuals in both races (though a small percentage of White respondents answered both treatment questions correctly).

From these preliminary results, we can see that the Supreme Court Justice item showed statistical significance which suggests an effect, however, a larger sample size is needed as well as more advanced analysis, including an item-by-item test for differential item functioning (per the Blackbox Transpose). The pilot, therefore, was inconclusive.

Figure 3: Name Recognition (Control)



Figure 4: Name Recognition (Treatment)



Race

3.6 Results

A sample size of n = 1062 was obtained. Data from respondents who did not fit the research criteria (i.e. those who were neither White nor Black) were removed. Additionally, the survey was updated to randomly send a portion of White respondents to the end of the survey (thus eliminating them from the analysis) in order to obtain an oversample of African American respondents before the n = 1000 threshold was met. After this process, the sample was reduced to n = 658. Of these 658, 39% were male and 61% were female. The average age of respondents was 39 with a median completed education at the undergraduate (BA or equivalent) level. Unlike the pilot study, this experiment obtained an adequate racial distribution—81 respondents were African American and the remaining 577 were white (12% and 88%, respectively). Though this is a rather large difference, it closely mirrors the racial distribution in the United States—in 2015 the U.S. Census Bureau estimated that 13.3% of the population was African American (U.S. Census Bureau).

Figures 5 and 6 show results from the designated anchors designed to assess the base knowledge of respondents in both control and treatment groups. Paired-sample t tests show no statistically significant difference between each base knowledge item and a plurality of respondents answered at least 2 knowledge items correctly—this was to be expected; the base knowledge items were the same in both control and treatment groups, so it would be troubling if there were significant variation between the two.

Figure 5: Base Knowledge (Control)



Race

Figure 6: Base Knowledge (Treatment)



Race

The remaining two items were designed to test whether racially-inclusive questions better measure the level of political knowledge of minority groups, in this case, African Americans. Paired-sample t tests conducted on each item (between control and treatment groups) indicate an effect on the first name recognition question regarding Supreme Court Justices; t = 4.63, df = 326, p =0.000005245) and also the second (Cabinet-level officials; t = 3.1033, df = 328, p = 0.00208). The statistical significance of both items paired with the substantively significant increase in correct item response across racial groups suggests that the inclusion of racially-specific items more accurately captures an individual's level of political knowledge. Despite this substantive increase, there are still a large number of individuals, both White and African American, that answered both questions incorrectly (Figures 7 and 8). However, is this because they selected "don't know" or did they attempt to answer and fail?

Figure 7: Name Recognition (Control)



Race





After separating "don't know" responses (coded as 99) from the individuals who attempted to answer the items and failed, we can see that the majority of individuals who answered 0 questions correctly was more due to the former than the latter. Figure 9 shows that the number of control-group respondents who didn't recognize Stephen Breyer (Q24, see Appendix 1) is almost the same amount as the number of those who correctly answered that he is a Supreme Court Justice. The remaining figures, 10 through 12, show lower rates of "don't know" answers, though the rates of respondents who incorrectly guessed are even lower.





Figure 10: Incorrect (0) vs. "Don't Know" (99) Responses, Secretary Item (Control)



Race





Figure 12: Incorrect (0) vs. "Don't Know" Responses, Secretary Item (Treatment)



Race

Before checking whether DIF was diminished in the treatment group, and whether this reduction is unique to the treatment group or the similar within the control group, we must first look to the raw perceptions data. These raw scale data provide an example of what inferences may be drawn by researchers without regard for any differential item functioning that may be present. The first two figures show White and Black control group respondents who place the Democratic Party on an 11-point Left-Right scale. The majority of both Black and White respondents placed the Democratic Party on the left-most side of the scale, close to 30 percent, indicating extreme values. Oddly enough, however, a few White and Black respondents positioned the Party towards the right side of the scale, perhaps indicating that they either misinterpreted the question or are simply unaware of the party's ideological leanings. This is a small yet ongoing trend.

Figure 13: White Respondents' Placement of Democratic Party on 11-Point Scale



Liberal-Conservative Scale

Figure 14: Black Respondents' Placement of Democratic Party on 11-Point Scale



Liberal-Conservative Scale

Instead of including a plethora of individual figures displaying L-R scale responses, the following table is a summary of *all* control and treatment raw scale responses—White and Black respondents' placement of the Democratic and Republican Party, Elizabeth Warren, and Marco Rubio (in percent).

Table 1 shows that a plurality of White and Black control-group respondents placed the Republican Party towards the right-most end of the scale (38.07 and 46.34 per cent, respectively), again indicating extreme views. One noteworthy trend considering the placement of both parties on the L-R scale across racial groups is that Black respondents—in both control and treatment groups placed Democrats further to the left and Republicans further to the right than White respondents.

		0	1	2	3	4	5	6	7	8	9	10
	Control Group (White)											
	Democratic Party	26.59	12.41	19.50	12.76	10.63	5.31	2.83	3.19	1.41	1.77	3.54
	Republican Party	2.84	2.49	1.42	1.06	2.49	3.91	2.84	6.04	19.92	18.86	38.07
	Elizabeth Warren	27.04	17.08	17.79	11.38	5.33	11.38	1.42	2.13	2.13	1.77	2.49
	Marco Rubio	1.41	1.41	2.48	1.41	4.60	8.51	3.19	12.05	22.34	23.40	19.14
	Control Group (Black)											
	Democratic Party	31.70	12.19	19.51	2.43	4.87	17.07	7.31	0.00	2.43	2.43	0.00
	Republican Party	7.31	7.31	4.87	0.00	2.43	7.31	0.00	4.87	7.31	12.19	46.34
	Elizabeth Warren	31.70	19.51	7.31	7.31	7.31	14.63	0.00	2.43	2.43	4.87	2.43
	Marco Rubio	4.87	2.43	0.00	2.43	0.00	12.19	0.00	12.19	19.51	12.19	34.14
	Treatment Group (White)											
	Democratic Party	27.01	11.92	20.00	10.87	10.87	7.01	2.10	3.15	1.75	2.10	3.15
	Republican Party	2.81	3.16	2.11	0.70	2.81	4.57	2.46	5.98	19.36	17.95	38.02
	Elizabeth Warren	26.40	17.95	15.49	11.97	5.63	11.61	1.05	2.46	2.46	2.46	2.46
	Marco Rubio	1.05	1.75	2.45	1.75	3.85	8.42	2.80	13.68	23.50	20.35	20.35
CEU eTD Collection	Treatment Group (Black)											
	Democratic Party	28.94	15.78	15.78	15.78	2.63	5.26	13.15	0.00	0.00	0.00	2.63
	Republican Party	7.89	2.63	0.00	2.63	0.00	2.63	2.63	5.26	10.52	18.42	47.36
	Elizabeth Warren	36.84	13.15	23.68	2.63	5.26	13.15	2.63	0.00	0.00	0.00	2.63
	Marco Rubio	7.89	0.00	0.00	0.00	5.26	13.15	2.63	0.00	10.52	34.21	26.31

Table 1: Respondents' Placements on 11-Point L-R Scale (in Percent)

This trend holds for the positioning of Elizabeth Warren and Marco Rubio. Black Respondents across control and treatment groups—placed Warren and Rubio further to the Left and Right (respectively) than White respondents.

While it is difficult to assess the variation between White and Black control and treatment groups, the inclusion of these raw data is important in displaying the scores that would normally be interpreted by social scientists without identifying whether or not differential item functioning is present. It is unclear whether inferences drawn from these data would impact perceptions about the levels of political knowledge across racial groups⁹—especially considering that a similar proportion of the majority of both groups of respondents, in control and treatment groups, placed both parties on the "correct"¹⁰ side of the scale, with only a small minority of respondents either misinterpreting the item instructions or lacking knowledge of candidate and party ideologies.

Interestingly enough, the base measures of political knowledge and the distribution of "correct" responses from the Left-Right scales suggest that White and Black respondents appear to know around the same amount about politics. Regarding absolute numbers, White respondents certainly have an advantage, but proportionally speaking, across treatment and control groups, responses from White and Black respondents are relatively similar. However, what happens when we correct for DIF?

⁹ Though, the higher percentage of Black respondents positioning parties and candidates to the extreme right and left perhaps adds to the picture.

¹⁰ By correct I mean that the majority of respondents placed parties and candidates on the side of the scale that correctly corresponds to the party or candidates actual position. Though we cannot definitively place these points on their true locations, they are spatially close to what is generally accepted as being correct.

Figure 15: Blackbox Transpose Corrected Estimates (White Respondents)



Figure 15 shows the results of the Blackbox Transpose on data from White respondents (in both control and treatment groups, blue and red points, respectively). The data displayed in the figure are the average raw scale responses (x axis) plotted against the corrected estimates (y axis). Like the x axis, the y axis may be read from left to right, liberal to conservative, on a scale from -1 to +1. As mentioned in Chapter 2, the Blackbox Transpose accounts for perceptual biases through the inclusion of stretch parameters to provide a closer estimation of the respondent's positioning

of a candidate or political party's true spatial location. A straightforward way to interpret this (and the following) graph would be to note deviations from the 45-degree line, as perfect correlation between the raw and corrected data would be linear; points far away from this line indicate that DIF is more prominent. Though some points appear to be close to this line, tests for Pearson's correlation between all raw and corrected estimates output negligible results (between r = -0.10 and r = 0.10), suggesting that we cannot use this scaling method to assess whether DIF was reduced.





Similar to Figure 15, Figure 16 shows the results of the Blackbox Transpose on data from Black respondents in both control and treatment groups (blue and red points, respectively). Also similar to Figure 15, though some points appear to be closer to and farther away from the 45-degree line, tests for Pearson's correlation between raw and corrected estimates output negligible results (between r = -0.10 and r = 0.10), suggesting that we cannot use this scaling method to assess whether DIF was diminished.¹¹

3.6.1 Results (Regression Analysis)

From these findings, one may think that race does not matter in an analysis of political knowledge. To see whether this is the case, two linear models were tested:

1: Base Knowledge (C) =
$$age + age^2 + gender + education + race + interestpolential education + race + interestpolenti$$

2: Base Knowledge $(T) = age + age^2 + gender + education + race + interestpol$

Both models are the same with the exception that the first looks at the combination of correct answers to the control group base knowledge items whereas the second looks at the treatment group base knowledge items. The models test whether there are significant relationships between these base knowledge items and respondent age, gender, education, race, and their interest in politics (it is safe to assume that this last independent variable will definitely have an effect on an individual's level of base knowledge—if someone has a high interest in something they are likely to teach themselves about it).

¹¹ Similarly, model fit for both Figures 15 and 16 were almost negligible.

Table 2 shows the multiple regression analysis of factors influencing control-group respondents' level of base political knowledge (model 1). There is a statistically significant relationship between the dependent variable Base Knowledge and the independent variables gender (p < 0.10), education (p < 0.01), race (p < 0.05) and the respondent's level of interest in politics (p < 0.01). Thus, we can reject the implied null hypothesis of race having no effect on an individual's level of base political knowledge.

	Dependent Variable:
	Base Knowledge (Control)
age	0.032
	(0.026)
age sq.	-0.0001
	(0.0003)
gender	-0.175^{*}
	(0.104)
education	0.178^{***}
	(0.063)
race	-0.354^{**}
	(0.155)
interestpol	-0.129^{***}
	(0.036)
Constant	2.171^{***}
	(0.621)
Observations	297
\mathbb{R}^2	0.203
Adjusted \mathbb{R}^2	0.187
Residual Std. Error	$0.860 \; (df = 290)$
F Statistic	12.347^{***} (df = 6; 290)
Note:	*p<0.1; **p<0.05; ***p<0.01
	48

Table 2: Multiple Regression Analysis of Factors Influencing Control-Group Respondents' Level of Political Knowledge (Model 1)

Table 3 paints an interesting picture. In looking at the factors influencing treatment-group respondents, it appears that race no longer has a significant effect on respondents' level of political knowledge. This will be further explored in the following section.

	Dependent Variable:
	Base Knowledge (Treatment)
age	0.014
	(0.023)
age sq.	-0.00002
	(0.0003)
gender	-0.408^{***}
	(0.102)
education	0.155^{***}
	(0.059)
race	-0.059
	(0.155)
interestpol	-0.122^{***}
	(0.037)
Constant	2.846^{***}
	(0.543)
Observations	306
\mathbb{R}^2	0.151
Adjusted \mathbb{R}^2	0.134
Residual Std. Error	$0.864 \; (df = 299)$
F Statistic	8.854^{***} (df = 6; 299)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 3: Multiple Regression Analysis of Factors Influencing Treatment-Group Respondents'Level of Political Knowledge (Model 2)

3.7. Discussion

Based on the above results, we can accept the null hypotheses corresponding to H₁ (The inclusion of African American-specific items will reduce the substantive gap in political knowledge between African American and White respondents) and H₂ (The inclusion of African American-specific items will show a visible reduction of differential item functioning when plotting average item responses against DIF-corrected estimates.). The inclusion of African American-specific items did not reduce the substantive gap in political knowledge between African Americans and White respondents. The treatment did show a substantive increase in the amount of correct item responses, however, this increase was relatively proportional between White and Black respondents. Additionally, paired-sample t tests showed that the treatment *did* have an effect, however, we cannot say how much of an effect it had aside from the substantive evidence. Regarding H₂, findings were inconclusive. Though the positioning of average raw scale responses to their corrected estimates showed substantively interesting results, we cannot use them to explain differential item functioning due to all correlation coefficients being negligible. This result is in opposition to Abrajano's findings—though perhaps this is due to the sample; the MTurk sample used for this experiment was smaller than the ANES sample that Abrajano utilized. However, we can reject the null hypothesis corresponding to H₃ (The inclusion of African American-specific items will not diminish the average reported knowledge of White respondents) as the inclusion of racially-specific items did not reduce the average reported knowledge of White respondents.

Though this might lead one to think that race has no effect on an individual's level of political knowledge, results from the above regression analysis indicate otherwise. There is a statistically significant relationship between race and political knowledge in the control group, but not in the

treatment group. We do not have sufficient evidence to claim that the treatment is responsible for "leveling the playing field" between the races, so to speak, however, further research should be conducted to see whether this is the case.

3.8. Limitations

The primary limitation experienced by this study was the disconnect in methods utilized by attempting to combine Abrajano and Pérez's research. The Blackbox Transpose only looks at responses on issue scales, whereas this experiment's predicted differential item functioning stemmed from the relationship between the base political knowledge items and L-R scale items. The base political knowledge items were included with the intention of creating designated "anchors" to compare L-R scale responses to, however, I was initially unaware that the Blackbox Transpose does not account for these anchors.

Another, albeit smaller, limitation to this study was the platform through which this experiment's sample was collected. Though Mechanical Turk was the most realistic and convenient platform to use for a Master's thesis, further work on this subject would benefit from administering a larger-scale survey that provides more of a nationally representative sample. I did not include locational data in my analysis, however, glancing over the average location of respondents showed that the majority of MTurk users taking my survey were based in major cities around the U.S., primarily on the West and East coasts, with a small portion in the Southwest and South. Additionally, the "bonus payment" system within MTurk provided some difficulties—the larger payment of \$0.50 for White and Black respondents was contingent upon pre-survey responses (based on race), therefore, MTurk requires researchers to advertise the study as having a base payment (in this case

\$0.10) with a potential bonus. Consequently, the study was advertised (and searchable) as paying only \$0.10, which is perhaps the cause of the slower-than-expected rate of survey responses.

In a similar vein, though my sample of African Americans respondents was oversampled and certainly large enough to conduct descriptive statistics and the Blackbox Transpose scaling procedure, an even larger sample would have provided more of a foundation for stronger, more robust analysis.

On a more technical note, CRAN—the platform for downloading R libraries—very recently archived 'basicspace', the library created to conduct the Aldrich-McKelvey scaling procedure and the Blackbox Transpose. It seems that its creators neglected to update the program, indicating that perhaps the Blackbox transpose is no longer a viable or frequently used enough method to test for differential item functioning. There are certainly costs and benefits to this measurement, however, it would be interesting to see whether a different scaling procedure would produce dissimilar results.

3.9. Conclusion

This paper sought to introduce the subject of political knowledge and the knowledge "gaps" existing between racial groups—particularly the gap between White and African American respondents—with the intention of conducting an experiment to test whether item bias could be partially corrected through the inclusion of African American-specific questions. Chapter 1 presented an argument against disenfranchising citizens with low levels of political knowledge—disenfranchisement, whether full or partial, is morally unjust and practically difficult to do due to the fragile state of measuring political knowledge, as this thesis has illustrated. Chapter 2 provided

a brief overview of what political knowledge actually is through categorizing various pieces of literature into three primary groups, explored how political knowledge is measured, and the costs and benefits of different types of tools used to measure political knowledge. Chapter 3 covered this paper's survey experiment—though its treatment had an initial effect, it is difficult to assess the extent of this effect. Measuring for differential item functioning using the Blackbox Transpose proved to be ineffective. Pérez's research suggests that the inclusion of political knowledge items relevant to the Latino community better captures their measured levels of political knowledge without diminishing the knowledge levels reported for White respondents. Through using similar logic, I have shown that we cannot necessarily say the same of the African American community, though regression analysis indicates some statistically significant relationship between race and political knowledge. However, a non-finding is a finding within itself-further research on the subject should dive deeper into the potential methodological and socio-historical reasons why African Americans may interpret knowledge items in a different way from majority-group respondents. Additionally, alternate scaling methods testing for differential item functioning should be used to see whether or not current items can be changed to more accurately measure the levels of political knowledge of not just the majority, but *all* Americans.

References

- Abrajano, Marisa. 2015. "Reexamining the 'Racial Gap' in Political Knowledge." The Journal of Politics 77 (1): 44–54. doi:10.1086/678767.
- Aldrich, John H., and Richard D. McKelvey. 1977. "A Method of Scaling with Applications to the 1968 and 1972 Presidential Elections." The American Political Science Review 71 (1)
- Althaus, Scott L. 1998. "Information Effects in Collective Preferences." *American Political Science Review* 92 (3): 545–58.
- Andersen, Kim, Camilla Bjarnøe, Erik Albæk, and Claes H. De Vreese. 2016. "How News Type Matters: Indirect Effects of Media Use on Political Participation Through Knowledge and Efficacy." *Journal of Media Psychology* 28 (3): 111–22.
- Armstrong, David, Ryan Bakker, Royce Carroll, Christopher Hare, Keith T. Poole and Howard Rosenthal. 2014. *Analyzing Spatial Models of Choice and Judgement with R*. Boca Raton: CRC Press.
- Bartels, Larry. 1996. "Uninformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science* 40 (1): 194–230.
- Berinsky AJ, G.A. Huber and G.S. Lenz (2012) Evaluating online labor markets for experimental research: Amazon.com's mechanical Turk. *Political Analysis* 20: 351–368.
- Biggers, Daniel R. 2011. "When Ballot Issues Matter: Social Issue Ballot Measures and Their Impact on Turnout." *Political Behavior* 33 (1): 3–25.
- Biggers, Daniel R. 2012. "Can a Social Issue Proposition Increase Political Knowledge? Campaign Learning and the Educative Effects of Direct Democracy." *American Politics Research* 40 (6): 998–1025.
- Bowler, Shaun, and Todd Donovan. 2002. "Democracy, Institutions and Attitudes about Citizen Influence on Government." *British Journal of Political Science* 32 (2): 371–390.
- Brady, Henry E. 1985. "The Perils of Survey Research: Inter-Personally Incomparable Responses." *Political Methodology*, 269–291.
- Brennan, Jason. 2011. "The Right to a Competent Electorate." *The Philosophical Quarterly* 61 (245): 700–724.
- Clifford, Scott, and Jennifer Jerit. 2016. "Cheating on Political Knowledge Questions in Online Surveys." *Public Opinion Quarterly* 80 (4): 858–87.
- Christiano, Thomas. 2004a. "An Argument for Democratic Equality" in Christiano (ed.) *Philosophy and Democracy*, Ch. 2.

- Christiano, Thomas. 2004b. "The Authority of Democracy." *Journal of Political Philosophy*, 12 (3): 266–290.
- DeBell, Matthew. 2013. "Harder Than It Looks: Coding Political Knowledge on the ANES." *Political Analysis* 21 (4): 393–406.
- Delli Carpini, Michael X., and Scott Keeter. 1993. "Measuring Political Knowledge: Putting Things First." *American Journal of Political Science* 37 (4): 1179–1206.
- Ferrín, Mónica, and Marta Fraile. 2014. "Measuring Political Knowledge in Spain: Problems and Consequences for the Gender Gap in Knowledge." *Revista Española de Investigaciones Sociológicas*.
- Galston, William A. 2001. "Political Knowledge, Political Engagement, and Civic Education." *Annual Review of Political Science* 4 (1): 217–234.
- Gay, Claudine. 2014. "Knowledge Matters: Policy Cross-Pressures and Black Partisanship." *Political Behavior* 36 (1): 99–124.
- Graber, Doris. 2004. "Mediated Politics and Citizenship in the Twenty-First Century." *Annual Review of Psychology* 55 (1): 545–71.
- Huff, Connor, and Dustin Tingley. 2015. "Who Are These People?' Evaluating the Demographic Characteristics and Political Preferences of MTurk Survey Respondents." Research & Politics 2 (3)
- Hutchens, Myiah, Jay Hmielowski, Bruce Pinkleton, and Michael Beam. 2016. "A Spiral of Skepticism? The Relationship Between Citizens' Involvement with Campaign Information to Their Skepticism and Political Knowledge." *Journalism & Mass Communication Quarterly* 1: 1073–90.
- Keeter, Scott. 2011. "Political Knowledge." In *Encyclopedia of Survey Research Methods*, edited by Paul J. Lavrakas, 587. Thousand Oaks, CA: SAGE Publications, Inc., 2008.
- King, Gary, and Jonathan Wand. 2007. "Comparing Incomparable Survey Responses: Evaluating and Selecting Anchoring Vignettes." Political Analysis 15 (1): 46–66.
- King, Gary, Christopher J. L. Murray, Joshua A. Salomon, and Ajay Tandon. 2004. "Enhancing the Validity and Cross-Cultural Comparability of Measurement in Survey Research." American Political Science Review 98 (1): 191–207.
- Krosnick, Jon A., Arthur Lupia, Matthew DeBell, and Darrell Donakowski. 2008. "Problems with ANES Questions Measuring Political Knowledge." *ANES Report, March.*
- Kuklinski, James H., and Paul J. Quirk. 2001. "Conceptual Foundations of Citizen Competence." *Political Behavior* 23 (3): 285–311.

- Ladha, Krishna K. 1992. "The Condorcet Jury Theorem, Free Speech, and Correlated Votes." *American Journal of Political Science* 36 (3): 617.
- Lau, Richard R., and David P. Redlawsk. 2001. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making." *American Journal of Political Science* 45 (4): 951.
- Lupia, Arthur. 1994. "Shortcuts versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections." *American Political Science Review*, 63–76.
- McKelvey, Richard D., and Peter C. Ordeshook. 1985. "Sequential Elections with Limited Information." *American Journal of Political Science* 29 (3): 480.
- Mondak, Jeffery 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.
- Necka, Elizabeth A., Stephanie Cacioppo, Greg J. Norman, and John T. Cacioppo. 2016. "Measuring the Prevalence of Problematic Respondent Behaviors among MTurk, Campus, and Community Participants." Edited by Jelte M. Wicherts. *PLOS ONE* 11 (6): e0157732.
- Pérez, Efrén O. 2015. "Mind the Gap: Why Large Group Deficits in Political Knowledge Emerge—And What to Do About Them." *Political Behavior* 37 (4): 933–54.
- Poole, Keith T. 1998. "Recovering a Basic Space from a Set of Issue Scales." *American Journal* of *Political Science* 42 (September): 954-93.
- Pzeworski, Adam and Henry Teune. 1966-67. "Equivalence in Cross-National Research." *Public Opinion Quarterly* 30 (Winter): 551-68.
- Rawls, John. 1958. A Theory of Justice. Harvard University Press.
- Schlozman, Daniel, and Ian Yohai. 2008. "How Initiatives Don't Always Make Citizens: Ballot Initiatives in the American States, 1978–2004." *Political Behavior* 30 (4): 469–89.
- Shaker, Lee. 2012. "Local Political Knowledge and Assessments of Citizen Competence." *Public Opinion Quarterly* 76 (3): 525–37.
- Thissen, David, Lynn Steinberg, and Howard Wainer. 1993. "Detection of Differential Item Functioning Using the Parameters of the Item Response Models." In Differential Item Functioning. ed. Paul W. Holland and Howard Wainer. Hillsdale, NJ: Lawrence Erlbaum.
- Todorov, Alexander, Anesu N. Mandisodza, Amir Goren, and Crystal C. Hall. 2005. "Inference of Competence from Faces Predict Election Outcomes." *Science* 303 (5728): 1623–26.

United States Census Bureau. https://www.census.gov/quickfacts/table/PST045216/00

- Weith, Paul T. 2015. "Information Effects Across the Globe: Does Political Knowledge Matter." Central European University.
- Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady. 1995. Voice and Equality: Civic Volunteerism in American Politics. Harvard University Press.

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

Appendix 1: MTurk Survey

Thank you for participating in this survey. Before you begin, we would like to inform you that participation in this survey is voluntary and you may stop at any time. Your responses will remain anonymous, and are being used for a study on political knowledge conducted by Jason Means under the supervision of Dr. Gabor Toka and Dr. Levente Littvay. Should you have any questions, please email Jason Means at surveys@ceu.edu. Please note that there are two versions of this survey: One longer, and one shorter. If you take the shorter survey you will receive \$0.10. If you take the longer survey you will receive \$0.10 and a bonus of \$0.40 (\$0.50 total). Survey selection is based on three questions on the next page. You will receive a randomized code at the end of the survey. Please enter this code into Mechanical Turk to complete your HIT and receive payment. By continuing this survey, you agree that you understand the above information. Again, thank you.

PRE-SURVEY:

Q2 Which of the following colors do you most prefer?

Purple (1)

Orange (2)

Green (3)

Q3 To what extent do you agree with the following statement: I am interested in politics.

Strongly Agree (1)

Agree (2)

Somewhat agree (3)

Neither agree nor disagree (4)

Somewhat disagree (5)

Disagree (6)

Strongly disagree (7)

Q4 What is your racial identity?

White (1)

Black or African American (2)

Hispanic (3)

American Indian or Alaska Native (4)

Asian (5)

Native Hawaiian or Pacific Islander (6)

Other (7)

CONTROL:

Q5 How many times can an individual be elected President of the United States under current laws?

1 time (1)

2 times (2)

3 times (3)

4 times (4)

5 times (5)

I don't know (6)

Q6 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 (1)

4 years (2)

6 years (3)

8 years (4)

10 years (5)

I don't know (6)

Q7 What is Medicare?

A program run by the U.S. Federal government to pay for old people's health care. (1)

A program run by state governments to provide health care to poor people. (2)

A private health insurance plan sold to individuals in all 50 states. (3)

A private, non-profit organization that runs free health clinics. (4)

I don't know (5)

Q8 On which of the following does the U.S. federal government currently spend the least? Foreign Aid (1) Medicare (2) National Defense (3) Social Security (4)

I don't know (5)

Q9 What job or political office does Stephen Breyer now hold?

Governor of Wisconsin (1)

Secretary of Energy (2)

Supreme Court Justice (3)

Senator (4)

I don't know (5)

Q10 What job or political office does Rick Perry now hold?

Governor of Wisconsin (1)

Secretary of Energy (2)

Supreme Court Justice (3)

Senator (4)

I don't know (5)

TREATMENT:

Q21 How many times can an individual be elected President of the United States under current laws?

1 time (1)

- 2 times (2)
- 3 times (3)
- 4 times (4)
- 5 times (5)

I don't know (6)

- Q22 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 (1)
- **4 years (2)**
- 6 years (3)
- 8 years (4)

10 years (5)

I don't know (6)

Q23 What is Medicare?

A program run by the U.S. Federal government to pay for old people's health care. (1) A program run by state governments to provide health care to poor people. (2) A private health insurance plan sold to individuals in all 50 states. (3) A private, non-profit organization that runs free health clinics. (4) I don't know (5)

Q24 On which of the following does the U.S. federal government currently spend the least? Foreign Aid (1) Medicare (2) National Defense (3) Social Security (4) I don't know (5)

Q25 What job or political office does Clarence Thomas now hold? Governor of Wisconsin (1) Secretary of Housing and Urban Development (2) Supreme Court Justice (3) Senator (4) I don't know (5) Q26 What job or political office does Ben Carson now hold? Governor of Wisconsin (1)

Secretary of Housing and Urban Development (2)

Supreme Court Justice (3)

Senator (4)

I don't know (5)

BOTH CONTROL AND TREATMENT:

- Q31 In politics people sometimes talk of left and right. Where would you place the Democratic Party on a scale from 0 to 10, where 0 means the left and 10 means the right?
- 0 (11)
- 1 (13)
- 2 (2)
- 3 (3)
- 4 (4)
- 5 (5)
- 6 (6)
- 7 (7)
- 8 (8)
- 9 (9)
- 10 (10)

Q32 Using the same scale, where would you place the Republican Party?

- 0 (2)
- 1 (11)
- 2 (12)
- 3 (3)
- 4 (4)
- 5 (5)

6	(6)
7	(7)
8	(8)
9	(9)

10 (10)

Q33 Where would you place Elizabeth Warren on this scale?

0 (1) 1 (11) 2 (2) 3 (3) 4 (4) 5 (5) 6 (6) 7 (7) 8 (8)

- Q34 Where would you place Marco Rubio on this scale?
- 0 (1)

9 (9)

10 (10)

- 1 (11)
- 2 (2)
- 3 (3)
- 4 (4)
5 (5)

6 (6)

- 7 (7)
- 8 (8)
- 9 (9)
- 10 (10)

Q35 What is your age?

Q26 What is your gender?

Male (1)

Female (2)

Q37 Which of the following political parties do you feel closest to?

Republican Party (1)

Democratic Party (2)

Other (Please Specify): (3) _____

Q38 Do you feel very close, somewhat close, or not very close to this party?

Very close (1)

Somewhat close (2)

Not very close (3)

Q39 What is your highest level of completed education?

Primary (1)

Secondary (High School or equivalent) (2)

Undergraduate (BA or equivalent) (3)

Graduate (MA or equivalent) (4)

PhD or equivalent (5)

FULL SURVEY RESPONDENTS:

- Num Your personalized code is XXXXXX. Please enter this code into Mechanical Turk to complete your HIT and receive bonus payment.
- Q40 Thank you for participating in our survey. If you have any feedback or comments, please enter them in the box below.

SHORT SURVEY RESPONDENTS:

- NumS Your personalized code is XXX. Please enter this code into Mechanical Turk to complete your HIT and receive payment.
- Q41 Thank you for participating in our survey. If you have any feedback or comments, please enter them in the box below.