

Using Experiments to Estimate Geographic Variation in Racially Polarized Voting

Marisa A. Abrajano* Christopher S. Elmendorf† Kevin M. Quinn‡

April 21, 2014

To the Reader: *This is a preliminary and incomplete writeup of what we hope will be a larger project. The data used in this paper are from an early pilot survey that utilized a convenience sample of mTurk workers. These data are not representative of the US voting age population, citizen voting age population, or state (citizen) voting age populations. Our primary goal in disseminating these early results of this pilot study is to elicit feedback on the study design. **Please do not distribute or cite without permission.***

1 Introduction

Political scientists have long been interested in racialized political identity and its expression at the ballot box. These phenomena are also centrally important to voting rights law. In most cases under the Voting Rights Act (VRA), plaintiffs must show that white and minority voters consistently prefer own-race candidates in biracial elections. In the lingo of the law, this is called racially polarized voting, and is understood to signify that white and minority voters belong to politically cohesive, mutually opposed communities (*Thornberg v. Gingles*, 478 U.S. 30 [1986]; Issacharoff 1995). Increasingly, the courts also ask whether race “caused” minority candidates’ lack of success (Greiner 2008; Elmendorf 2012). Judges view racially polarized voting as more innocuous if it arises in a race-blind manner, e.g., as a byproduct of partisan voting under conditions where most minority candidates self-select into one political party, than if it reflects disparate treatment

*UC San Diego Department of Political Science

†UC Davis School of Law

‡UC Berkeley School of Law

of otherwise similar white and minority candidates.

Despite considerable disciplinary interest in the subject of minority voting rights, the body of work which political scientists have produced is often unable to reliably answer the central questions in voting rights cases. Methodologists have made important contributions to the art and science of ecological inference (EI), the standard statistical tool of vote-dilution litigation, but recent research highlights the difficulty of recovering race-correlated voting patterns in jurisdictions with more than two racial groups and non-trivial residential integration (Greiner and Quinn 2010; Greiner 2011). EI performs much better if the analyst supplements aggregate with individual-level data (Haneuse and Wakefield 2008; Glynn et al. 2008; Glynn and Wakefield 2010; Greiner and Quinn 2010), but individual-level data are “almost never available” in voting rights cases (Grofman 1992, p. 217). And even if the analyst can produce reliable EI estimates, it remains very difficult to estimate with observational data the number of additional votes that minority candidates would have garnered had they been perceived to be white (Greiner 2008).

Recognition of this point underlies the recent profusion of experimental research seeking to identify the causal effect of a candidate’s race, ethnicity, or skin tone on voter beliefs and preferences regarding the candidate or associated policies. Experimental research on race and politics has generated many important insights, but in a couple of respects it falls short of what is needed for VRA purposes. First, to the extent that experimentalists worry about external validity, they generally design their studies with the goal of generalizing to the national population. Yet for voting rights law and policy, the question of how political behavior *varies* across regions, states, and smaller geographic units is more important. Vote-dilution cases always concern particular locales, and courts give little weight to evidence concerning political behavior in the nation as a whole (*Grove v. Emison*, 507 U.S. 25 [1993]).

Second, standard experiments shed little light on *how much* the race of a candidate generally affects voter support for the candidate. For experiments conducted with hypothetical candidates, the how-much question is hard to answer because the typical design holds constant all or nearly

all attributes of the candidate other than his race. This strategy avoids confounds but leaves open the question of whether the size or even the direction of the estimated “race effect” is conditional on certain non-varied attributes of the candidates. Similarly, it is hard to know whether the experimental race effect would dissipate—or perhaps strengthen—if subjects were exposed to additional, campaign-generated information about the candidates. For experiments involving real candidates in real elections it is, paradoxically, even harder to establish the real-world magnitude of the effect of candidate race. An obvious problem is that the race of real candidates is fixed, so the only treatment effects that researchers can estimate concern peripherally related matters such as representations of the candidate’s skin tone (light or dark), or of the race of supporters pictured in a campaign photo. Estimated treatment effects are also likely to be biased downward owing to pretreatment of experimental subjects by the campaign itself (Druckman and Leeper 2012). For example, if a prominent black candidate has many black supporters, some experimental subjects are likely to have figured this out prior to being assigned to experimental conditions.

Perhaps it is not surprising, then, that experimental research has played no role to date in vote dilution litigation. Nor were experimental studies featured in the legislative record developed in 2006 when Congress reauthorized the VRA’s “preclearance” provisions, which require certain states and political subdivisions to obtain approval from the federal government before implementing changes to their election laws and procedures. Congress’s extension of the preclearance regime was subsequently overturned by the Supreme Court on the ground that Congress had failed to establish that the problem of race discrimination in voting is more severe in jurisdictions covered by the preclearance regime than elsewhere (*Shelby County v. Holder*, 113 S.Ct. 2612 [2013]). The *Shelby County* decision makes it all the more urgent for political scientists to develop better estimates of the geography of racial discrimination in the electoral process.

The present paper develops a research design for studying racially polarized voting and associated questions about racially disparate treatment experimentally. Our design—which supports credible causal inferences about the effects of candidates’ apparent race and racial interest group

endorsements—sheds considerable light on the “where” question and makes some headway on the “how much” question. Building on Hainmueller et al. (2014), we use a conjoint design in which multiple candidate attributes are simultaneously randomized. Instead of explicitly describing a candidate’s race with words in a table, however, we signal race using visual images and racially identifiable names. To the extent that voters have a visceral reaction to candidates’ race yet also subscribe to norms against discriminating on the basis of race (Mendelberg 2001; Kinder and Sanders 1996), photographs as opposed to words are likely to better elicit the reactions voters have to actual minority candidates. We also randomize the candidates’ endorsements, providing both partisan and racial interest group signals. This enables us to identify presumptive “candidates of choice” of the minority community in much the same manner as courts do. It also helps us to gauge whether citizens treat race as a proxy for partisanship, or whether racial associations have further and perhaps more invidious effects on voter decision-making.

We analyze the data using multilevel regression with poststratification, generating state-specific estimates of polarization between whites and Latinos, whites and African Americans, and Latinos and African Americans. We also provide nonparametric estimates of the effects of candidate race and racial interest group endorsements, conditional on office, for the nation as a whole.

Within our convenience sample of mTurk workers we find little evidence that white voters penalize minority candidates solely due to the candidates’ race. We do find evidence that minority voters are more likely to vote for a coethnic candidate than a non-coethnic candidate. We also find evidence of relatively large levels of racially polarized voting within our full convenience sample of mTurk workers, and suggestive evidence of even higher levels of racially polarized voting in a few large states where we have a large number of respondents.

2 Candidate Race and Voter Choice

Political scientists have devoted enormous effort to understanding why and how voters take account of candidates’ race when deciding whom to support. We know that voters may respond

differently to explicit and implicit racial appeals (e.g., Mendelberg 2001); we know that white voters may react to stereotype-conforming information about minority candidates by stereotyping the candidate in other, more socially acceptable ways (Berinsky et al. 2011); and we know that exogenous variation in the salience or proximity of persons of other racial groups can affect voter support for candidates of that group (Enos 2010, 2011). It is also understood that voters may take account of candidates' race for either statistical or taste-based reasons (Becker 1971; Arrow 1972; McDermott 1998b; Marisa A. Abrajano and Nagler 2005; Karl 2014), and that it is very difficult to decompose race-treatment effects into statistical and taste-based components (Pager and Shepherd 2008).

Political scientists do not, however, have a very good grasp on how much the race of a candidate affects voter support for the candidate, or of how this effect varies across racial groups and political jurisdictions. Several researchers have tried to make headway on these questions using Barack Obama as a case study. Ansolabehere et al. (2010) show that the difference between Obama's estimated vote share among whites in 2008 and Kerry's vote share in 2004 was greatest in the South and in Appalachia. Jackman and Vavrek (2011), using survey data on head-to-head matchups of various Democratic and Republican candidates for President, estimate that an average white Democratic candidate for President would have done about three percentage points better than Obama in 2008. Other researchers have tried to leverage variation in the race and ethnicity of Members of Congress, or state and city-level representatives, to estimate race effects (e.g., Citrin et al. (1990); Browning and Tabb (2002); Highton (2004); Hajnal (2007); Branton et al. (2012); Ansolabehere and Fraga (2013)).

Observational studies cannot, however, tell us very much about variation across subnational jurisdictions in the typical effect of a candidate's apparent race on voter support for the candidate. To obtain state-specific estimates of the effect of candidate race, one needs to be able to sample from the population of voters in each state and elicit evaluations of exchangeable white and minority-race candidates from exchangeable respondents. But this minimal condition is (plausibly) met only in the Obama-based studies. In studies based on vote choice in congressional elections

or respondents' evaluations of their congressional representative, it is redistricters and campaign strategists who in effect determine whether a particular respondent is asked to evaluate a white candidate, a minority candidate, or both. Moreover, any geographic variation in "race effects" that such studies might document could be due to non-racial differences in the candidates being evaluated by voters. Louisiana, for example, has only one black Member of Congress, Cedric Richmond. An observational estimate of the effect of black race on white Louisianans' support for Members of Congress would require assuming both (1) that white voters in Cedric Richmond's district are similar to white voters elsewhere in the state and (2) that Cedric Richmond is similar to the other seven members of Louisiana's delegation but for his race. To compare the effect of candidate race in Louisiana with the effect of race in Mississippi, one would have to assume as well (3) that the black and white members of the Louisiana delegation are similar to those of the Mississippi delegation.

Even if one conditions on observables like partisanship and ideology, these assumptions are rather heroic. White voters in Cedric Richmond's district may well have a higher latent probability of supporting a black candidate than co-partisan or ideologically similar voters elsewhere in the Louisiana—after all, his was the only district to elect a black candidate. This would bias downward the estimated effect of candidate race on white voters in Louisiana. If Richmond is an exceptionally good or bad representative, relative to white members of Louisiana's delegation or black members of Mississippi's delegation, this will also bias estimates of state-specific race effects and between-state differences. Because most states have only a few minority members in their congressional delegations, this is not a small concern. Finally, strategic anticipation of race discrimination by potential candidates and campaign donors may well result in large, unmeasured differences between white and minority candidates and legislators (Greiner (2008)), confounding both within-state and between-state estimates of the effect of candidate race. Black candidates who run for office in states or districts where black candidates have a reasonable prospect of securing many white votes are likely to be importantly distinct from black candidates who run in states where white votes are hard to win.

Obama-based studies of the effect of candidate race avoid some of the difficulties with studies based on data from congressional, state, or municipal elections. The candidates (treatments) being compared are the same across subjects in different states and regions, and the exit poll and reweighted opt-in samples on which the Obama studies rely are plausibly representative of the population of voters. Yet it goes almost without saying that President Obama is not a typical case (Tessler and Sears 2010; King and Smith 2011). The effect of race in Obama's case is the impact of race as it pertains to a light-skinned black man who was born to a white mother and black father, raised by his white mother and white grandparents, educated at Columbia and Harvard, and who ran in a relatively high-information election. Obama-based studies tell us very little about the treatment effect of candidate race in lower profile elections with ordinary black candidates, and of course they say nothing about the effect of race on voter support for Latino or Asian American candidates. Nor do the Obama-based studies allow one to isolate the effect of Obama's race from other aspects of his background, personal character, or policy positions that differentiate him from other Democratic candidates for President.

Another way to tackle the question of geographic variation in voter discrimination against minority-race candidates is to estimate variation in some underlying voter characteristic that is associated with support for minority candidates. Pursuing this approach, Elmendorf and Spencer (2014c; 2014a) show that blacks are stereotyped most negatively in the Deep South, but that geographic patterns in negative stereotyping of Latinos and Asian-Americans are less clear cut. Their strategy does not, however, yield a measure of actual disparate treatment on the basis of candidate race, and the survey responses on which it relies may be distorted to some extent by social desirability biases (Huddy and Feldman 2009).

The remaining option is to use survey experiments. Yet conventional survey experiments have serious limitations as a tool for estimating the magnitude of candidate-race effects and associated geographic variation. Because a proliferation of experimental conditions would dramatically reduce statistical power, experimentalists typically vary only one or two features of the choice scenario

presented to subjects. This leaves open the question of whether the treatment effect would be robust to variation in other features of the choice scenario. For example, might the effect of switching a candidate’s apparent race from white to black vary depending on the candidate’s education, income, or political party? Similar considerations of statistical power discourage the disaggregation of survey responses by geography. We make some headway on these problems with a conjoint design carefully tailored to the legal issues in vote dilution cases, and by using multilevel regression with poststratification to back out state-specific estimated treatment effects. This design enables us to estimate candidate-race effects averaged over a distribution of other candidate traits, and to compare the treatment effect of varying candidate race with the effects of variation in other familiar candidate attributes.

3 Study Design and Implementation

Because our study is meant to inform vote-dilution litigation under Section 2 of the Voting Rights Act, we briefly explain the legal issues before turning to the design. Vote-dilution cases proceed in two stages. At the first stage, plaintiffs must satisfy the so-called *Gingles* conditions: proving that plaintiff-race voters are politically cohesive; that other voters vote as a bloc against the plaintiff group’s “candidates of choice”; and that a majority-minority district could be created to enable the election of such candidates (*Thornburg v. Gingles*, 478 U.S. 30 [1986]; *Bartlett v. Strickland*, 556 U.S. 1 [2009]). Findings about racially polarized voting (RPV) have almost universally been based on estimates of white and minority vote shares for candidates the court regards as likely candidates of choice of the minority community (Katz et al. 2005). Courts make rough judgments about whether a candidate is plausibly a “candidate of choice” on the basis of the candidate’s race and his or her support—or lack of support—among minority interest groups and elites (Katz et al. 2005; Persily 2007; Hebert et al. 2010).

If plaintiffs surmount the *Gingles* hurdle, the court proceeds to a murky “totality of circumstances” analysis which governs the final determination of liability (*Johnson v. De Grandy*, 512

U.S. 997 [1994]). Courts and commentators disagree about how to understand the totality-of-circumstances inquiry. One recurring issue is “rough proportionality” between the number of ex-tant majority-minority districts in the jurisdiction and the minority’s population share, but other considerations are prominent too.

For example, many courts ask whether race is the “cause” of polarized voting (Katz et al. 2005; Greiner 2008; Elmendorf 2012; Elmendorf and Spencer 2014a).¹ Sometimes the causation question is phrased in terms of the race of candidates (Did minority candidates receive fewer votes than they would have received had they been white?); sometimes in terms of the race of voters (Did voters’ race “cause” them to vote the way that they did?). The most doctrinally coherent and empirically intelligible version of the causation question is to ask whether the plaintiff group’s lack of electoral success resulted at least in part from disparate treatment of plaintiff-race persons on the basis of their apparent race. For example, do plaintiff-race candidates in the defendant jurisdiction receive less support from white voters than they would have received had they been perceived to be white? This gloss on the causation question reconciles Section 2 with modern constitutional law (Elmendorf 2011), which defines race discrimination as disparate treatment on the basis of race.² And, unlike the question of whether voters’ race causes them to vote as they do, the disparate-treatment question is at least potentially answerable within the potential outcomes framework for causal inference.³

Some courts have suggested that the crux of the causation question is whether minority candidates lose because they are Democrats rather than because of their race (e.g., *League of United Latin American Cities v. Clements*, 999 F.2d 831 [5th Cir. 1993]; Ryan 2009). What the judges who pose this question mean by it is not entirely clear. Perhaps it’s just a rhetorical flourish on the fundamental matter of disparate treatment on the basis of race. A judge sees evidence of racially

¹Some courts incorporate the causation question into the *Gingles* conditions, but more often it is taken up at the totality-of-circumstances stage (Katz et al. (2005)).

²(*Washington v. Davis*, 426 U.S. 229 [1976]; *Personnel Administrator of Mass. v. Feeney*, 442 U.S. 256 [1979])

³To be sure, there are daunting barriers to trying to estimate the effect of perceived candidate race using observational data (Greiner 2008), but experiments offer a viable alternative.

polarized voting in biracial, partisan elections, but she suspects that white Democrats would have lost by similar margins. She asks, “Did the minority candidates lose because of their party or because of their race?” as an informal way of getting at the question of whether the minority Democrats who ran for office would have done any better among white voters had voters perceived them to be white Democrats.⁴

Another way of reading the “race or party” question is that courts mean to suggest that disparate treatment on the basis of race is unworthy of a legal remedy if race is merely being used to make inferences about partisanship or ideology, inferences which in turn motivate the disparate treatment. Disparate treatment is conceded but then excused. Modern equal protection law generally rejects the proposition that statistical discrimination merits less constitutional scrutiny than other kinds of discrimination on the basis of “suspect classifications” (such as race), yet the Supreme Court has intimated that statistical race discrimination for partisan reasons may be less constitutionally objectionable than other forms of racially disparate treatment, at least if the discrimination is not too pronounced or severe (*Easley v. Cromartie*, 532 U.S. 234 [2001]; Elmendorf and Spencer 2014c).

Easley was decided under the Equal Protection Clause rather than the VRA, but its notion that *some* race discrimination in politics is inevitable and must be tolerated harkens to yet another gloss on the totality-of-circumstances inquiry in Section 2 cases. On this view, the core issue is how thoroughly the minority community has been excluded from full participation in civic life (Issacharoff 1995). Proponents of this approach say that liability under Section 2 is warranted where racially polarized voting and/or disparate treatment is particularly severe; where there is a history of de jure discrimination whose effects still linger in the form of depressed minority turnout; where racial appeals and attacks are commonplace in political campaigns; and where no major, organized political factions actively campaigns for minority votes. Viewed holistically, these and

⁴Critics have pointed out that even if minority Democrats perform as well as otherwise similar white Democrats, it does not follow that minority candidates have not been hindered by race discrimination, because voters’ party identification itself is in an important sense post-treatment: voters may be aware of the association between racial minorities and the Democratic Party when their party identity forms or evolves (Karlan and Levinson 1996; Valentino and Sears 2005; Greiner 2008; Tesler 2013).

other factors should guide the courts as they try to mark the liability-determinative line between “ordinary politics” and extraordinary “racialized politics” (cf. *Whitcomb v. Chavis*, 403 U.S. 124 [1971]; *White v. Regester*, 412 U.S. 755 [1973]).

This paper does not advance an account of how Section 2 should be interpreted. Our goal, rather, is to design and execute an experiment whose results will be pertinent under any of the preceding glosses on the statute. Our design speaks to the following VRA questions:

1. Do white and minority (black or Latino⁵) voters polarize along racial lines in contests where one of the leading candidates is the sort of candidate whom courts would treat as a presumptive minority “candidate of choice”? (This is the threshold question of racially polarized voting.)
2. Are black and Latino voters jointly politically cohesive in such contests? (This goes to the viability of so-called “coalitional” claims, brought jointly by two or more minority groups [Hopkins 2012; Browning and Tabb 2002].)
3. Do minority-race candidates receive less support from voters not of the candidate’s race than otherwise similar white candidates, and if so, is this effect “large” relative to the effects of other candidate attributes which may be favored by some voters and disfavored by others? (This goes to disparate treatment on the basis of race, and to the question of whether race functions as an ordinary or extraordinary factor in vote choice.)
4. Do voters treat minority candidates’ race as a proxy for Democratic partisanship, or does race have effects on vote choice over and above whatever it may signal about the candidate’s partisanship? (This goes to the confounded race-or-party question.)
5. Is the penalty for minority race (if any) conditional on the type of elective office at issue, e.g., legislative vs. executive office, or “political” vs. judicial office? (This question has not received attention from courts or litigants, but is clearly germane to the VRA objective of providing minorities with an equal opportunity to participate in the political process and elect candidates of their choice.⁶)
6. Are the preceding effects (if any) geographically uniform, or are they more pronounced in some political subdivisions of the United States than in others? (This question must be answered because vote dilution claims are always brought against a state or local government, and the courts have given little weight to evidence concerning typical voting patterns in the nation as a whole. See *Grove v. Emison*, 507 U.S. 25 (1993).)

⁵In future work we hope to incorporate Asian-American candidates and voters.

⁶Also, if voters tend to discriminate more in elections for offices that are not generally elected through single-member districts (e.g., Mayor and Judge), this should motivate the development of alternative, non-district-based remedies under Section 2, such as changes in the timing of elections or interventions in the informational environment (e.g., partisan vs. nonpartisan ballots).

3.1 The Experimental Setup

We adapt and extend Hainmueller et al. (2014)’s fully randomized conjoint design for stated-preference experiments. Respondents are presented with K choice tasks. Each task requires stating a preference among J alternatives. The alternatives are constituted by L discrete-valued attributes. Each attribute in turn has D_l levels ($D_l \geq 2$), with l indexing the attribute; attribute levels are randomized with equal probability. In our design, the alternatives are candidates ($J = 2$); the attributes are Race, Education, Military Service, Endorsement, and Other Information; and the choice task is answering the question, “If you were voting in this election for [Office], which candidate do you think you would prefer?”

The screen shot in Figure 1 shows how we presented candidates to respondents. Table 1 presents the levels for each attribute.



In creating attributes and levels we were guided by the several considerations. First, we wanted to focus on estimands corresponding to the “VRA questions” enumerated above. As we will explain in more detail shortly, this dictated not only the attribute Race but also Endorsement, with placebo, racial, ideological, and partisan levels (see Table 1). The attributes Education, Military Service, and Other Information provide a point of reference for race-treatment effect sizes. Additionally, they make the purpose of our study less apparent to respondents, and they enable respondents to discriminate against minority candidates without making their discrimination obvious to the researcher.⁷ The Other Information attribute also yields a test of whether results from naturalistic vignette experiments can be replicated using conjoint designs.⁸

⁷With candidates defined by five attributes, a vote against a candidate might be due to any number of considerations other than his apparent race.

⁸Conjoint designs do not convey information in a naturalistic manner. It is possible that voter responses to the same tidbit of information may be quite different in cases where the information is acquired from a news story about a real candidate than in cases where the information concerns a hypothetical candidate and is presented in a table (Hainmueller et al. 2014, p. 27). To assess this question, we created an Other Information level with the goal of replicating the finding of a recent experiment that employed real-world news stories and actual candidates (Berinsky et al. 2011).

Berinsky et al. (2011) hypothesized that racially resentful whites would respond to sex-scandal news stories about a black male candidate by stereotyping the candidate in a “socially acceptable” way (as an extreme liberal); that this effect would be more pronounced for black candidates than for white candidates; and that the effect would dissipate if the story referred explicitly to the race of the candidate (black) and his consort (white). Using fake news stories

The following candidates are running for **City Council**:

Candidate	Robert Gonzalez 	Robert Rivers 
Military Service	No military service	Served in military.
Endorsed By	The National Association for the Advancement of Colored People (NAACP), the largest African-American civil rights and advocacy organization in the United States	The Chamber of Commerce, the largest association of business groups in the United States
Education	Graduate from an Ivy League university	Graduate from a state university (with honors)
Other Information	Volunteers with parent-teacher association	Charged with violating campaign finance laws during previous run for office

If you were voting in this election for City Council, which candidate do you think you'd vote for?

- Robert Gonzalez
- Robert Rivers
- I'd be equally likely to vote for both candidates

>>>

Figure 1: Presentation of candidate profiles to respondents

Attribute	Levels
Race	Black White Latino
Education	“Did not graduate from college” “Graduate from a state university” “Graduate from a state university (with honors)” “Graduate from an Ivy League university” “Graduate from an Ivy League university (with honors)”
Endorsement	“The National Association for the Advancement of Colored People (NAACP), the largest African-American civil rights and advocacy organization in the United States” “The National Council of La Raza (NCLR), the largest Hispanic-American civil rights and advocacy organization in the United States” “The Democratic Party” “The Republican Party” “The AFL-CIO, the largest association of labor unions in the United States” “The Chamber of Commerce, the largest association of business groups in the United States” “The Coalition for Sound Government”
Military Service	“Served in military; honored for valor in combat” “Served in military” “No military service”
Other Information	“Honored as whistle-blower for exposing corruption at a public agency” “Serves as board chairman of a not-for-profit hospital” “Was the starting quarterback for college football team” “Volunteers with parent-teacher association” “Charged with violating campaign finance laws during previous run for office” “Accused of sexual harassment by several former employees”

Table 1: Attributes and Levels

We randomize one feature of the profiles, Office, at the level of candidate pairs, varying whether the candidates are portrayed as running for Mayor, City Council, Congress, or Judge. This makes it possible to estimate variation in disparate treatment and racial polarization conditional on office. Though the vast majority of cases brought under Section 2 have concerned legislative offices, we hypothesize that African American and Latino candidates face greater opposition from white voters when they run for judicial or executive offices. Negative stereotypes about minorities' competence (Sigelman et al. 1995; Visalvanich 2014; Kinder and Dale-Riddle 2012) are likely to have a greater effect on vote choice when the candidate seeks an office that would empower him to make governmental decisions unilaterally rather than as part of a group. As well, conscious or subconscious associations between criminality and race / ethnicity (Gilliam and Iyengar 2000; Chavez 2008, 2001) may make white voters particularly reluctant to support minority candidates for positions in criminal justice administration, such as judgeships. If these conjectures are borne out by the data, they would call for a rethinking of the Voting Rights Act's present focus on legislative elections and single-member-district remedies.

In their pilot conjoint study of the effect of candidates' personal attributes on vote choice in presidential elections, Hainmueller et al. (2014) represent all attribute levels with words in a table. We follow this convention except for Race, whose levels we convey using photographs and names at the top of the screen, above the table of candidate attributes. This distinction is important.

The race-and-politics literature cautions that individuals may not reveal their actual racial preferences in response to questions that ask about them explicitly, either because respondents are unaware of racial preferences that they manifest behaviorally (Ashburn-Nardo et al. 2003; Perez 2002; Livingston 2002), because of social desirability biases (Huddy and Feldman 2009), or with accompanying photos, they found evidence corroborating these hypotheses.

We convey related information through the Other Information level, "Accused of sexual harassment by several former employees." We measure respondents' perceptions of a candidate's ideology in the final head-to-head matchup (after eliciting vote choice), and we measure explicit racial attitudes in a survey conducted several weeks after the conjoint study. We can therefore establish whether "accused of sexual harassment," when presented in a conjoint table, has average and conditional effects on perceived candidate ideology similar to those found in a leading naturalistic experiment.

because explicit racial prejudice questions fail to tap contemporary symbolic racism (Kinder and Sanders (1996); Sears and Henry (2005)). Conjoint designs, like list experiments, can mitigate the problem of social desirability bias in that they obtain information about socially sensitive preferences without requiring respondents to state those preferences expressly, but they cannot shed light on subconscious or implicit discrimination on the basis of candidates' racial appearance unless appearance itself is experimentally manipulated.⁹

Experimental studies of race effects often use stock photographs precoded for similarity in terms of likability, attractiveness, and the like, to ensure that the treatment effect of the minority likeness is really a race effect, rather than a byproduct of some other feature of the picture. Yet judgments of the likability, attractiveness, or competence of the person in a photograph may well be endogenous to the coder's racial stereotypes (Kinder and Sanders 1996; Sears and Henry 2005), as coders cannot report on such attributes without also seeing the race or apparent race of the person in the picture. Insofar as coders regard minority-race candidates negatively, designing experimental conditions for balance on perceived competence or likeability is tantamount to comparing mediocre white candidates with extraordinary minority candidates. The risk of bias is severe.

A good solution to this problem would be to photograph persons drawn at random from the pool of "potential candidates," defined as persons who possess certain objective attributes that are common among candidates for public office. Because this population is essentially unobservable, we use a proxy: currently serving state legislators. This is a conservative solution; since, to the extent that minority candidates are stereotyped as less likable, attractive, competent, etc., than their white counterparts, minority state legislators are likely to be more visually appealing relative to the pool of potential minority candidates than are white state legislators relative to the pool of white potential candidates.

We hired mTurk workers to precode male state legislator photographs for age, race / ethnicity,

⁹Political scientists generally regard photographs as the most implicit of racial signals (Mendelberg, 2001) and race-correlated names presumably function as implicit signals as well as it seems entirely innocuous for a researcher to identify a candidate by name when asking about him or her.

and quality as a campaign photo.¹⁰ We asked for campaign-photo quality ratings as a way to minimize the risk of imbalance with respect to non-racial signals about professionalism, seriousness, or competence communicated by the photographs.¹¹ Each image was rated by an average of 37 mTurk workers. On the basis of the codings we excluded very young and very old legislators, as well as legislators who were not clearly identified as black, Latino, or white. We then used the GenMatch package in R to find black and white matches for the remaining Latinos, matching on perceived age and photo quality and winnowing the pool down to eight well-matched blacks, whites, and Latinos (24 photos in all). QQ plots and other balance statistics are provided in the Supplemental Information.

To further signal race / ethnicity, we portray candidates with names as well as photographs. We matched the surnames of current state legislators to the 1,000 most common surnames in the United States, excluded non-matches, and selected from the remaining surnames the eight most racially unique white, black, and Latino names.¹² We then created three pools of presumptively racial/ethnic first names, selecting for each pool the first names of all state legislators whose surnames matched one of the eight we had chosen for the racial group in question.¹³ We randomly assigned first names to surnames (within racial groups), subject to the constraint that no first-name / surname pair match an actual legislator's name.

3.2 Estimands

Before discussing our estimands we introduce a moderate amount of notation. Throughout this paper we consider two-candidate contests. As a convention, we will call one of the candidates the

¹⁰Photos were downloaded from Project VoteSmart's website. We limit our study to male legislators because we do not want gender to confound our results, and because the greater number of male state legislators makes it easier to obtain matches across racial groups. More on this below.

¹¹Like judgments of likeability or attractiveness, judgments of photo quality are made by coders who are aware of the apparent race of the person in the photo, so there is some risk of bias from coders' racial stereotypes. But we think the risk of bias is less severe when the coder is asked to evaluate an attribute of the picture itself, rather than an attribute of the person in the picture.

¹²We used ethnic uniqueness ratings derived from the 2000 census by [cite].

¹³As the first name, we used the "preferred name" field in the Project VoteSmart database.

*focus candidate*¹⁴ and the other candidate the *challenger candidate*. Let

$$Y_{ik} = \begin{cases} 1 & \text{if respondent } i \text{ voted for the focus candidate in choice task } k \\ 0 & \text{if respondent } i \text{ did not vote for the focus candidate in choice task } k. \end{cases}$$

In much of what follows, we will assume that we can pool information across (at least some subset of) choice tasks. As a result we will typically drop the k subscript. Further, in much of what follows we will be interested not in a particular individual i but in a randomly selected individual, possibly from some well-defined demographic subgroup (Black voters, Hispanic/Latino voters, voters who self-identify with the Republican Party, etc.). We will thus oftentimes drop the i subscript.

While there are many voter attributes in our study, for present purposes we only need to introduce notation for voter race / ethnicity. We let V_i denote voter i 's race / ethnicity. V_i can take one of three values in our study: b (black), h (Hispanic/Latino), and w (white). As noted above, we will oftentimes consider a randomly chosen voter and thus drop the i subscript on V .

For present purposes we are primarily interested in two of the randomly assigned candidate attributes: focus candidate race / ethnicity, and the focus candidate's endorsement. Let C denote the focus candidate's race. C can take three values: b (black), h (Hispanic/Latino), and w (white). Further, let E denote the endorsement that the focus candidate received. While E takes seven levels in our experiment we only need to introduce notation for five of those levels. For reasons

¹⁴Note that under assumption 2 of Hainmueller et al. (2014) we can, in essence, treat each of the two candidates in a particular choice task as a focus candidate for purposes of estimation.

that will become more apparent below we chose to label these levels as:

$$E = \begin{cases} p & \text{if Coalition for Sound Government (placebo)} \\ b & \text{if NAACP} \\ h & \text{if NCLR} \\ d & \text{if Democratic Party} \\ r & \text{if Republican Party} \\ \vdots & \end{cases}$$

3.2.1 Disparate Treatment on the Basis of Race

Our measure of disparate treatment of minority candidates is straightforward: the average marginal effect of switching focus candidate Race from level c' to level c , conditional on the race of the voter. Formally, we can write the estimand as:

$$\begin{aligned} \delta_{cc'|v} &\equiv \mathbb{E}[Y(C = c)|V = v] - \mathbb{E}[Y(C = c')|V = v] \\ &= \mathbb{E}[Y|C = c, V = v] - \mathbb{E}[Y|C = c', V = v] \end{aligned}$$

Some words of explanation are in order here. The notation $Y(C = c)$ denotes a *potential outcome*. It is the value of Y that would occur when the focus candidate race / ethnicity is set to level c for a randomly chosen respondent and choice task. Note that the expectations are taken over both the joint randomization distribution of all other focus candidate attributes and all challenger candidate attributes as well as the distribution of respondents who have race / ethnicity v . Finally, note that the second line follows from the first because of the random assignment of candidate attributes and the no interference assumptions mentioned above. Sample estimates of the population expectations in the second line can be calculated in a variety of relatively straightforward ways such as by linear regression or by simply calculating sample means of Y within the appropriate

subgroups (Hainmueller et al. 2014).¹⁵

$\delta_{cc'|v}$ is an example of an *Average Marginal Component Effect* or AMCE (Hainmueller et al. 2014). At various points below we will refer to other types of causal effects that are defined by averaging over candidate attributes and voter characteristics as AMCEs.

Disparate treatment is always relative; the question is whether candidates who appear to be of race / ethnicity c are treated better or worse than candidates of race / ethnicity c' by voters who identify as members of group v . Usually the comparison of interest is between minority and white candidates, in which case $c' = w$, but in some applications one might be interested in disparate treatment of one class of minority candidates relative to another. The candidate race attribute takes on three levels in our design (black, white, and Latino / Hispanic), so we can estimate disparate treatment of candidates of any of three groups relative to any other.

Because our design randomizes candidate traits and levels are chosen with equal probability, the focus candidate's race is uncorrelated in expectation with levels for his other attributes, and with all attribute levels of the challenger candidate. The estimand $\delta_{cc'|v}$ is therefore an unconfounded measure of disparate treatment on the basis of apparent race.

Our design also allows us to say something about whether race effects are large, not just whether they are statistically significant. We do this by contrasting the treatment effect of minority race with the treatment effect of varying the levels of other attributes. For example, we can compare the effect of minority race (relative to white race) with the effect of not serving in the military (relative to serving, or serving and being honored for valor), with the effect of not graduating from college (relative to graduating from a state college or Ivy League institution), and with the effect of having been charged with campaign finance violations or accused of sexual harassment (relative to serving on the PTA, or being quarterback of the college football team). Such comparisons provide a rough sense of whether race / ethnicity operates as an ordinary factor in vote choice or as an

¹⁵To be clear, the “population” referenced here is the population of subjects in the experiment. Whether the treatment effects generalize to national or state populations depends on the representativeness of the sample, and / or reweighting.

extraordinary trump.

3.2.2 Minority Political Cohesion and White Bloc Voting

Estimating group political cohesion is more complicated than estimating disparate treatment, as there are several plausible estimands. As noted above, the courts have inferred minority political cohesion and white bloc voting from vote shares for white and minority candidates in elections that feature a candidate whom the court regards as a presumptive minority “candidate of choice.” These are generally understood to be minority-race candidates and/or candidates strongly backed by minority elites. When two or more minority groups jointly bring a vote dilution claim, their joint political cohesion is assessed similarly—the courts ask whether voters of each plaintiff group back the other group’s “candidates of choice.”

This judicial practice motivates the candidate Race and Endorsement attributes in our design. The Endorsement attribute has seven levels, two of which correspond to endorsements from prominent Latino and African American interest groups (see Table 1). In the context of our experiment, a presumptive candidate of choice of the African American community is one who physically appears to be of black origin and who is endorsed by the National Association for the Advancement of Colored People (NAACP). Similarly, presumptive Latino candidates of choice are candidates who physically appear to be of Latino descent and who receive the National Council of La Raza (NCLR) endorsement. It would not make sense to define candidate of choice in this experimental setting by race or ethnicity alone, since our randomization scheme implies that a black or Latino candidate is no more likely than his opponent to receive the NAACP / NCLR endorsement.

Minority political cohesion and white bloc voting estimands can be defined in terms of the average marginal effect of jointly switching the level of Race (from white to black, or Latino) and the level of Endorsement (from the placebo, Coalition for Sound Government, to NAACP or NCLR), conditional on the race of the respondent. Indeed, one can build most of the core concepts of vote dilution law from an expression for the degree to which minority-endorsed candidates of

minority group c are preferred to white placebo candidates by voters of racial group c' :

$$\begin{aligned}\kappa_{cc'} &\equiv \mathbb{E}[Y(C = c, E = c)|V = c'] - \mathbb{E}[Y(C = w, E = p)|V = c'] \\ &= \mathbb{E}[Y|C = c, E = c, V = c'] - \mathbb{E}[Y|C = w, E = p, V = c']\end{aligned}$$

The expectations are taken over the joint randomization distribution of all challenger candidate attributes and all focus candidate attributes other than C and E , as well as the distribution of voter characteristics conditional on voter race / ethnicity. The second equality follows from the first because of randomization of treatment assignment and the usual non-interference assumptions.

In the expression for $\kappa_{cc'}$, c is always a minority population, whereas c' may be the white population, the minority population c , or another minority population. When $c' = c$, $\kappa_{cc'} = \kappa_{cc}$, which expresses within-group minority political cohesion, i.e., the extent to which minority voters of group c prefer a co-ethnic, co-ethnic-endorsed focus candidate to a white, placebo-endorsed focus candidate after averaging over all other focus candidate attributes, all challenger candidate attributes, and all group c voter attributes.

When $c' \notin \{c, w\}$, $\kappa_{cc'}$ captures the degree to which voters of a *different* minority group prefer minority c 's presumptive candidates of choice over white, placebo-endorsed candidates. For example, if $c = b$ and $c' = h$, then $\kappa_{cc'} = \kappa_{bh}$ represents the difference on average between Latino / Hispanic support for black, NAACP-endorsed candidates and white, placebo-endorsed candidates. Two minority groups are *mutually politically cohesive* if they each prefer the other group's presumptive candidates of choice to white, placebo-endorsees. Formally, mutual political cohesion can be expressed as an average:

$$\varsigma_{cc'} = \frac{\kappa_{cc'} + \kappa_{c'c}}{2}$$

Continuing with the black and Latino example, $\varsigma_{bh} = \varsigma_{hb} = \frac{\kappa_{bh} + \kappa_{hb}}{2}$. Blacks and Latinos are

mutually politically cohesive if, on average, blacks prefer Latino, NCLR-endorsed candidates to white placebo candidates, and Latinos prefer black, NAACP-endorsed candidates to white placebos.

Most federal courts recognize coalitional claims under the Voting Rights Act—claims brought jointly by two or more minority groups—provided that the plaintiff groups are jointly politically cohesive. Joint political cohesion exists when each plaintiff group displays within-group cohesion, and the groups are also mutually cohesive as defined above.

When $c' = w$, $\kappa_{cc'} = \kappa_{cw}$ represents the extent to which *white voters* prefer the presumptive candidates of choice of minority group c to white, placebo-endorsed candidates. When whites disfavor minority-race, minority-endorsed candidates, $\kappa_{cc'} = \kappa_{cw} < 0$. Therefore, the condition that the law describes as “white bloc voting” corresponds to large values of $(-\kappa_{cw})$.

The above estimands are formally different from the observational measures of minority political cohesion and white bloc voting that have been used in Section 2 litigation, but our estimands nonetheless capture what courts seem to be driving at. Judges usually restrict their analysis to biracial elections where one candidate is presumed to be a minority candidate of choice. Voting is deemed polarized—signifying both minority political cohesion and white bloc voting—if the estimated difference between white and plaintiff-group vote shares for the minority candidate is large.¹⁶ In contrast to our measures, current judicial practice does not make an explicit comparison between vote shares of the minority candidates of choice and vote shares that would have been realized by certain counterfactual candidates. However, since the racial polarization inquiry is intended to reveal whether the minority community has “clear political preferences that are distinct from those of the majority” (*Gomez v. City of Watsonville*, 863 F.2d 1407, 1415 [9th Cir. 1988]), it would make little sense to measure racial polarization or cohesion without considering who is the alternative to the presumptive minority candidate of choice. If a minority race, minority-endorsed candidate were running against another minority candidate with similar issue positions and experience, the

¹⁶Some courts use quantitative rules of thumb for what constitutes legally significant polarization, e.g., at least 60% or 65% support for the minority candidate among minority voters, and no more than 35% or 40% support among other voters. But most courts of appeal have declined to establish any such cutoffs (Crayton (2011)).

absence of racial bloc voting certainly would not signify that the minority community lacked “clear political preferences . . . distinct from those of the majority.” Courts presumably have an intuitive grasp of this, and the law gives judges discretion to discount voting patterns in particular elections that the judge believes are unlikely to be probative of racial polarization (Katz et al. 2005, 668-670). So in practice, the candidate opposing the minority candidate of choice in most elections that courts will credit is likely to be a white candidate who is not particularly supportive of the minority community’s concerns. $\kappa_{cc'}$ makes this intuition precise, defining the comparison candidate as a white candidate who is average in all respects except that he has no endorsement that signals issue positions or partisanship. As such, he is a placeholder for whatever the voter believes to be typical about white candidates.

However, $\kappa_{cc'}$ is not the only plausible measure of group cohesion given our conjoint design. One could also measure cohesion and polarization relative to comparison candidates who are defined by the absence of attribute levels which mark a candidate as a presumptive minority candidate of choice, rather than comparison candidates who are white empty vessels. Consider two such definitions:

$$\check{\kappa}_{cc'} = \mathbb{E} [Y|C = c, E = c, V = c'] - \mathbb{E} [Y|C \neq c, E \neq c, V = c']$$

$$\tilde{\kappa}_{cc'} = \mathbb{E} [Y|C = c, E = c, V = c'] - \mathbb{E} [Y|C = w, E \notin \{b, h\}, V = c']$$

In $\check{\kappa}_{cc'}$, the comparison candidate is in effect the average candidate who is not of or endorsed by the minority group c . In $\tilde{\kappa}_{cc'}$, the comparison is the average white candidate without a minority group’s endorsement—one who is equally likely (given our randomization) to have been endorsed by the Democratic Party, the Republican Party, the AFL-CIO, the Chamber of Commerce, or the

placebo group.

We disfavor $\tilde{\kappa}_{cc'}$ because the comparison candidate is with equal probability white or minority, and because the comparison candidate will have a minority interest group endorsement in about 1/6 of the cases. This means that the measure of racial polarization between whites and blacks, for example, would be quite sensitive to how each group evaluates Latino and NCLR-endorsed candidates. Put differently, $\tilde{\kappa}_{b,w}$ conflates black-white polarization with black-Latino cohesion and white opposition to Latino candidates. $\tilde{\kappa}_{cc'}$ does not give a clear picture of dyadic relations between the three groups considered pairwise.

$\tilde{\kappa}_{cc'}$ does not suffer from this problem because the comparison candidates are white and lack the backing of minority interest groups. But $\tilde{\kappa}_{cc'}$ is likely to be quite sensitive to other levels of the endorsement attribute, and to the manner of their randomization. If the balance of endorsements (excluding the minority interest group endorsement) tilts conservative and levels are randomized with equal probability, the average comparison candidate will be a white conservative. If the endorsements are mostly from liberal groups, the average comparison candidate will be a white liberal. Since racial minorities tend to be liberal, we would expect $\tilde{\kappa}_{cc'}$ to be considerably higher in conjoint designs with a preponderance of conservative endorsements. Relatedly, $\tilde{\kappa}_{cc'}$ is less likely than $\kappa_{cc'}$ to capture respondents' subjective perceptions of where the typical white candidate stands relative to the interests and concerns of white and minority communities. This is so because the comparison candidate in $\tilde{\kappa}_{cc'}$ usually has an ideological or partisan endorsement, which means that voters' prior beliefs about white candidates will have less of an impact on what voters believe (after reading the candidate's profile) about the values and issue positions of particular comparison candidates under $\tilde{\kappa}_{cc'}$ relative to $\kappa_{cc'}$.

One could also carry the “empty vessel” idea a step further and define cohesion in terms of support for minority candidates with the placebo endorsement relative to white candidates with the placebo endorsement, rather than using minority candidates with a co-ethnic interest-group endorsement:

$$\begin{aligned}\dot{\kappa}_{cc'} &\equiv \mathbb{E} [Y(C = c, E = p)|V = c'] - \mathbb{E} [Y(C = w, E = p)|V = c'] \\ &= \mathbb{E} [Y|C = c, E = p, V = c'] - \mathbb{E} [Y|C = w, E = p, V = c']\end{aligned}$$

This estimand has a certain logic insofar as respondents assume that minority-race candidates with the placebo endorsement have values, priorities, and issue positions similar to typical candidates of their race. On the other hand, if the name of the placebo endorser (“Coalition for Sound Government”) happens to connote that its leadership is white, respondents might suppose that minority candidates with the placebo endorsement are atypical and not fully committed to serving the minority community. Also, if political factions whose objectives conflict with those of a politically cohesive minority community have strategic incentives to field minority candidates occasionally, voters’ inferences about where a candidate stands based solely on his or her race may not capture the true extent of racial cohesion / polarization within the jurisdiction.¹⁷

Yet another option is to define racial group cohesion in terms of the average marginal effect of switching Endorsement from placebo to the pertinent minority interest group, disregarding candidate race entirely:¹⁸

$$\begin{aligned}\dot{\kappa}_{cc'} &\equiv \mathbb{E} [Y(E = c)|V = c'] - \mathbb{E} [Y(E = p)|V = c'] \\ &= \mathbb{E} [Y|E = c, V = c'] - \mathbb{E} [Y|E = p, V = c']\end{aligned}$$

However, if respondents are heterogeneous in their knowledge of the NAACP and the NCLR, or in their sense of whether these groups speak for local minority interests and preferences, this estimand may not capture the true extent of racial group cohesion / polarization in the community.

¹⁷Such incentives might arise if, for example, a subset of white voters who are ideologically aligned with the white majority value symbolic expressions of racial tolerance, or if a predominantly white political faction has reason to groom local minority candidates for (eventual) contests at other levels of government.

¹⁸Thanks to Justin Levitt for suggesting this possibility.

On balance, we recommend and will employ the estimands derived from the treatment effect of switching Race (from w to minority c) and Endorsement (from p to c , or from something other than b or h to c). These estimands account for both treatments that signal whether the focus candidate is likely to cater to the interests and preferences of the minority community, and they also correspond more closely than the others to current judicial practice in vote dilution cases.¹⁹

3.2.3 Race vs. Party

To address the judicial conjecture that minority candidates lose because they are Democrats rather than because of their race, we add political party (Democratic, Republican) and general ideological (Chamber of Commerce, AFL-CIO) levels to the Endorsement attribute. This enables us to estimate effects of candidate race conditional on the candidate having a political party endorsement.²⁰ If race is a proxy for partisanship, then the effect of race should be much stronger when the candidate has the placebo endorsement than when the candidate has a partisan endorsement. Define $\delta_{cc'|v}^{-P}$ as the conditional marginal component effect of candidate race when the focus candidate does not have a party endorsement, and $\delta_{cc'|v}^P$ as corresponding effect when he does:

$$\delta_{cc'|v}^{-P} \equiv \mathbb{E}[Y|C = c, E \notin \{d, r\}, V = v] - \mathbb{E}[Y|C = c', E \notin \{d, r\}, V = v]$$

$$\delta_{cc'|v}^P \equiv \mathbb{E}[Y|C = c, E \in \{d, r\}, V = v] - \mathbb{E}[Y|C = c', E \in \{d, r\}, V = v]$$

¹⁹A very different approach would be to infer racial group political cohesion / polarization from a conjoint study in which attributes of hypothetical *citizens* rather than hypothetical candidates are randomized, and the choice task is estimating one’s own political distance from the focus citizen. Researchers could then model perceived political distance as function of respondent and focal citizen demographics, estimate the perceived distance from a given post-stratification cell to every other cell, and calculate the average perceived distance between racial groups in a geographic unit as a population-weighted average of perceived distances between post-stratification cells. This approach would have the advantage of not being sensitive to respondents’ suppositions about what is typical of minority-race candidates or certain interest groups, but it is much farther removed from current judicial practice than the approach we pursue here. Also, it is not clear that *perceived* distance between citizens of different racial groups materially affects the quality of representation for minority communities. If citizens are actually closer than they think they are, enterprising politicians should be able to figure this out and build a winning cross-racial coalition.

²⁰Of course, a Bayesian respondent with strong priors concerning the association between a candidate’s race and his partisanship or ideology may not believe that the distribution of partisanship and ideology is really the same for the white and minority candidates in our study. But courts have been generally reluctant to treat statistical discrimination as legitimate, even if the stereotype (“prior”) on which the discriminator relies has some grounding in reality (Sigelman et al. 1995; McDermott 1998a).

The “voters treat race as a proxy for partisanship” hypothesis implies:

$$|\delta_{cc'|v}^{-P}| - |\delta_{cc'|v}^P| > 0,$$

for any partisan subset of v .

Our design also ensures that the AMCE’s of racial interest group and political party endorsements are directly comparable. In each case, the reference condition (the placebo endorsement) is the same, and the average is taken over the same distribution of other candidate characteristics.

So the race-proxies-partisanship hypothesis further implies:

$$\begin{aligned} \text{sign} \{ \mathbb{E} [Y(E = b)|V = v] - \mathbb{E} [Y(E = p)|V = v] \} &= \text{sign} \{ \mathbb{E} [Y(E = d)|V = v] - \mathbb{E} [Y(E = p)|V = v] \} \\ \text{sign} \{ \mathbb{E} [Y(E = h)|V = v] - \mathbb{E} [Y(E = p)|V = v] \} &= \text{sign} \{ \mathbb{E} [Y(E = d)|V = v] - \mathbb{E} [Y(E = p)|V = v] \} \end{aligned}$$

and

$$\begin{aligned} |\mathbb{E} [Y(E = b)|V = v] - \mathbb{E} [Y(E = p)|V = v]| &< |\mathbb{E} [Y(E = d)|V = v] - \mathbb{E} [Y(E = p)|V = v]| \\ |\mathbb{E} [Y(E = h)|V = v] - \mathbb{E} [Y(E = p)|V = v]| &< |\mathbb{E} [Y(E = d)|V = v] - \mathbb{E} [Y(E = p)|V = v]| \end{aligned}$$

for any partisan subset of v . The first equations state that the AMCEs of the racial interest group endorsements are in the same direction as the AMCE of the Democratic Party endorsement; the second equations say that the racial endorsement effects are smaller in magnitude, which follows from the hypothesis that race is a noisy signal of Democratic partisanship.²¹

It would be somewhat surprising if voters did not treat candidate race and racial interest group endorsements as a proxy for Democratic partisanship, given that blacks and Latinos back Democrats over Republicans by wide margins today. Arguably more important than the question of whether voters treat race as a proxy for partisanship is (1) whether candidate race effects nonetheless persist in the presence of the partisan endorsement, and (2) whether the effects of candidate race

²¹The race-proxies-partisanship hypothesis also implies analogous expressions about the signs (opposite) and magnitudes (larger) of the treatment effect of Republican Party endorsements, relative to the treatment effect of racial interest group endorsements.

and racial interest group endorsements are larger among voters who subscribe to racial stereotypes which the Constitution disallows as the basis for state action (Elmendorf and Spencer 2014c). To address this question, we intend to elicit racial attitudes in a survey administered several weeks after the experiments, and to incorporate racial attitudes into an individual-level model predicting the response surface of the treatments (cf. Grimmer et al. 2013).²²

3.3 Implementation

3.3.1 Website Interface and Randomization

Because standard survey software packages such as Qualtrics have limited capacity for producing custom graphics with randomized components, we paid a technology consulting firm to create a simple web-based interface linked to a database of candidate attributes.²³ We used Qualtrics to call this website, which generates sixteen fully randomized candidates and eight graphical displays of information about successive pairs (see Figure 1).

Attribute levels were fully and independently randomized.²⁴ Hainmueller et al. (2014) note that researchers may wish to constrain the randomization so as to prevent implausible combinations of attribute levels. In a design such as ours, for example, one might wish to exclude endorsements by the Republican Party where Race takes on the level “black” (because black Republicans are uncommon), or to preclude a candidate who is running for Judge from having the Education level “did not graduate from college.” We did not impose such constraints, however. We allow for black Republicans because the effect of candidate race holding political party constant has been of substantive interest to the courts. As for would-be judges who did not graduate from college, we concluded that this would serve as a useful manipulation check. If respondents failed to disfavor non-college-educated candidates for judge relative to non-college-educated candidates for other

²²We administer the racial attitudes survey several weeks after the experiment to minimize the risk that answers will be contaminated by the candidates that the respondent saw in the conjoint experiments.

²³We hired the firm Seraph Analytics.

²⁴Because we randomized candidate attributes independently, it was possible for candidates of the same levels to have one or more attributes be pitted against each other. For example, both candidates in a matchup could have the same military service, or the same endorsement.

offices, this probably indicates that they are not reading or assimilating the Office treatment.

As noted previously, our decision to randomize all attribute levels with equal probability means that a simple plug-in estimate of $\delta_{cc'|v}$ is unconfounded as a measure of disparate treatment on the basis of race. But this randomization scheme also means that we do not have much statistical power to estimate minority political cohesion and white bloc voting, because focus candidates only infrequently have the relevant combinations of levels for the Race and Endorsement attributes.

Two variables were randomized at the levels above the individual candidate. “Office” was randomized at the level of the candidate pair. And “template” was randomized at the level of the respondent. We used five templates to vary the order in which attributes are presented, permitting us to detect potential violations of the no-treatment-order-effect identification assumption (Hainmueller et al. 2014).

3.3.2 Sample

Our initial results are based on a pilot sample of mTurk workers. We fielded 50 state-specific human intelligence tasks (HITs), seeking fifty respondents from each state. Four days later, after more than a majority of the state-specific surveys had been completed, we posted additional HITs for African-American and Latino workers, seeking 500 respondents from each minority group but not imposing geographic constraints on these workers.²⁵ After analyzing the mTurk results and making any necessary adjustments to the survey and website, we will field the survey on a census-matched sample from Survey Sampling International (SSI).

Because our mTurk sample is much younger and more liberal than the citizen voting age population of the United States, these data cannot support inferences to the national electorate, let alone the electorate of any state. However, they offer some initial insights regarding our experimental set-up and may be suggestive of future results when we survey a representative sample of the citizen voting age population.

²⁵We staggered these HITs because we wanted to minimize the risk that white mTurk workers would discover the minority-oversample HITs, a discovery which might affect their responses to the survey.

3.4 Estimation

Both our measure of disparate treatment, $\delta_{cc'|v}$, and our measure of group political cohesion, $\kappa_{cc'}$, are average marginal component effects (AMCEs). Hainmueller et al. (2014) prove that under certain non-interference conditions,²⁶ least squares regression and subclassification both offer nonparametric, unbiased estimate of the AMCE of varying an attribute’s level, where the average is taken over the joint distribution of the levels for other attributes.

We rely on this insight in much of what follows. The estimates that are presented in this paper are all based on linear regressions specified in a manner consistent with the results in Hainmueller et al. (2014). Standard errors are calculated using cluster-robust standard errors where the respondent ID is used as the clustering variable.

The basic insight that the estimands of interest in this paper can be represented as differences of particular, estimable, conditional expectations also opens up the possibility of constructing state-level estimates of disparate treatment, minority political cohesion, mutual cohesion, and white bloc voting. The idea is to use a multilevel model²⁷ to estimate, say, disparate treatment for a particular combination of voter race, voter age, voter education, and voter sex, and state citizenship. These demographic-specific estimates of disparate treatment could then be weighted by the appropriate census counts and added up within voter race / ethnicity categories within each state to arrive at estimates of disparate treatment specific to various voter race / ethnicity categories within states. Using such a methodology to construct point estimates is a particular case of multilevel regression and post-stratification (MRP) and should be a relative straightforward extension of Gelman and Little (1997); Lax and Phillips (2009); Ghitza and Gelman (2013); see also Elmendorf and Spencer (2014b).

²⁶The conditions are extensions of the familiar Stable Unit Treatment Value Assumption, or SUTVA.

²⁷In principle, any method of estimating a conditional expectation could be used. Machine learning methods would appear to be another attractive choice for this part of the analysis. For example, see Grimmer et al. (2013); Green and Kern (2012).

4 Results

4.1 Manipulation Checks

Respondents appear to have understood the treatments. Following the last matchup we asked a series of questions about the candidates in the pair, concluding with a question about race and ethnicity. Respondents' perceptions of candidate race / ethnicity were correct in 94% of the observations of Latino candidates, 96% of white candidates, and 97% of black candidates.

Evidence that respondents understood and assimilated other attributes is apparent from Figure 2, which reports the average marginal component effects for each attribute, conditional on respondent party identification. Liberal endorsements had a positive impact on vote choice for Democratic voters and a negative impact for Republicans; the opposite holds true for conservative endorsements. Education, Military Service and Other Information were all treated as valence traits, affecting vote choice more or less as we expected. Subtle variations that fit with conventional wisdom are also apparent. The effect of military service, for example, is slightly greater for Republican than for Democratic and Independent respondents.

4.2 Stability of Treatment Effects

The conjoint design presupposes that effects of attribute levels on vote choice are independent of (1) where the attribute is presented in the candidate's profile, (2) the timing of the choice task, i.e., whether it comes early or late in the survey, and (3) the particular sequence of candidates the respondent has seen prior to the choice task at hand (Hainmueller et al. 2014). We find, however, that treatment effects vary with the timing of the choice task. Wald tests of the interaction between attribute levels and dummies for the choice task (matchup 1 through matchup 8) indicate that the dummies are jointly significant more often than one would expect by chance. Curious patterns come to light if one treats choice task as a continuous variable and regresses vote on levels for an attribute, choice task, and an interaction term. The negative effect of campaign finance violations

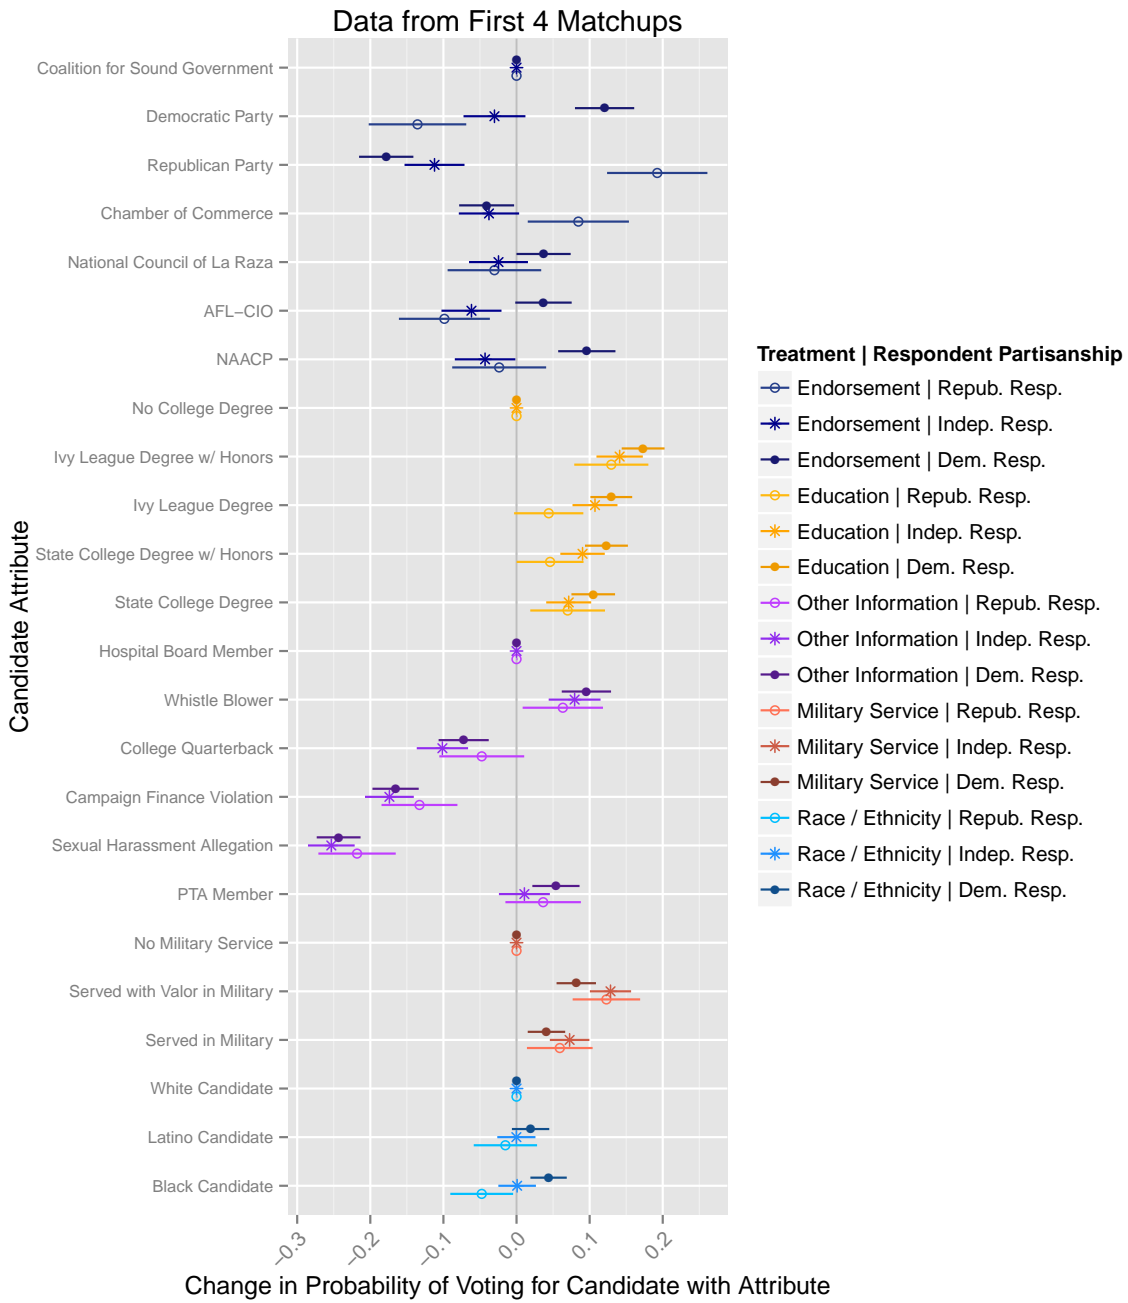


Figure 2: *Average Marginal Component Effects of Candidate Attributes By Party of Survey Respondent.* Symbols represent point estimates and horizontal lines represent 95% confidence intervals based on standard errors clustered by respondent. The effects in this figure reflect the change in probability of voting for a candidate with the attribute in question averaged over all other attributes and all competing candidate types, but conditional on the party ID of the survey respondent.

and sexual harassment increase in later choice tasks, as if respondents become fed up after seeing a number of misbehaving candidates. And the marginal benefit of serving in the military “with valor” as opposed to merely serving disappears. Serving with valor perhaps seems less extraordinary after it becomes apparent that roughly a third of the candidates under consideration have served in this manner.

In light of these timing effects, we report marginal component effects using only the first four matchups from each respondents, rather than all eight. We think this strikes a reasonable balance between the risk of bias and the benefits of greater power.

4.3 Nonparametric Estimates

This section reports nonparametric estimates of the treatment effects of substantive interest, i.e., the estimands put forth and justified in 3.1 above. Because these are, at best, inferences to the population of U.S. mTurk workers, they are not legally relevant. At most they are suggestive of what we may find when we survey a representative population.

4.3.1 Disparate Treatment

Figure 3 shows the treatment effect of candidate race / ethnicity conditional on respondent race / ethnicity. To provide a sense of magnitudes, it also presents the average marginal effect of other attributes conditional on the race of the respondent. There is essentially no evidence in this sample that white voters treat minority candidates worse than otherwise similar white candidates.²⁸ Treatment effects of candidate race occur only among co-ethnic minorities. Black voters are about 10 percentage points more likely to vote for a black candidate than a white candidate, and Latino voters prefer Latino candidates over white candidates by about the same margin. (The point estimates for the effects of the NAACP and NCRL endorsements on black and Latino voters are also about 10 percentage points.)

²⁸Likewise, Hainmueller et al. (2014) find no effect of candidate race on an mTurk sample asked to consider hypothetical candidates for President. We thought our race treatment would induce more or an effect because we used photographs and names rather than words to signal race, but it did not.




Figure 3: *Average Marginal Component Effects of Candidate Attributes By Race of Survey Respondent.* Symbols represent point estimates and horizontal lines represent 95% confidence intervals based on standard errors clustered by respondent. The effects in this figure reflect the change in probability of voting for a candidate with the attribute in question averaged over all other attributes and all competing candidate types, but conditional on the race of the survey respondent. For instance, an NAACP endorsement applied to a candidate with a random race, military service record, education, and other information, compared to an endorsement by the fictional Coalition for Sound Government applied to the same candidate will increase the probability of a black respondent voting for that candidate in a race against a randomly chosen opponent by about 0.10 (10 percentage points).

4.3.2 Minority Political Cohesion and White Bloc Voting

In this section we report estimates of within group cohesion ($\kappa_{cc'}$) as well as between minority group mutual cohesion for our full mTurk convenience sample and for respondents from two of the states with the largest numbers of respondents (including minority respondents). These states are California and Texas. Figures 4-6 present these results.

One should not put too much emphasis on the exact values of these estimates due to the non-representativeness of the mTurk sample upon which these estimates are based. Nonetheless, we do see that, even within this very young and politically liberal sample, there is some evidence of both within group cohesiveness and minority white polarization. It is also interesting to note that there is much higher black cohesiveness in Texas than in either the full sample or in the California sample. There is weak evidence that black - Latino mutual cohesion is highest in California—in large part because of the willingness of California Latino voters to vote for black candidates at higher rates than Latino voters elsewhere.

4.4 Parametric Estimates of State-Specific Average Treatment Effects

[forthcoming]

5 Conclusion

[We have no conclusions as yet.]




Figure 4: *Estimates of Political Cohesion for Full mTurk Sample.* Dots represent nonparametric point estimates and horizontal lines represent 95% confidence intervals based on standard errors clustered by respondent.




Figure 5: *Estimates of Political Cohesion for California mTurk Workers*. Dots represent nonparametric point estimates and horizontal lines represent 95% confidence intervals based on standard errors clustered by respondent.




Figure 6: *Estimates of Political Cohesion for Texas mTurk Workers.* Dots represent nonparametric point estimates and horizontal lines represent 95% confidence intervals based on standard errors clustered by respondent.

References

- Ansolabehere, Stephen, Nathaniel Persily, and Charles Stewart III. 2010. "Race, region, and vote choice in the 2008 election: Implications for the future of the Voting Rights Act." *Harvard Law Review* pp. 1385–1436.
- Ansolabehere, Stephen D, and Bernard L Fraga. 2013. "Just a Matter of Party? The Impact of Race on House Incumbent Evaluations." Working Paper.
- Arrow, Kenneth J. 1972. *Racial Discrimination in Economic Life*. D.C. Heath.
- Ashburn-Nardo, Leslie, Megan L. Knowles, and Margo J. Monteith. 2003. "Black Americans' Implicit Racial Associations And Their Implications for Intergroup Judgment." *Social Cognition* 21(3):61–87.
- Becker, Gary S. 1971. *The Economics of Discrimination*. University of Chicago Press.
- Berinsky, Adam J, Vincent L Hutchings, Tali Mendelberg, Lee Shaker, and Nicholas A Valentino. 2011. "Sex and Race: Are Black Candidates More Likely to be Disadvantaged by Sex Scandals?" *Political Behavior* 33(2):179–202.
- Branton, Regina P, Erin C Cassese, and Bradford S Jones. 2012. "Race, Ethnicity, and US House Incumbent Evaluations." *Legislative Studies Quarterly* 37(4):465–489.
- Browning, Dale R. Marhsall, Rufus, and David H Tabb. 2002. *Racial Politics in American Cities*. Pearson.
- Chavez, Leo. 2001. *Covering Immigration: Popular Images and the Politics of the Nation..* University of California Press.
- Chavez, Leo. 2008. *The Latino Threat: Constructing Immigrants, Citizens, and the Nation*. Stanford University Press.
- Citrin, Jack, Donald Philip Green, and David O Sears. 1990. "White Reactions to Black Candidates: When Does Race Matter?" *Public opinion quarterly* 54(1):74–96.
- Crayton, Kareem U. 2011. "Sword, Shield, and Compass: The Uses and Misuses of Racially Polarized Voting Studies in Voting Rights Enforcement." *Rutgers L. Rev.* 64:973.
- Druckman, James N, and Thomas J Leeper. 2012. "Learning More from Political Communication Experiments: The Importance of Pretreatment Effectsical communication experiments: pretreatment and its effects." *American Journal of Political Science* 56(4):875–896.
- Elmendorf, Christopher S. 2011. "What Kind of Discrimination Does the Voting Rights Act Target." *U. Pa. L. Rev. PENNumbra* 160:357.
- Elmendorf, Christopher S. 2012. "Making Sense of Section 2: Of Biased Votes, Unconstitutional Elections, and Common Law Statutes." *University of Pennsylvania Law Review* 160:377–455.
- Elmendorf, Christopher S, and Douglas M Spencer. 2014a. "After Shelby County: Getting Section 2 of the VRA to Do the Work of Section 5." Working Paper.
- Elmendorf, Christopher S., and Douglas M. Spencer. 2014b. "The Geography of Racial Stereotyping: Evidence and Implications for VRA Preclearance after *Shelby County*." *California Law Review* 102.

- Elmendorf, Christopher S, and Douglas M Spencer. 2014c. “The Geography of Racial Stereotyping: Evidence and Implications for VRA Preclearance After Shelby County.”
- Enos, Ryan D. 2010. “What tearing down public housing projects teaches us about the effect of racial threat on political participation.” *Unpublished*. Retrieved from [http://ryandenos.com/papers/chicago threat. pdf](http://ryandenos.com/papers/chicago%20threat.pdf) .
- Enos, Ryan D. 2011. “Testing the elusive: a field experiment on racial threat.” *Unpublished manuscript (Harvard University)* .
- Gelman, Andrew, and Thomas C. Little. 1997. “Poststratification Into Many Categories Using Hierarchical Logistica Regression.” *Survey Methodology* 23.
- Ghitza, Yair, and Andrew Gelman. 2013. “Deep interactions with MRP: Election turnout and voting patterns among small electoral subgroups.” *American Journal of Political Science* 57(3):762–776.
- Gilliam, Franklin D, and Shanto Iyengar. 2000. “Prime Suspects: The Influence of Local Television News on the Viewing Public.” *American Journal of Political Science* 3(44):243–265.
- Glynn, Adam N, and Jon Wakefield. 2010. “Ecological Inference in the Social Sciences.” *Statistical Methodology* 7(3):307–322.
- Glynn, Adam N., Jon Wakefield, Mark S. Handcock, and Thomas S. Richardson. 2008. “Alleviating Linear Ecological Bias and Optimal Design with Subsample Data.” *Journal of the Royal Statistical Society Series A* 171.
- Green, Donald P., and Holger L. Kern. 2012. “Modeling Heterogeneous Treatment Effects in Survey Experiments with Bayesian Additive Regression Trees.” *Public Opinion Quarterly* 76:491–511.
- Greiner, D. James. 2008. “Causal Inference in Civil Rights Litigation.” *Harvard Law Review* pp. 533–598.
- Greiner, D James. 2011. “Re-Solidifying Racial Bloc Voting: Empirics and Legal Doctrine in the Melting Pot.” *Indiana Law Journal* 86:447–497.
- Greiner, D James, and Kevin M Quinn. 2010. “Exit Polling and Racial Bloc Voting: Combining Individual Level and Ecological Data.” *The Annals of Applied Statistics* 4(4):1774–1796.
- Grimmer, Justin, Solomon Messing, and Sean J Westwood. 2013. “Estimating Heterogeneous Treatment Effects and the Effects of Heterogeneous Treatments with Ensemble Methods.” unpublished.
- Grofman, Bernard. 1992. “Expert Witness Testimony and the Evolution of Voting Rights Case Law.” In *Controversies in Minority Voting: The Voting Rights Act in Perspective* (Bernard Grofman, and Chandler Davidson, editors), Brookings Institution Press.
- Hainmueller, Jens, Daniel J Hopkins, and Teppei Yamamoto. 2014. “Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments.” *Political Analysis* 22(1):1–30.
- Hajnal, Zoltan. 2007. *Changing White Attitudes Toward Black Political Leadership*. Cambridge University Press.
- Haneuse, Sebastian, and Jonathan Wakefield. 2008. “The Combination of Ecological and Case-Control Data.” *Journal of the Royal Statistical Society, Series B* 70.

- Hebert, J. Gerald, Paul M. Smith, Martina E. Vandenberg, and Michael B. DeSanctis. 2010. *The Realist's Guide to Redistricting: Avoiding the Legal Pitfalls*. ABA Publishing, second edition.
- Highton, Benjamin. 2004. "White voters and African American Candidates for Congress." *Political Behavior* 26(1):1–25.
- Hopkins, Chelsea J. 2012. "Minority Coalition's Burden of Proof under Section 2 of the Voting Rights Act, The." *Santa Clara L. Rev.* 52:623–653.
- Huddy, Leonie, and Stanley Feldman. 2009. "On Assessing the Political Effects of Racial Prejudice." *Annual Review of Political Science* 12:423–447.
- Issacharoff, Samuel. 1995. "Groups and the right to vote." *Emory Law Journal* 44:869–910.
- Jackman, Simon, and Lynn Vavrek. 2011. "How Does Obama Match-Up? Counterfactuals and the Role of Obama's Race in 2008." (unpublished).
- Karl, Timothy J. Ryan, Kristyn L. 2014. "Statistical Discrimination or Prejudice? Examining When and Why Minority Candidates Pay a Penalty." Paper presented at the annual meeting of the Midwest Political Science Association.
- Karlan, Pamela S, and Daryl J Levinson. 1996. "Why Voting Is Different." *Cal L. Rev.* 84:1201–1232.
- Katz, Ellen, Margaret Aisenbrey, Anna Baldwin, and Emma Cheuse. 2005. "Documenting Discrimination in Voting: Judicial Findings Under Section 2 of the Voting Rights Act Since 1982: Final Report of the Voting Rights Initiative, University of Michigan Law School." *U. Mich. J.L. Reform* 39:643.
- Kinder, Donald R., and Allison Dale-Riddle. 2012. *The End of Race? Obama, 2008, and Racial Politics in America*. Yale University Press.
- Kinder, Donald R., and Lynn M. Sanders. 1996. *Divided by Color: Racial Politics and Democratic Ideals*. University of Chicago Press.
- King, Desmond S., and Rogers M Smith. 2011. *Still a House Divided: Race and Politics in Obama's America*. Princeton University Press.
- Lax, Jeffrey R, and Justin H Phillips. 2009. "How should we estimate public opinion in the states?" *American Journal of Political Science* 53(1):107–121.
- Livingston, Robert. 2002. "The role of perceived negativity in the moderation of African Americans' implicit and explicit racial attitudes." *Journal of Experimental Social Psychology* 28(4):405–413.
- Marisa A. Abrajano, R. Michael Alvarez, and Jonathan Nagler. 2005. "A natural experiment of race-based and issue voting: The 2001 city of Los Angeles elections." *Political Research Quarterly* 58(2):203–218.
- McDermott, Monika. 1998a. "Race and Gender Cues in Low-Information Elections." *Political Research Quarterly* 4(51):895–918.
- McDermott, Monika L. 1998b. "Race and gender cues in low-information elections." *Political Research Quarterly* 51(4):895–918.
- Mendelberg, Tali. 2001. *The race card: Campaign strategy, implicit messages, and the norm of equality*. Princeton University Press.

- Pager, Devah, and Hana Shepherd. 2008. "The Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit, and Consumer Markets." *Annual Review of Sociology* 34:181.
- Perez, Efren O. 2002. "Implicit attitudes: meaning, measurement, and synergy with political science." *Politics, Group and Identities* 2(1):275–207.
- Persily, Nathaniel. 2007. "The Promise and Pitfalls of the New Voting Rights Act." *The Yale Law Journal* 117:174–254.
- Ryan, Elizabeth M. 2009. "Causation or Correlation? The Impact of LULAC v. Clements on Section 2 Lawsuits in the Fifth Circuit." *Michigan Law Review* pp. 675–709.
- Sears, David, and PJ Henry. 2005. "Over Thirty Years Later: A Contemporary Look at Symbolic Racism." In *Advances in Experimental Psychology* (Mark P. Zanna, editor), Elsevier Academic Press.
- Sigelman, Carol K., Lee Sigelman, Barbara J. Walkosz, and Michael Nitz. 1995. "Black Candidates, White Voters: Understanding Racial Bias in Political Perceptions." *American Journal of Political Science* 1(39):243–265.
- Tesler, Michael. 2013. "The Return of Old-Fashioned Racism to White Americans Partisan Preferences in the Early Obama Era." *The Journal of Politics* 75(01):110–123.
- Tessler, Michael, and David O Sears. 2010. *Obama's Race: The 2008 Election and the Dream of a Post-Racial America*. University of Chicago Press.
- Valentino, Nicholas A, and David O Sears. 2005. "Old times there are not forgotten: Race and partisan realignment in the contemporary South." *American Journal of Political Science* 49(3):672–688.
- Visalvanich, Neil. 2014. "When Does Race Matter? Exploring White Responses to Minority Candidates in a "Post-Racial" Era."