Voter Discrimination in Democratic Elections

Suzanne K. Barth Nikolas Mittag Kyung H. Park*
Wellesley College CERGE-EI Wellesley College

October 25, 2016

Abstract

The study of voter discrimination is complicated by the possibility that voters spurn minority candidates due to unobserved candidate characteristics besides race. This paper exploits low-level statewide elections in which voters are plausibly ill-informed about candidates but can still infer race via the informational content in their names. Using nearly two decades of election results from the state of Texas, we find considerable evidence of minority disadvantage in democratic elections. Voter bias affects *both* vote share and the selection of minority candidates.

JEL Codes: J15, D72, P47

*Corresponding email: *kyung.park@wellesley.edu*. We would like to thank Kristin Butcher, Kirabo Jackson, Robin McKnight, seminar participants at the University of Chicago (Harris), Wellesley College for helpful comments. We gratefully acknowledge financial support from the Wellesley College SSSRP. We thank Victoria Angelova and Adam Martonosi for providing excellent research assistance.

1 Introduction

While the 2008 election of President Obama ushered in tidal waves of optimism and hope of a post-racial America, economic research exposed a potential underbelly of Obama's win. For example, Stephens-Davidowitz (2012) found a robust negative correlation between Obama's vote share and a proxy of racial prejudice.[1] However, the interpretation of these findings is not crystal clear. An ever-present challenge to identifying the effect of racial animus is that the candidate's race could be correlated with any number of factors that enter the voter's information set. In 2008 and 2012, an alternative explanation is simply that those electorates disfavored Obama due to some aspect of his policy platform rather than the color of his skin. Our first objective is to provide plausibly causal evidence on the impact of candidate ethnicity on voting behavior in democratic elections.

The second objective of our paper is to draw much needed attention to the possibility that voter discrimination can have *indirect* impact on policy outcomes via selection effects. Existing literature focuses primarily on the "demand" side of democratic elections; that is, whether the candidate's race affects voter choices (Stephens-Davidowitz (2012)) or whether the voter will turnout at all (Washington (2006)). Far less is known about how voter bias affects the "supply" of candidates even though, in equilibrium, voter discrimination could affect the types of minority candidates who decide to run. The relative inattention to candidate selection is surprising. In economics, se-

[1] In particular, the study finds that Obama lost approximately 4 percentage points in vote share, the rough equivalent of home state advantage, due to racial animus both in 2008 and 2012.

lection effects are ubiquitous and in political economy, a growing literature warns that ignoring the equilibrium response of candidates runs the risk of mis-characterizing elections (Ashworth and De Mesquita (2008), Prato and Wolton (2015)).

Our aim is to meld these two objectives into a comprehensive analysis of voter discrimination. The citizen-candidate model as in Besley and Coate (1997) anchors all of the paper's empirical work. We choose this theoretical framework as our starting point because it places symmetric emphasis on voter behavior and candidate entry. Our modest extension is to allow citizens to harbor racial animus towards minority candidates. This generates three intuitive predictions. First, when the candidate is a minority, her political party is expected to lose vote share in comparison with when both candidates are white. Second, in response, minority candidates should sort away from elections in which prejudice is more likely to operate. Third, if minority candidates run in prejudicial districts, then we should expect these candidates to be more moderate on average.[2] With these hypotheses in hand, our paper pivots to empirical analysis that tests each prediction in turn.

In our analysis of *voter* behavior, there are two key aspects of the research design. First, we exploit the fact that in the United States, elections for a wide range of political positions are held contemporaneously; for example, state-level positions such as the Railroad Commissioner, Comptroller of Public Accounts, and Commissioner of Agriculture are decided on the same day as

[2]The latter intuition is similar to Borjas and Bronars (1989) who finds that self-employed minority workers charge lower prices. In this application, minority candidates maintain competitiveness by "compensating" biased voters with favorable policy.

the nation's Presidency. While this structure reduces coordination costs, it raises the prospect that in "down-ticket" statewide elections, voters may have only sparse information on candidate characteristics. If voters spurn minority candidates for Railroad Commissioner, for example, then this may implicate racial discrimination to the extent that policy positions and other traits of prospective Railroad Commissioners are not well-known.[3]

Second, while the candidate's race is not explicitly listed on the ballot, previous literature has established that names can reliably predict one's race or ethnicity (Bertrand and Mullainathan (2004)). Our paper combines election results from the universe of political contests held in the state of Texas during 1992 to 2010 with Census Genealogy records. The genealogy records provide the probability of race or ethnicity conditional on surname; for example, roughly 92.7%, 90.81%, and 93.81% of persons in the U.S. with the surname of "Rodriguez", "Garcia", and "Hernandez" self-identify as Hispanic. Thus, for at least a subset of candidates, voters should be able to deduce the candidate's race or ethnicity with high degree of accuracy due to the informational content in names. This allows us to quantify the extent to which citizens trade off party allegiance for candidate race and ethnicity.

We find that in "down-ticket" state-level elections, the county-level vote share for the Democratic party decreases by roughly 5.1 percentage points when the Democratic candidate has a distinctively Hispanic surname and increases by 5.8 percentage points when the Republican candidate has a distinctively Hispanic surname. These estimates are highly robust to the inclusion

[3]Indeed, throughout the empirical analysis, we will present findings that support this view.

of a wide array of county-level characteristics and specifications that include year fixed effects.[4] In addition, descriptive statistics show only moderate voter roll-off and relatively low online search rates for "down-ticket" statewide candidates which undercut the narrative that these patterns are driven by the policy preferences of well-informed voters.

Two key pieces of evidence speak to potential mechanisms more directly. First, we construct a proxy of county-level racial prejudice based on Stephens-Davidowitz (2012) and show that the Hispanic disadvantage is more pronounced in counties that are predicted to have higher levels of racial animus. Second, specialization among lower level offices allows for more explicit tests of the whether the effects are driven by policy concerns; for example, in Texas, the Railroad Commission exclusively regulates the oil and gas industry. Because distinctively Hispanic names can signal policy preferences[5] or socio-economic status (Fryer Jr and Levitt (2004)) in addition to ethnicity, we may expect counties with more at stake in elections for the Railroad Commissioner to be particularly wary of candidates with Hispanic sounding names. We find no such heterogeneous effects.

How then does voter discrimination affect the selection of minority candidates? Several pieces of evidence strongly imply that voter bias plays an

[4]This is important because candidates are not randomly assigned. These results alleviate concerns that the estimates are driven by selection of minority candidates into statewide elections based on local characteristics or statewide trends.

[5]If minority candidates are partial to policies that diverge considerably from the median voter's preferences (Pande (2003), Chattopadhyay and Duflo (2004)), then minority candidates may receive lower vote share, on average, but for policy-related reasons that are orthogonal to race. To borrow language from the literature on statistical discrimination (Arrow (1972), Arrow (1973), Coate and Loury (1993)), minority candidates can underperform because race provides an informative signal regarding the candidate's preferences that are imperfectly observed.

influential role in candidate entry. We find that in local elections, there is a sharp and discontinuous increase in the probability of observing Hispanic candidates in voting blocs whose share of Hispanic residents exceeds 50%. That this discontinuity holds across different types of elections, including judicial contests, implies that the result is not an artifact of strategic gerrymandering since judicial district boundaries are seldom changed. This discontinuity is also difficult to reconcile with other models of candidate entry; for example, one in which minority candidates simply prefer to serve their own constituents. Instead, this result implies that minority candidates are acutely aware of and endogenously respond to voter preferences for candidate race.

In addition, our results are consistent with the narrative that voter bias has impact on the types of minority candidates who seek political office. Using ideology scores from Bonica (2014), we find that the average Hispanic candidate is more moderate in comparison with non-Hispanic candidates which cuts against the common view of minority candidates as more extreme. The degree of policy moderation is differentially stronger in voting blocs that are more ethnically diverse. An important caveat is that these estimates are not statistically precise and thus, we interpret them with caution. However, their direction and magnitudes are broadly consistent with the model's prediction that minority candidates who run in less favorable districts will "compensate" biased voters with more moderate policy in order to preserve electability.

Together, these findings have potentially important implications. Several studies have identified policy interventions that could conceivably alleviate the persistence of racial inequality (Neal (2006), Lundberg and Startz (1998),

Chetty et al. (2015)). However, the implementation of policy first requires the election of like-minded public officials. Voter discrimination undermines minority representation in a political process that crafts salient public policies, such as EITC, Head Start, Food Stamps, and Sentencing Guidelines. This is especially problematic in light of recent work that demonstrates the absence of racial diversity in public institutions can dramatically tilt outcomes unfavorably for minorities.[6] Given that these and other consequential programs have disproportionate impact on minorities, it seems plausible that voter discrimination could have far-reaching impact on racial inequality.

The rest of the paper is organized as follows. In the Section 2, we provide a simple theoretical framework that structures all of the empirical work that follows. In Section 3, we will introduce the data and descriptive statistics that provide more context for the informational environment in "down-ticket" statewide elections. Section 4 shows the econometric model and discusses why it produces *conservative* estimates of the effect of race. Sections 5 and 6 presents the empirical results on voter and candidate behavior, respectively. In Section 7, we interpret these results and we conclude in Section 8.

2 Theoretical Framework

We will now outline a theoretical framework that modestly extends the citizen-candidate model as in Besley and Coate (1997) in order to formalize our intuition as to how racial considerations might impact both voter and can-

[6]For example, Anwar et al. (2012) finds that black defendants are substantially less likely to be convicted when there is even one minority in the jury pool in comparison with jury pools that are all white.

didate behavior. In particular, we choose the citizen-candidate model as our starting point because it provides a simple way to conceptualize endogenous candidate entry. Recent theoretical work has forcefully argued that ignoring the candidate's response to voter behavior can lead to implications that severely unwind when equilibrium responses are incorporated (Ashworth and De Mesquita (2014), Prato and Wolton (2015)). Our priors are that voter discrimination may also have important *indirect* effects on the selection of minority candidates.

To begin, we consider a baseline model that abstracts from racial considerations. There is a policy space $\Omega_x = [-1, 1]$. Each citizen has a most preferred policy denoted as x_i^* which are distributed uniformly across Ω_x. We can think of citizens whose $x_i^* < 0$, $x_i^* > 0$, and $x_i^* = 0$ as Democrats (D) and Republicans (R), respectively. Any citizen can choose to enter an election but the cost of running is given by δ. A candidate cannot credibly commit to any platform that deviates from x_i^*, and thus, the winner implements his ideal policy which we denote as $\bar{x}$. In the event of a tie, the winner of the election is determined by a coin toss. There is no additional benefit to holding elected office apart from the right to set policy which precludes equilibrium in which both candidates share the same ideal point.

In addition, citizens differ in level of informativeness. A random fraction of citizens, η, observe candidate ideal points perfectly, whereas those who are uninformed vote based on statistical expectations. Given that $x_i^* \sim U[-1, 1]$, the expected policy positions of the Democratic and Republican candidates are $\mu_D = -\frac{1}{2}$ and $\mu_R = \frac{1}{2}$, respectively. Citizens have preferences over policy;

in particular, they prefer policy to be closer to their ideal point. If citizen i does not run for election, then his utility associated with policy $\bar{x}$ can be written as:

$$u_i(\bar{x}|x_i^*) = -|\bar{x} - x_i^*|$$

Even though the identification strategy hinges critically on voters being ill-informed in "down-ticket" statewide elections. We can think of these citizens as being uninformed in the sense that influential candidate-specific factors, such as character (Kartik and McAfee (2007)), valence (Stone and Simas (2010)), productivity (Rogoff (1990)), and incumbency (Ashworth and De Mesquita (2008)) appear nowhere in the citizen's utility function.[7] Instead, voters have preferences over only policy and the candidate's race. This is conceivably a close approximation to the true informational environment to the extent that voters lean primarily on party labels and candidate names in statewide "down-ticket" elections.[8]

Citizens have preferences over the race of the winning candidate, r, as well. The variable r is a binary variable that equals 1 if the winner of the election is

[7] In addition, the comparative static of how informativeness affects voter choices is not the central focus of this paper. Unlike Snyder Jr and Strömberg (2010) and Ferraz and Finan (2008), our research design does not use exogenous sources of variation in voter information sets for identification. In other words, our main interest is on the effects of race in a low information environment rather than a cross-partial that yields the differential effects of race in low versus high information environments. This lessens the value of parameterizing informativeness in our theoretical model. However, this does raise a question as to whether our estimates generalize to different informational environments. We will be careful to discuss this point in the interpretation of our results.

[8] As it turns out, allowing for heterogeneously informed voters does not yield any additional insights. This is because, in equilibrium, the two candidates must be located on opposite sides and equally distant from the median voter even when voters are heterogeneously informed. Otherwise, the candidate who is further away from median loses the election because informed voters prefer the closer candidate, while the votes of uninformed voters perfectly offset.

a minority and 0 otherwise. The fraction of citizens with distaste for minority officials is given by α. The α parameter is statistically independent of x_i^* such that "liberal" citizens are no more likely to be prejudiced than "conservatives". The disutility from minority officials is denoted as $\gamma > 0$ and, for simplicity, we restrict γ such that it does not vary across persons.[9]

We begin our analysis by graphically illustrating an example of a two-candidate equilibrium in which both candidates are white. Figure 1 depicts the ideal points of the Democratic and Republican candidate as μ_D and μ_R, respectively. In this configuration, no other citizen has an incentive to enter the election. Citizens with $x_i^* < -\frac{1}{2}$ will not enter because they would lose with vote share of no more than $\frac{1}{4}$, take votes away from μ_D, and thus, ensure a win for μ_R. In contrast, by not running, these citizens can save δ and play a lottery that leads to a more favorable policy outcome with probability $\frac{1}{2}$. Similar calculations show that citizens with $x_i^* \in (-\frac{1}{2}, \frac{1}{2})$ and $x_i^* \in (\frac{1}{2}, 1)$ do not have incentives to enter the election as well. In addition, it is straightforward to show that neither μ_D nor μ_R will drop out of the election as long as the cost of running is not too excessive ($\delta < \frac{1}{2}$). Since no actor has incentive to deviate, this constitutes a two-candidate equilibrium.

Now consider the comparative static in which the Democratic candidate, μ_D^M, is a minority and all else is held constant. The superscript denotes that the candidate is a minority. The presence of a minority candidate changes

[9] We acknowledge that this is a fairly strong assumption on the distribution of prejudice. Existing literature on labor market discrimination has found that different percentile points in the distribution of prejudice have profoundly disparate impact on the black-white wage gap (Becker (1971), Charles and Guryan (2008)). If similar forces operate in this context, then this assumption should push us towards finding null results.

Figure 1: Two-Candidate Equilibrium

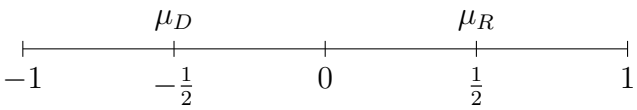

the decision calculus for citizens who harbor racial animus. Only those voters with $x_i^* < \frac{-\gamma}{2}$ will vote for μ_D^M due to racial animus. In contrast, any citizen with $x_i^* < 0$ would have voted for μ_D^M in the absence of racial prejudice. The immediate implication of this result is that the political party loses vote share when their candidate is a minority.[10] This constitutes the first prediction that we will test empirically.

While racial animus impacts voting behavior, there are additional *indirect* effects on candidate entry. In this example, the minority candidate can clearly be made better off by choosing not to run since he loses with probability 1 and running for office is costly. In other words, minority candidates should avoid seeking office in elections where the electorate is expected to have higher levels of prejudice. This constitutes the second prediction that we will examine empirically.

Finally, racial animus does not necessitate that the set of minority can-

[10]Formally, it is straightforward to show that the minority candidate receives a vote share given by:

$$P(vote for \mu_D^M) = \frac{1}{2}\left(1 - \frac{\gamma\alpha}{2}\right) < \frac{1}{2}$$

which indicates that the electoral disadvantage due to racial animus rises in γ and α. We understand that the political party could gain vote share to the extent that there is a large minority population that prefers minority candidates. However, note that if $\gamma < 0$ for a large fraction of voters, then this should push us against finding evidence of a minority disadvantage.

didates who can ever win election is an empty set. Figure 2 illustrates an example of a two-candidate equilibrium in which a minority candidate runs for office and wins with the probability of $\frac{1}{2}$. In this case, the minority candidate can win because his preferred policy is more moderate in comparison with μ_R. The median voter is now indifferent between the two candidates because even if he has distaste for minority officials, he is compensated by more favorable policy. Thus, the third prediction of the model is that *conditional on entry*, minority candidates should be more moderate, on average.

Figure 2: Two-Candidate Equilibrium with Minority Candidate

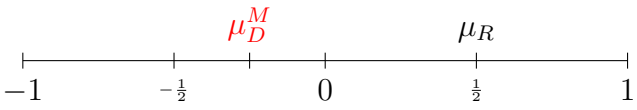

Several aspects of this framework are overly simplistic. A richer model might allow voter discrimination to be driven by uncertainty over the policy preferences of minority candidates. To the extent that minorities hold more extreme preferences, the median would eschew minority candidates not because of racial animus but because minorities are statistically more likely to implement disfavored policies. Our modeling choices should not be interpreted as a dismissal of statistical discrimination as a viable alternative explanation. In fact, our empirical work will take seriously the possibility that voting behavior towards minority candidates might reflect policy considerations rather than racial animus.

The value of this simple theoretical framework is that it focuses on our anal-

ysis around three testable implications. These are 1) when the candidate of a political party is a minority, the party is expected to lose vote share in comparison with when both candidates are white, 2) in response, minority candidates should be less likely to run in elections in which the electorate is expected to be more prejudiced, and 3) conditional on entry, minority candidate should be associated with more moderate policy preferences, on average. While these predictions are fairly intuitive, it should be noted that the literature on voter discrimination focuses disproportionately on the effects of candidate race on *voter* behavior (Stephens-Davidowitz (2012), Washington (2006)). This simple theoretical framework highlights that racial animus could have consequential effects on policy via the *selection* of political contestants as well.

3 Data and Descriptives

3.1 Data

Our primary database includes county-level vote totals for each candidate, the candidate's incumbency status, political party, and name, the number of registered voters at the county level, and the type of election for the universe of elections held during 1992 to 2010 in the state of Texas.[11] We assign each candidate to a race or ethnic group by merging information on the distinctiveness of surnames from the U.S. Census Genealogy records. We use different

[11]These data are much more comprehensive in comparison with existing political science research that studies the role of heuristics in voter choice (McDermott (1997), Bartels (1996), Matson and Fine (2006)). These studies analyze the effects of heuristics in U.S. Congressional or Presidential elections or focus on an election in a specific county.

subsets of the elections data to test the three empirical predictions generated by our theoretical framework. We examine voter behavior using only *contested* general elections for *"down-ticket" statewide offices*. The first restriction comports with the focus on two-candidate equilibrium in our theoretical model. The second is motivated by a desire to focus on elections in which candidate-specific factors are less likely to influence vote choices.

When we pivot towards the study of candidate selection, we focus the analysis on general elections held at the local level. We aggregate census tract level data from the National Historical Geographic Information System to the voting bloc level using crosswalks available at the Missouri Census Data Center. These crosswalks link census tracts to counties, Texas Senate, state House of Representative, and U.S. House of Representative districts. The measures of local voting bloc level characteristics include educational attainment, age, and racial composition among others. To examine the hypothesis that minority candidates are more moderate, on average, we use measures of candidate ideology that are computed by Bonica (2014).

3.2 Voter Informativeness

Are voters actually less informed about candidates who run in "down-ticket" statewide elections? We obtain Google Trends for politicians in both high and low level state offices. To the extent that online search behavior reflects the demand for information, this exercise should reveal whether voters are differentially informed about high and low level politicians. Panel (a) of Figure 3 shows search trends across time for John Cornyn (U.S. Senator from

Texas) and three recent Railroad Commissioners. These politicians are chosen in a way that grossly understate the disparity in search rates across offices. Including former Governors George Bush and Rick Perry would only exacerbate the difference given their high profile forays into Presidential politics. The Railroad Commissioners that we exclude held office prior to the availability of Google search rates. In spite of these choices, the plot clearly shows that the Google search rates are substantially higher for the U.S. Senator in comparison with the Railroad Commissioners.[12]

Panel (b) of Figure 3 plots the Google search trends for Greg Abbott who is an interesting case study because of his experience in both low and high level state offices. In the earlier part of his career, Abbott served as a Texas Supreme Court Justice and state's Attorney General. In July 2013, Abbott announced that he would run for the Governor of Texas which is demarcated in the figure by the vertical grey bar. Abbott would eventually win the general election by a margin of 21 percentage points and was inaugurated as Governor in January of 2015. The interesting feature of this graph is that prior to announcing his candidacy for Governor, search trends for Greg Abbott is close to zero. It is not until his general election for Governor that we observe a dramatic increase in search rates for Greg Abbott. It is rather interesting that the level of office strongly predicts the within-Greg Abbott variation in search rates over time.

[12]Specifically, Google Trends provides data on search rates as a share of the maximum search rate over the time horizon. These data are provided at the weekly level which we have aggregated to the monthly level. Google Trends allows users to compare trends across a handful of terms. This explains why our figure includes trends for a limited number of politicians. However, these politicians are chosen in a way that severely understate the differences. Finally, a value of 0 does not imply that no searches are conducted for a given search term. If the search rate does not exceed a certain threshold, Google reports the search rates as zero.

While this is only one data point, it is consistent with the view that voters are more attuned to politicians who are involved in higher level offices.

3.3 Voter Roll-off

The preceding descriptives support the notion that voters are relatively less informed "down-ticket" statewide elections. However, it is possible that voters abstain from elections in which they know little. This type of selection would imply that voters are well-informed conditional on voting. In this section, we examine the degree to which voters roll-off in low level statewide elections.

To facilitate visual presentation, we will focus on two high and two low level offices - President, U.S. Senate, Railroad Commissioner, and Comptroller of Public Accounts. Panel (a) of Figure 4 shows the statewide vote totals in general elections separately for each of the four offices by year. The striking feature of this graph is that there is much more variation in vote totals across midterm vs non-midterm elections than there is across offices within a given year. With the exception of 2000, vote totals for low level offices are surprisingly comparable to those of the Presidency and the U.S. Senate.

Panel (b) of Figure 4 shows vote totals in the elections for Railroad Commissioner in general versus primary elections. There is substantial variation in vote totals across general and primary elections. This is interesting because it seems plausible that primary voters represent a subset of voters who are relatively more informed than the average. In this case, the figure is consistent with the narrative that those who vote in "down-ticket" statewide contests in the general election are relatively less informed.

3.4 Candidate Race and Ethnicity

If voters are not be particularly well-informed about candidate-specific characteristics in "down-ticket" elections, then informational heuristics may play a more salient role in the voter's decision calculus. Importantly, voters may be able to perceive a candidate's ethnicity based on the informational content in their names. We assign a candidate's perceived race and ethnicity based on the following procedure. We begin by defining all candidates as white such that white is the default racial group. We then re-assign a candidate to the race r group if the conditional probability, $P(r|surname)$, exceeds 0.80. As an example, according to the U.S. Genealogy records, 95.93% of persons with the surname of Nguyen identify as Asian, and thus, any candidate with the surname Nguyen is categorized as Asian. The rationale is that if more than 80% of persons in the U.S. with a given surname self-identify with the race r group, then it is likely that voters will associate a candidate with this surname with the race r group as well.[13]

Table 1 shows the top 20 surnames among all candidates separately by each racial group. The names are ranked according to their joint probability, $P(r, surname)$, in order to take into account both the distinctiveness of the surname as well as its prevalence. Our approach appears successful in identifying Hispanic and Asian candidates. All of the top 20 candidates that we categorize as Hispanic or Asian have surnames that voters could conceivably

[13]Clearly, this approach is not perfect in the sense that someone whose surname is identified as Hispanic 79% may be perceived as Hispanic by voters but categorized as white by our procedure. We discuss this form of measurement error when we specify our statistical model.

perceive as sounding distinctively Hispanic or Asian.

Table 1: Candidate Surnames as a Signal of Race and Ethnicity

Rank	White	Black	Hispanic	Asian	Unmatched
1	Smith	Washington	Garcia	Nguyen	Greytok
2	Johnson	Muldrow	Rodriguez	Tran	Hinckson
3	Miller	Grays	Martinez	Chen	Kohlhausen
4	Brown	Beckles	Hernandez	Wong	Magnis
5	Harper	Winkfield	Lopez	Le	Worldpeace
6	Jones	Amadi	Gonzalez	Liu	Yokie
7	Williams		Perez	Vu	Nuchia
8	Davis		Sanchez	Cheng	Malazzo
9	Anderson		Ramirez	Vo	Naishtat
10	Wilson		Torres	Hoang	Touzel
11	Martin		Flores	Chow	Cranberg
12	Taylor		Rivera	Yoo	Sarpalius
13	Moore		Gomez	Yao	Arashvand
14	Thompson		Diaz	Yau	Dorrycott
15	White		Reyes	Hsiao	Morovich
16	Clark		Morales	Sinha	Berriozabal
17	Thomas		Cruz	Chae	Alverez
18	Hall		Ortiz	Sakai	Fastuca
19	Holm		Gutierrez	Mahajan	Markantonis
20	Baker		Chavez	Shinoda	Deotte

Note: Census Genealogy records show both 1) the prevalence of a given surname $P(surname)$ and 2) the distinctiveness of the surname $P(race|surname)$. A candidate is categorized as the race r group if $P(r|surname) > 0.80$. All other surnames are categorized as Whites. Names are ranked according to $P(r, surname)$. Among the surnames that are not matched to the Census Genealogy records, a random subset are shown in the column labeled as Unmatched.

Our approach does a poor job of identifying African-American candidates as the true number of African-American candidates is likely to exceed six. However, this may not pose serious problems for our analysis. An overview of profiles for African-American members of Congress shows that distinctively black names are relatively scarce among U.S. Congressional Representatives.[14]

[14] Fryer Jr and Levitt (2004) provide examples of distinctively black names such as Tyrone, DeShawn, Reginald, Shanice, Precious, Kiara, and Deja. Profiles of African-American

This is interesting because the prevalence of nondescript names is consistent with African-American politicians being *positively* selected since distinctively black names negatively predict socio-economic status (Fryer Jr and Levitt (2004)). Moreover, the scarcity of distinctively black first names implies that voters may also mistake African-American candidates for whites. In this case, our categorization will accurately capture voter perceptions of candidate race and ethnicity.[15]

The next step is to assess how much variation there is across racial groups. The estimated effects of race will be noisy unless minority candidates are sufficient in number. We organize the data into candidate-elected office-year cells and count the number of times we observe candidates from each racial group. Table 2 shows these counts across all general elections separately by elected office. The tabulation shows that the variation in perceived race is predominantly driven by Hispanics. While our estimates will not explicitly speak to the effects of having a black or Asian sounding name on voter choices, it seems possible that prejudice against Hispanic candidates generalizes across racial groups.

4 Empirical Model

Our theoretical model predicts that either political party will lose vote share when that party's candidate is a minority. The following specification

members of Congress are available online at history.house.gov.

[15]In addition, if voters discriminate against black candidates and we mis-classify African-American sounding names as whites, then our estimates of the vote share received by white candidates will be understated. This will have the impact of attenuating our estimates of the electoral disadvantage for Hispanic candidates.

Table 2: Counts of Candidate Racial Group by Election Type

General Elections					
High Information Elections	White	Black	Hispanic	Asian	Fraction Hispanic
President	47				0.000
U.S. Senator	43		3		0.065
U.S. Representative	834	1	124	3	0.129
Governor	22		1	1	0.042
Statewide "Low" Information Elections					
Attorney General	14		2		0.125
Lieutenant Governor	12		4		0.250
State Treasurer	2				0.000
Railroad Commissioner	27		7		0.206
Comptroller of Public Accounts	13		2		0.133
Commissioner of the General Land Office	12		3	1	0.188
Commissioner of Agriculture	17				0.000
Court of Criminal Appeals Presiding Judge	4		1		0.200
Court of Criminal Appeals Judge	56		4		0.067
Supreme Court Chief Justice	12				0.000
Supreme Court Justice	67		7		0.095
Local "Low" Information Elections					
State Senator	381		39		0.093
State Representative	2266	2	311	11	0.120
District Attorney	431		44		0.093
Criminal District Attorney	268		10		0.036
District Judge	2098	1	282	4	0.118
Family District Judge	199		22	1	0.099
Court of Appeals Chief Justice	66		7		0.096
Court of Appeals Judge	319		38	2	0.106
Member, State Board of Education	144		22		0.133
Criminal District Judge	70		5	1	0.066

Note: Census Genealogy records show both 1) the prevalence of a given surname $P(surname)$ and 2) the distinctiveness of the surname $P(race|surname)$. A candidate is categorized as the race r group if $P(r|surname) > 0.80$. All other surnames are categorized as Whites. The table shows counts of racial groups across all general election-by-year cells.

allows us to test this intuition:

$$Demvs_{cet} = \beta_0 + \beta_1 Dem * Hisp_{et} + \beta_2 Rep * Hisp_{et} + \beta_3 Dem * Inc_{et} + \beta_4 Rep * Inc_{et} + \gamma X_{ct} + \gamma_e + \eta_t + \epsilon_{cet} \quad (1)$$

The c, e, and t subscripts reflect that the election results are organized at the county, elected office (e.g. Railroad Commissioner, Comptroller of Public Accounts, and etc.), and year level. $Demvs_{cet}$ is the vote share in county c

for the Democratic candidate running for elected office e in year t. $Dem * Hisp_{et}$ and $Rep * Hisp_{et}$ are indicator variables for whether the Democratic or Republican candidate is Hispanic and $Dem * Inc_{et}$ and $Rep * Inc_{et}$ are indicator variables for whether the Democratic or Republican candidate is an incumbent, respectively. The race indicators vary only at the election-by-year level and we cluster standard errors accordingly. The γ_e and η_t reflect elected office and year fixed effects. Finally, variables in the vector X will vary across specifications but, in general, represent county-specific characteristics.

The main parameters of interest are β_1 and β_2. If voters discriminate against minority candidates, then we would expect support for the Democratic party to fall (e.g. $\hat{\beta}_1 < 0$) when the Democratic candidate is Hispanic and to rise (e.g. $\hat{\beta}_2 > 0$) when the Republican candidate is Hispanic. These capture the voter's willingness to trade off party allegiance to vote for a candidate with a less distinctively Hispanic name instead. Thus, they can be interpreted as a *marginal rate of substitution* type of parameter. In the absence of voter discrimination, we would expect both parameters to be close to zero.

The β_3 and β_4 parameters convey the extent to which incumbency status affects voter choice. In our model, we assume that voter utility does not explicitly depend on salient candidate-specific characteristics including incumbency status, and thus, we expect estimates of β_3 and β_4 to both be close to zero. It is important to note that these expectations are at odds with existing literature that finds highly robust evidence of an incumbency advantage that holds even in "low-information" elections. Conventional wisdom would predict support for the Democratic candidate to rise (e.g. $\hat{\beta}_3 > 0$) when the

Democratic candidate is an incumbent and to fall (e.g. $\hat{\beta}_4 < 0$) when the Republican candidate is an incumbent.[16] Thus, these estimates present a nice opportunity to either falsify or corroborate the key underlying assumption of our theoretical framework.

Finally, we should note that our key variables of interest $Dem * Hisp_{et}$ and $Rep * Hisp_{et}$ are plausibly measured with error. For example, a surname that is identified as Hispanic 79% of the time is categorized as white even though voters may perceive such a candidate to be Hispanic. This type of measurement error will attenuate our estimates of β_1 and β_2. To see this clearly, let us denote the voter's true perception of a Democratic and Republican candidate's ethnicity as h_D^* and h_R^*, respectively, and the measurement error as u. Because $Dem * Hisp_{et}$ and $Rep * Hisp_{et}$ are binary variables, the measurement error is mean reverting such that $cov(h_D^*, u) < 0$ and $cov(h_R^*, u) < 0$. Under the assumption that the measurement error is uncorrelated with ϵ in equation 1, then the probability limits of $\hat{\beta}_1$ and $\hat{\beta}_2$ are given by the expressions:

$$plim\hat{\beta}_1 = \beta_1 \left(\frac{var(h_D^*) + cov(h_D^*, u)}{var(h_D^*) + 2cov(h_D^*, u) + var(u)} \right)$$

$$plim\hat{\beta}_2 = \beta_2 \left(\frac{var(h_R^*) + cov(h_R^*, u)}{var(h_R^*) + 2cov(h_R^*, u) + var(u)} \right)$$

This implies that our estimates will be *conservative* as long as $-cov(h_D^*, u) <$

[16] Given that existing research finds robust evidence of incumbency advantage, it may appear as though our assumption is destined to be discredited. However, there are reasons to think otherwise. The literature has yet to examine whether the incumbency advantage differs across midterm and non-midterm in "down-ticket" statewide elections. Consistent with our descriptive statistics, our priors are that voter informativeness and hence the incumbency advantage may vary substantially across these different types of elections.

$var(u)$ and $-cov(h_R^*, u) < var(u)$.

5 Results on Voter Behavior

5.1 Main Findings

Table 3 shows the main results. Column 1 shows estimates from the baseline model that includes party-by-incumbency status indicators, elected office fixed effects, and year fixed effects. This results imply that the candidate for the Democratic Party loses 3.3 percentage points in vote share when the Democratic candidate has a Hispanic sounding surname. This point estimate is very precisely estimated achieving statistical significance at the 1% level. Given that the sample average of the county-level vote share for the Democratic party is 0.347, this constitutes a 9.8% decline in party allegiance. In addition, the Democratic party receives a 1.2 percentage point gain in vote share when the Republican party fields a candidate with a Hispanic sounding name. While this point estimate is not statistically significant at conventional levels, the direction is still consistent with voters trading off party allegiance in exchange for a candidate of the preferred racial group.

Columns 2 to 7 show point estimates from specifications that sequentially include more and more county-level characteristics. These covariates include measures of the county's racial composition, educational attainment, median income, unemployment rates, age composition, and total population. The point estimates demonstrate remarkable stability across columns. Although they appear to be identical, they are not and differ only beyond the thou-

sandths place. This robustness is to be expected. Characteristics of candidates who enter statewide elections are mechanically orthogonal to county-level characteristics. In column 8, we replace the county-level covariates with county fixed effects. It is reassuring that using either cross-county or within-county variation yields similar results.

Figure 3: Google Trend Searches by Election Type 2004-Present

(a) U.S. Senate vs Railroad Commissioner

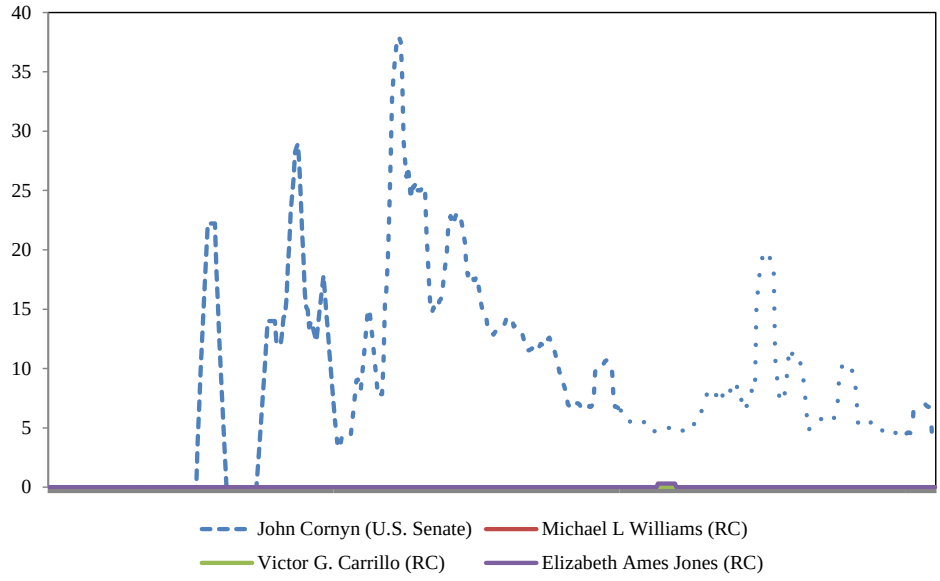

(b) Transition from Low to High Information Office

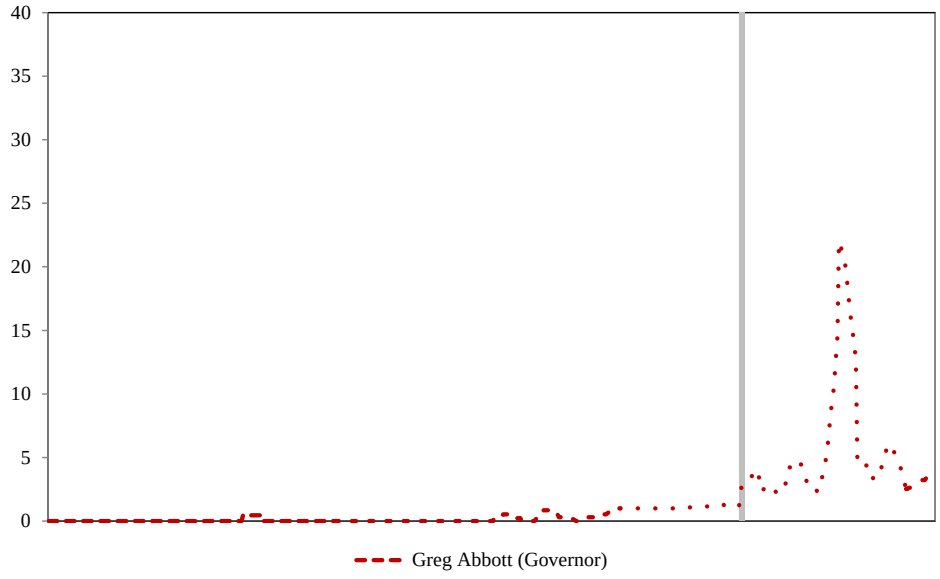

Notes: This figure graphs Google search rates as a share of the maximum search rate from 2004 to present. We average the weekly data to the month level. The vertical grey bar in Panel (b) denotes the date at which Greg Abbott announces his candidacy for state Governor.

Figure 4: Vote Totals Across Election Type

(a) Vote Totals in General Election

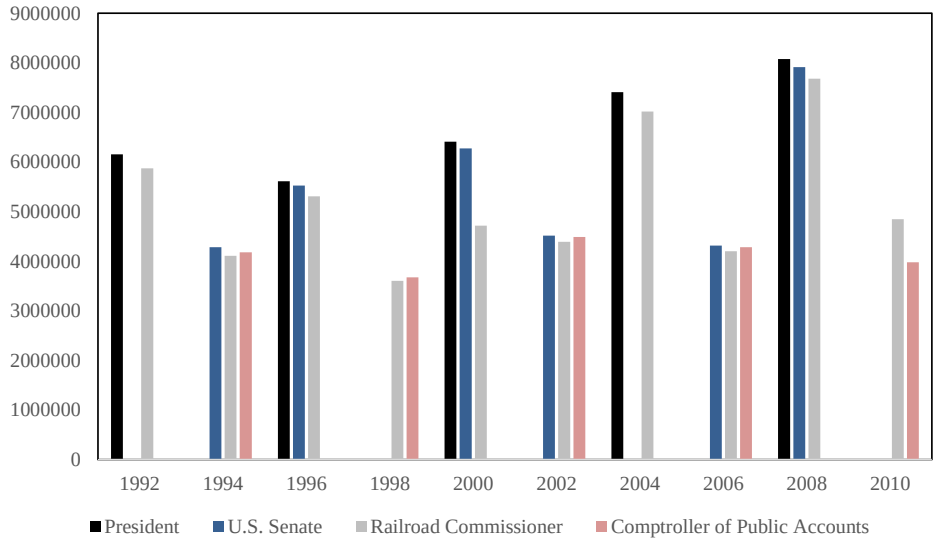

(b) Vote Total for Railroad Commissioner: General vs. Primary Elections

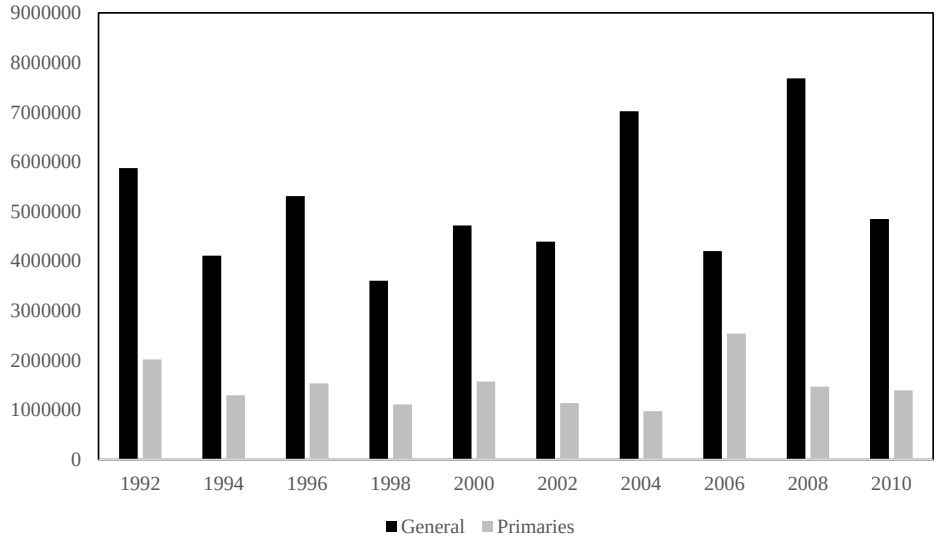

Notes: The election results are collected from the Texas Secretary of State website. In Panel (a), we show vote totals separately across these four types of elections by year. In Panel (b), we show vote totals in elections for the Railroad Commissioner across general and primary elections.

Table 3: Voting Response to Race Cues in Low-Level Statewide Elections

Dep Var: County Level Democratic Candidate Vote Share (Mean is 0.347)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Democratic*Hispanic	-0.033***	-0.033***	-0.033***	-0.033***	-0.033***	-0.033***	-0.033***	-0.033***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Republican*Hispanic	0.012	0.012	0.012	0.012	0.012	0.012	0.012	0.012
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Controls:								
Incumbent	Y	Y	Y	Y	Y	Y	Y	Y
Elected Office Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Year Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Fraction Hispanic	N	Y	Y	Y	Y	Y	Y	N
Fraction College or More	N	N	Y	Y	Y	Y	Y	N
Unemployment Rate	N	N	N	Y	Y	Y	Y	N
Median Household Income (in 2010 $)	N	N	N	N	Y	Y	Y	N
Fraction Age 65+	N	N	N	N	N	Y	Y	N
Population (in 1,000's)	N	N	N	N	N	N	Y	N
County Fixed Effects	N	N	N	N	N	N	N	Y
R-squared	0.289	0.456	0.501	0.569	0.587	0.588	0.596	0.860

Note: N = 20,065. These regressions restricts the sample to statewide low information elections which include elections for Attorney General, Lieutenant Governor, State Treasurer, Railroad Commissioner, Comptroller of Public Accounts, Commissioner of General Land Office, Commissioner of Agriculture, Court of Criminal Appeals, Supreme Court Justice. County level characteristics are computed using the 1990, 2000, and 2010 Census and intermediate years are linearly interpolated. Standard errors are clustered at the elected office-by-year level.

One curiosity is the differential effect sizes across political parties. Our theoretical framework does not predict differential effects for Democratic versus Republican Hispanic candidates. This puzzle is resolved in Table 4. Column 1 shows the main results that pool across election types whereas columns 2 and 3 show results separately for midterm and non-midterm elections, respectively. The effect of having a Hispanic sounding name is driven entirely by elections in which the nation's Presidency is at stake. In Presidential elections, the Democratic and Republican parties lose 5.1 and 5.8 percentage points, respectively, when their party's candidate is Hispanic. In contrast, in midterm elections, the estimates are much more modest and statistically insignificant.

The estimates associated with $Dem * Inc_{et}$ and $Rep * Inc_{et}$ are fascinating. First, the fact that incumbency is more advantageous for Democratic candidates is intuitive. Because Texas traditionally leans conservative, Democratic candidates who win election are more likely to be infra-marginal politicians. Second, incumbents hold an advantage only in midterm elections; in Presidential elections, the incumbency advantage vanishes completely. This result sharply cuts against conventional wisdom that suggests the incumbency advantage holds ubiquitously (Ansolabehere and Snyder Jr (2002)). Instead, this pattern is consistent with Presidential elections eliciting a subset of voters who are interested in national politics but less informed about candidates in statewide "down-ticket" elections. These results provide compelling affirmation of our research design.

Figure 5 shows the county-level distribution in residual Democratic vote

Table 4: Effects by Midterm vs. Non-Midterm Elections

Dep Var: Democratic Candidate Vote Share			
		Election Type:	
	Overall	Midterm	Non-Midterm
Candidate Race:	(1)	(2)	(3)
Democratic*Hispanic	-0.033***	-0.020	-0.051***
	(0.012)	(0.014)	(0.011)
Republican*Hispanic	0.012	-0.004	0.058**
	(0.013)	(0.011)	(0.023)
Incumbency Status:			
Democratic*Incumbent	0.054**	0.080***	-0.007
	(0.021)	(0.025)	(0.005)
Republican*Incumbent	-0.027***	-0.030***	-0.003
	(0.010)	(0.011)	(0.015)
Controls:			
Elected Office Fixed Effects	Y	Y	Y
Year Fixed Effects	Y	Y	Y
County Fixed Effects	Y	Y	Y
Observations	20,065	13,716	6,349
R-squared	0.860	0.870	0.876

Note: These regressions restricts the sample to statewide low information elections which include elections for Attorney General, Lieutenant Governor, State Treasurer, Railroad Commissioner, Comptroller of Public Accounts, Commissioner of General Land Office, Commissioner of Agriculture, Court of Criminal Appeals, Supreme Court Justice. Standard errors are clustered at the elected office-by-year level.

share separately across midterm versus non-midterm elections.[17] Panel (a) shows that in midterm elections, the distributions of county-level Democratic vote share is similar regardless of whether or not a Hispanic candidate runs for office. However, in non-midterm elections, Panel (b) shows stark shifts in the distribution depending on the candidate's ethnicity. When the Democratic

[17]The residual variation is estimated from a regression of Democratic vote share on incumbency status, the full set of county-level characteristics, elected office, and year fixed effects.

candidate is Hispanic, the distribution of county-level Democratic vote share shifts to the left signifying a loss in own-party support. When the Republican candidate is Hispanic, the distribution shifts to the right which reflects a gain in support for the opposing party. This clearly shows that our regression estimates are not driven by a handful of "outlier" counties but instead reflect broader changes in voting behavior.

Figure 5: Distribution of County-level Residualized Democratic Vote Share

(a) Midterm Elections

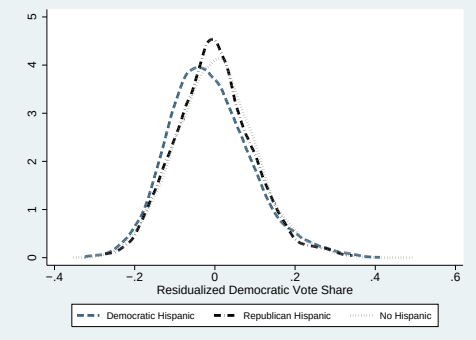

(b) Non-Midterm Elections

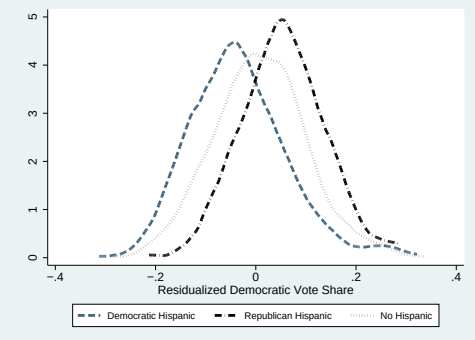

Notes: These figures plot the distribution of residual county-level Democratic vote share from regressions of Democratic vote share on incumbency status, county-level characteristics, office and year fixed effects.

Table 5 shows results in which we replace our dependent variable of county-level Democratic vote share with county-level turnout. This exercise is motivated by Washington (2006) who finds that additional black Democratic candidates on the ballot increases turnout by both black and white voters. It is interesting that we find no such evidence of increased turnout in response to Hispanic candidates in our data. All of the point estimates are close to zero. This is true across midterm and non-midterm elections as well as in low and high fraction Hispanic counties. For example, in non-midterm elections,

turnout is only 0.7 percentage points higher when the Democratic candidate is Hispanic in comparison with all white elections. While a few of the estimates are statistically significant, the magnitudes constitute negligible increases in turnout.

Table 5: Effects on County-level Voter Turnout in Low-Level Statewide Elections

Dep Var: Share of Registered Voters who Turnout (Mean is 0.443)							
		Election Type		Quartile of Fraction Hispanic			
	Overall	Midterm	Non-Midterm	Bottom	2nd	3rd	Top
Democratic*Hispanic	0.003	-0.001	0.007	0.000	0.002	0.003	0.005***
	(0.002)	(0.002)	(0.004)	(0.002)	(0.002)	(0.002)	(0.002)
Republican*Hispanic	0.001	0.004***	-0.004	-0.001	0.001	0.002	0.005**
	(0.002)	(0.002)	(0.010)	(0.003)	(0.003)	(0.002)	(0.002)
Controls:							
Incumbency	Y	Y	Y	Y	Y	Y	Y
County Fixed Effects	Y	Y	Y	Y	Y	Y	Y
Elected Office Fixed Effects	Y	Y	Y	Y	Y	Y	Y
Year Fixed Effects	Y	Y	Y	Y	Y	Y	Y
Observations	20,065	13,716	6,349	5,096	4,976	5,015	4,978
R-squared	0.885	0.829	0.878	0.912	0.912	0.883	0.832

Note: These regressions restricts the sample to low level statewide elections which include elections for Attorney General, Lieutenant Governor, State Treasurer, Railroad Commissioner, Comptroller of Public Accounts, Commissioner of General Land Office, Commissioner of Agriculture, Court of Criminal Appeals, Supreme Court Justice. The regressions include controls for whether the Democratic or Republican candidate is an incumbent, county fixed effects, office fixed effects, and year fixed effects. County level demographic data is constructed using NHGIS data. Standard errors are clustered at the elected office-by-year level.

What explains the disparity across Washington (2006) and the results of this paper? One consideration is that Washington (2006) studies the elections of state Governors, U.S. Senators, and U.S. House of Representatives whereas this analysis focuses on statewide low-level elections. If voters are ill-informed about the candidates who run for Railroad Commissioner, state Attorney General, and the Commissioner of the General Land Office, then candidate-specific attributes, such as race and ethnicity, are unlikely to influence voter decisions on turnout. However, race or ethnicity can affect vote choices to the extent

voters can infer ethnicity from the informational content in the names listed on the ballot. Thus, the absence of effects on voter turnout is also consistent with the narrative that voters are ill-informed in statewide low-level elections.

5.2 Mechanisms

Taste-based Discrimination

We now shift our attention to assess potential mechanisms. The usual suspect is taste-based discrimination. Our strategy will be to construct a proxy of racial prejudice and then examine whether or not the effects are heterogeneous across counties with different levels of predicted prejudice. We preface this analysis with two important caveats. First, we recognize the existence of cleaner tests of taste-based discrimination in the literature (Charles and Guryan (2008), Knowles et al. (2001), Anwar and Fang (2006)); however, adapting these methods to this context requires data that does not exist (to our knowledge).[18] Second, the validity of this exercise hinges critically on the reliability with which our proxy measures racial prejudice. While we will do what we can to validate our measure, we emphasize the need to interpret the following evidence with caution.

Our proxy of racial prejudice is based on Stephens-Davidowitz (2012) (hereafter SD) who uses Google search rates for racially offensive terms as a proxy for racial animus. It is reassuring that SD finds that these search rates posi-

[18] As an example, Knowles et al. (2001) conduct an innovative test for police bias that is predicated on the intuition that if police search rationally (e.g. without tastes for discrimination), then the rates of successful car searches will be the same across the racial groups. This test is not easily applied to elections since it requires measures of the elected officials' performance in order to quantify a successful election. These data are not readily available especially for statewide low-level politicians.

tively correlate with conventional measures of prejudice constructed with the General Social Survey, are more prevalent in the South, as well as among older, less educated, and less racially diverse populations. These patterns comport with prior research on the correlates of racial attitudes. A complication is that the google search rates are unavailable at the county level. Thus, we take the author's estimates from a regression of racially charged search rates on demographics and then project them onto county-level characteristics in order to construct a county-level measure of predicted racial animus.[19]

Figure 6 maps the variation in county-level predicted prejudice in Texas. The map illustrates that prejudice is expected to be higher in East Texas with pockets of high levels of prejudice in North Central Texas. This is generally consistent with the discourse found in various on-line forums that debate which areas of Texas are most racist.[20] In addition, the map displays circular dots that denote counties with known chapters of the Ku Klux Klan.[21] The map shows that counties with active KKK chapters are associated with above median levels of predicted prejudice. The exceptions are almost all in populous metropolitan areas including San Antonio, Austin, Dallas-Fort Worth, and Houston, where the KKK is less likely to be representative. In addition to the validation exercises in SD, these spatial patterns provide additional

[19] Specifically, we use the estimates in column (3) of Table 3 in SD.

[20] See for example, forums on city-data.com.

[21] This information is compiled from various sources including the United White Knights of the Ku Klux Klan website, Southern Poverty Law Center, and case law involving the KKK. As an example of the latter, *State of Texas v. Knights of Ku Klux Klan* is a case in which the KKK argues that the Texas Department of Transportation infringed on their right to free speech by denying their application to adopt a highway. The KKK's intent was to adopt a portion of a highway that passed by a housing complex mostly populated by African-Americans as a means of intimidation.

Figure 6: Spatial Patterns of Predicted Racial Animus

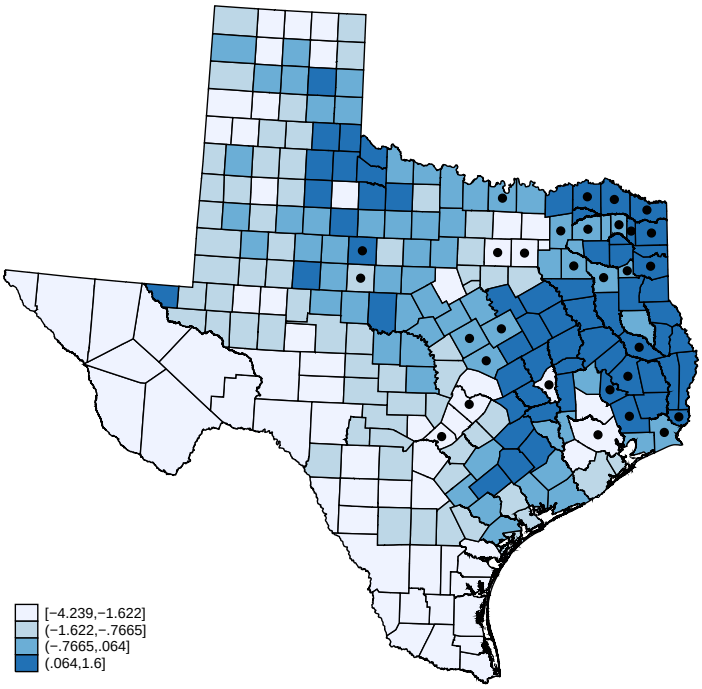

Notes: We construct a measure of predicted prejudice by taking a weighted average of the share of population older than 65, share with a bachelor's degree, fraction Hispanic, and fraction black. The weights are point estimates from regressions of Google search rates of racially insensitive terms on these covariates which are in Stephens-Davidowitz (2012). The black dots in Panel (b) are counties that are known to have active chapters of the Ku Klux Klan.

reassurance that our measure may reflect more signal than noise.

Table 6 shows results from the baseline regression model that is estimated separately for each quartile of the county-year predicted prejudice distribution. The quartiles are computed using the unweighted county-year distribution of predicted prejudice. The patterns show clear evidence of heterogeneous effects with respect to predicted racial animus. In the bottom quartile, Democratic vote share falls by 2.2 percentage points when the Democratic candidate is

Table 6: Effects by Predicted Racial Prejudice

Dep Var: Democratic Candidate Vote Share				
	Quartile of Predicted Prejudice			
	Bottom	2nd	3rd	Top
Democratic*Hispanic	-0.022**	-0.058***	-0.059***	-0.064***
	(0.009)	(0.014)	(0.010)	(0.011)
Democratic*Republican	0.018	0.078***	0.069***	0.070**
	(0.019)	(0.025)	(0.023)	(0.025)
Observations	1,599	1,575	1,600	1,575
R-squared	0.786	0.586	0.692	0.660

Note: These regressions restricts the sample to low level statewide elections which include elections for Attorney General, Lieutenant Governor, State Treasurer, Railroad Commissioner, Comptroller of Public Accounts, Commissioner of General Land Office, Commissioner of Agriculture, Court of Criminal Appeals, Supreme Court Justice. The regressions include controls for whether the Democratic or Republican candidate is an incumbent, the full set of county characteristics, office fixed effects, and year fixed effects. We construct a measure of predicted prejudice by taking a weighted average of the share of population older than 65, share with a bachelor's degree, fraction Hispanic, and fraction black. The weights are point estimates from regressions of Google search rates of racially insensitive terms on these covariates which are in Stephens-Davidowitz (2012). Standard errors are clustered at the elected office-by-year level.

Hispanic but this deficit is 2.5 to 3 three times larger in the next three quartiles. The patterns are similar when the Republican candidate is Hispanic. For example, in the bottom quartile, the Democratic party gains 1.8 percentage points in vote share when the Republican candidate is Hispanic, but in the next three quartiles the analogous estimate is roughly 4 times larger. It is important to emphasize that all of these estimates are from regressions that control for county level fraction Hispanic. Thus, ethnic composition does not account for these results even though fraction Hispanic is a strong predictor of predicted prejudice.

Statistical Discrimination

A lingering concern is that our measure of predicted prejudice correlates with unobservables that explain the Hispanic disadvantage. Perhaps these counties eschew Hispanic candidates because voters associate Hispanics with less favorable policy positions. In other words, the loss in vote share may due to political interest rather than tastes for discrimination.

An attractive feature of our research design allows us to examine this hypothesis more directly. While the Attorney General, Lieutenant Governor, and other statewide officials are responsible for a wide array of policy domains, the Railroad Commissioner heads a regulatory agency that specializes in Texas' oil and gas industry. This specialization allows us to plausibly quantify the heterogeneity in political interest associated with elections for Railroad Commissioner across counties. Specifically, counties with relatively high per-capita employees in the oil and gas sector may be more affected by the outcome of Railroad Commissioner elections than other counties. Thus, we can test the alternative hypothesis by restricting attention to Railroad Commissioner elections and examining whether the Hispanic disadvantage is differentially stronger in counties with high per-capita oil and gas employees.[22]

Panel A of Table 7 shows the effects of having a Hispanic sounding surname separately by quartile of per-capita energy sector employees. The striking feature of these estimates is their remarkable stability across quartiles. In counties with the least energy sector workers, Democratic Hispanic candidates lose 7 percentage points in vote share which is nearly identical to the 6.9 percentage point loss in counties with the most energy sector workers. The

[22]We gather data on county-level oil and gas employees from the U.S. Census Annual County and Business Patterns series.

Table 7: Elections for Railroad Commissioner

Dep Var: Democratic Candidate Vote Share				
Panel A: Effects by Political Interest				
	Quartile of Per Capita Energy Workers			
	Bottom	2nd	3rd	Top
Democratic*Hispanic	-0.070*	-0.054*	-0.058	-0.069
	(0.031)	(0.026)	(0.038)	(0.039)
Republican*Hispanic	0.048	0.038	0.051	0.051
	(0.042)	(0.035)	(0.044)	(0.048)
Observations	576	567	576	567
R-squared	0.533	0.588	0.548	0.593
Panel B: Effects by Predicted Prejudice				
	Quartile of Predicted Racial Prejudice			
	Bottom	2nd	3rd	Top
Democratic*Hispanic	-0.013	-0.080**	-0.084**	-0.093**
	(0.032)	(0.026)	(0.029)	(0.030)
Republican*Hispanic	0.009	0.045	0.053	0.059
	(0.027)	(0.034)	(0.040)	(0.058)
Observations	577	569	573	567
R-squared	0.767	0.475	0.488	0.434

Note: These regressions focus only on elections for the Railroad Commissioner. The regressions include controls for whether the Democratic or Republican candidate is an incumbent and county characteristics. We construct a measure of predicted prejudice by taking a weighted average of the share of population older than 65, share with a bachelor's degree, fraction Hispanic, and fraction black. The weights are point estimates from regressions of Google search rates of racially insensitive terms on these covariates which are in Stephens-Davidowitz (2012). Standard errors are clustered at the year level.

Republican party loses roughly 4 to 5 percentage points when their candidate is Hispanic and this effect is stable across all 4 quartiles. We note that these estimates are not precisely estimated and thus, we cannot statistically rule out that there exists a gradient across quartiles. However, the similarities in the point estimates undercut the narrative that political interest drives the main

results.

Panel B shows estimates separately for each quartile of our predicted prejudice measure. Again, we observe a stark relationship between Hispanic disadvantage and predicted racial prejudice. There is no evidence that the electorate in the least prejudiced counties vote differently when faced with candidates with Hispanic sounding surnames. In contrast, in counties with higher levels of predicted prejudice, we find that both Democratic and Republican parties lose vote share when their respective candidate is Hispanic. For example, in the top quartile, the Democratic and Republican parties lose roughly 9 and 5 percentage points, respectively, when their candidate has a Hispanic sounding surname. That we find heterogeneous effects with respect to predicted prejudice and none in relation to political interest further supports that prejudice and not policy is driving our main results.

6 Results on Candidate Selection

The results show evidence of voter bias against Hispanic candidates in democratic elections. Our theoretical framework predicts that, in equilibrium, 1) minority candidates should eschew elections in which prejudice may be more salient and 2) conditional on entry, minority candidates should be more moderate on average. We now assess these predictions by focusing on local elections but in the absence of exogenous variation in racial animus, we emphasize that the results should be interpreted as descriptive evidence.

Candidate Sorting

Panel (a) of Figure 7 graphs exposure to Hispanic candidates in "down-ticket" local elections separately for each type of elected office. These include contests for state Senate, state House of Representatives, district judges, and more. Voting bloc boundaries differ across different types of elected office. The x-axis partitions fraction Hispanic at the voting bloc level into 10 evenly spaced increments. We focus on Hispanic composition because it is observed and strongly correlates with actual racial prejudice which is arguably not as easily observed by potential candidates. The y-axis plots the voting bloc's exposure to Hispanic candidates. Exposure is defined as the fraction of elections in which a Hispanic candidate runs for office within a given voting bloc. As an example, if we observe a point "SS" with the (x,y) coordinates of (0.4,0.2), then this implies that in state Senate districts whose fraction Hispanic is between 30 and 40%, then the electorate observes a Hispanic candidate in 20% of all state Senate elections.

The notable feature of this graph is the sharp discontinuity at 0.5. In voting blocs where Hispanics constitute a minority of the population, Hispanics represent only a small fraction of the candidates overall. As soon as we cross the 50% threshold, however, there is a sizable increase in exposure to Hispanic candidates. Voting blocs whose fraction Hispanic is between 0.5 and 0.6 observe a Hispanic candidate in 35% of all elections. This is an increase of 31 percentage points in comparison with voting blocs in the left adjacent partition. From there, the likelihood of observing a Hispanic candidate steadily rises with the fraction Hispanic of the voting bloc.

A few remarks are in order. First, the increase in exposure to Hispanic can-

didates may be expected since legislators purposely re-draw district boundaries in order to increase minority representation. However, redistricting cannot fully account for these patterns. While district boundaries for state senate and state house of representative are re-drawn after each decennial census, the "one-person, one-vote" requirement does not extend to other types of office. Bloc boundaries associated with elections for district court judges, appellate court judges, criminal district attorneys, and family district judges are redrawn less frequently, and when they are, the intent is to even the "judicial burdens" (i.e. caseloads) across courts. That the same discontinuous increase holds across all offices implies that redistricting cannot explain away this result.

Second, the discontinuity is difficult to reconcile with other sensible models of candidate entry. For example, perhaps Hispanic candidates prefer to serve in majority Hispanic blocs because they can satisfy the demands of own-group constituents more efficiently. Alternatively, a positive relationship between exposure and fraction Hispanic is arguably mechanical; voting blocs with larger pools of potential minority candidates will naturally yield Hispanic candidates at higher rates. However, neither of these explanations predict a discontinuity across the 50% threshold. Instead, the discontinuous jump at 50% is highly suggestive of strategic selection into elections on the basis on racial composition. Because racial composition strongly predicts racial prejudice, it seems plausible that racial considerations are playing an important role in candidate behavior.

Panel (b) of Figure 7 characterizes the strategic entry of Hispanic candidates more fully. The figure shows exposure separately by the candidate's

Figure 7: Exposure to Hispanic Candidates in Local Elections

(a) By Elected Office

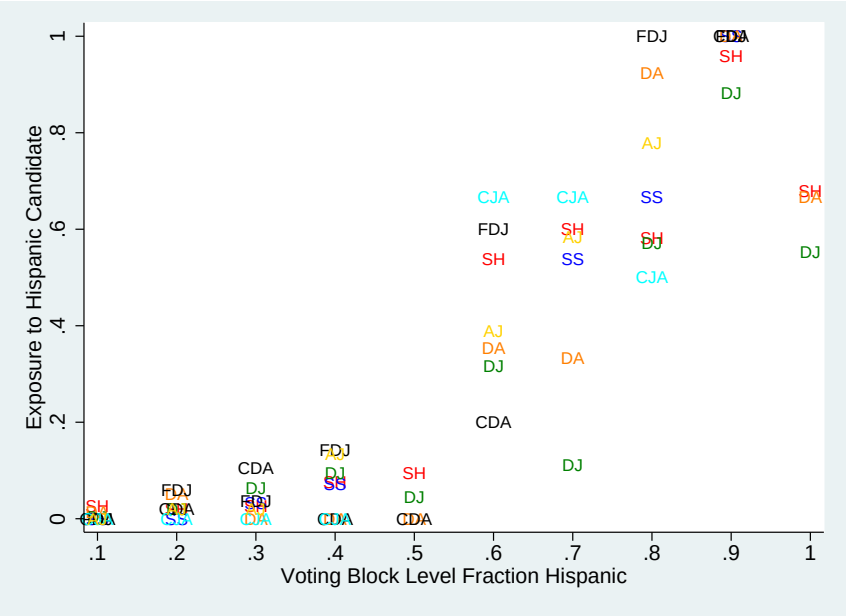

(b) By Political Party

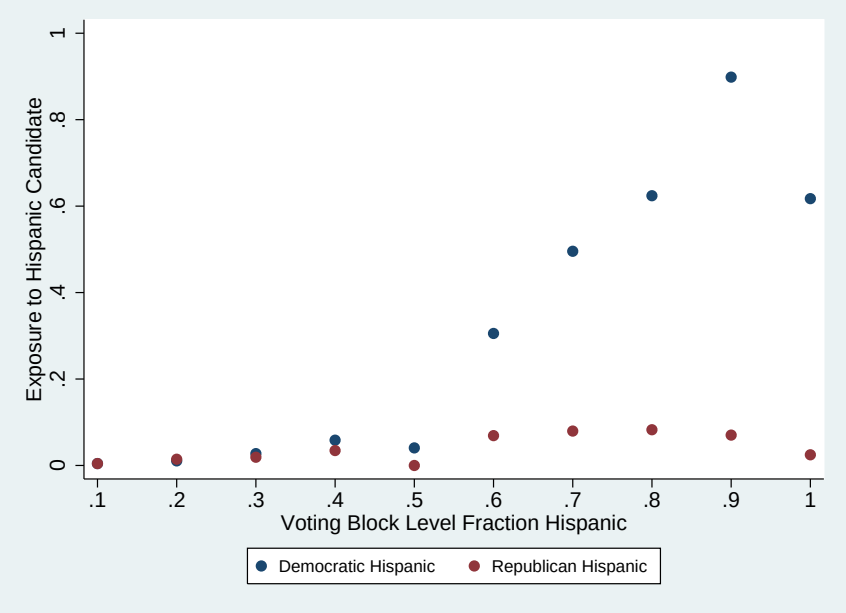

Note: This figure shows separately for different offices the fraction of election-by-year cells in which a Hispanic Democratic or Republican candidate is a candidate in the general election by the racial composition of the voting block. "SS", "SH", "DJ", "DA", "CDA", "FDJ", "CJA", and "AJ" stand for State Senate, State House of Representative, District Court Judges, District Attorney, Criminal District Attorney, Family District Judge, Chief Justice of Appellate Court, and Appellate Judge, respectively. State rep and state senate only use data from 2002 onward due to the fact redistricting and crosswalk are available only after then.

political party. The notable feature of this plot is that the pattern is driven entirely by Hispanic Democratic rather than Republican candidates. For Democratic Hispanic candidates, we see low exposure until the 50% threshold, after which exposure jumps up and then steadily increases. Exposure to Republican Hispanic candidates is essentially uniformly flat across the support. It is worth noting that this is antithetical to our model which predicts symmetric effects across parties. One possible explanation may be that the assumption that x_i^* is distributed uniformly is too rigid. Majority Hispanic districts may disproportionately lean left which would limit the potential gains to strategic selection for Republican Hispanic candidates.

Table 8 shows explicitly the effects of selection in conventional regressions of vote share on candidate ethnicity. Columns (1) and (2) show results from regressions that focus on local low level elections that exclude and then include county-level characteristics as controls. Unlike the results from statewide elections, column (1) shows that local Democratic Hispanic candidates actually gain 3.8 percentage points in vote share in comparison with whites. While Republican Hispanic candidates continue to lose vote share, this is not surprising given that Republican Hispanics exploit own-group preference less (see Figure 7). Interestingly, the effects of a Hispanic sounding name are completely explained away when we include county-level observables in column (2). As a point of comparison, columns (3) and (4) show analogous results for statewide low level elections. The main takeaway is that analysis centered on local elections would understate the role of candidate ethnicity due to endogenous candidate entry.

Table 8: Low-level Statewide vs. Local Elections

Dep Var: Democratic Candidate Vote Share				
	Non-Midterm Elections:			
	Local		Statewide	
Candidate Race:	(1)	(2)	(3)	(4)
Democratic*Hispanic	0.038***	-0.007	-0.051***	-0.051***
	(0.013)	(0.009)	(0.011)	(0.011)
Republican*Hispanic	0.052**	0.016	0.058**	0.058**
	(0.022)	(0.015)	(0.023)	(0.023)
Democratic*Incumbent	0.071***	0.061***	-0.007	-0.007
	(0.017)	(0.015)	(0.005)	(0.005)
Republican*Incumbent	-0.059***	-0.050***	-0.003	-0.003
	(0.015)	(0.013)	(0.015)	(0.015)
Controls:				
Elected Office Fixed Effects	Y	Y	Y	Y
Year Fixed Effects	Y	Y	Y	Y
County-Level Characteristics	N	Y	N	Y
Observations	476	476	6,349	6,349
R-squared	0.413	0.535	0.247	0.597

Note: Elections for state house of representatives and state senate positions prior to 2002 and elections for state board of education are excluded due to unavailability of district-level characteristics. All other local elections for district attorney, criminal district attorney, district judge, family district judge, court of appeals judge, and criminal district judge are included. Standard errors are clustered at the elected office-by-year level.

Finally, we call attention to the difference across local and statewide elections in the estimates of incumbency advantage. It is interesting that the incumbency advantage prevails only locally even though all of the elections included in this analysis are "down-ticket" contests. Several factors may contribute to this result. Perhaps the cost of running a political campaign is convex with the size of the coalition of voters required to win. In this case,

candidates of statewide offices may lack the requisite resources to mount meaningful informational campaigns. Interest groups may prefer to allocate contributions to high rather than low level statewide offices.[23] Either way, the discrepancy in incumbency advantage implies that voters are likely better informed in local versus statewide elections. This provides further rationale for our earlier focus on statewide elections to test the effects of candidate ethnicity on voting behavior.

Policy Moderation

In this section, we examine the prediction that minority candidates should be more moderate on average. For this analysis, we use measures of candidate ideology provided by Bonica (2014). This measure is available for all candidates who run for U.S. House of Representatives including those who lost election. The measure is based on the campaign contributions that a candidate receives rather than the elected politician's voting record once in office. This is relevant because the prediction is that minority candidates are more moderate *conditional on entry*.

To validate the data, Panel (a) of Figure 8 shows the Bonica (2014) ideal point estimates for high profile politicians. It is apparent that Democratic candidates locate towards the left and that Republican candidates locate towards the right of the ideological scale. Within party, we see variation across candidates that comport with individual reputations. To a first approximation, it seems plausible that these scores credibly measure candidate ideology.

Panel (b) plots the ideology scores of all U.S. House of Representative can-

[23] Alternatively, in localized voting blocs incumbents can perhaps forge beneficial political networks more easily than is possible at the state level.

Figure 8: Ideology of U.S. Congressional Representatives by Race

(a) Ideological Scores for Well-Known Officials

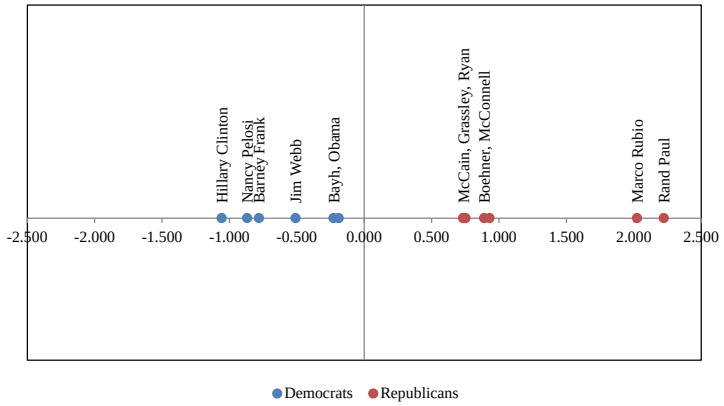

(b) Ideological Scores for U.S. Representatives from Texas by Race

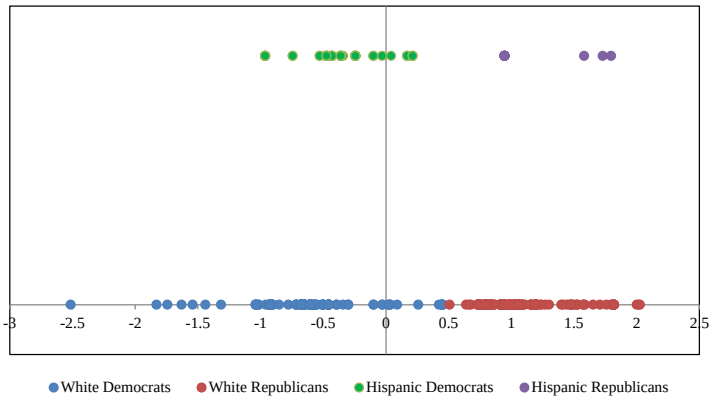

Notes: These data are from Bonica (2014). The publicly available data provides ideal point estimates for candidates in high level elections. Panel (a) plots ideal point estimates associated with well-known politicians. Panel (b) plots ideal points for Texas candidates who run for U.S. House of Representatives. Candidates are categorized as Hispanic using the Census Genealogy records.

didates from Texas during our sampling frame separately by ethnicity. Again, our focus on U.S. House of Representatives is not by choice but constrained by public availability of the data. Race is assigned using the U.S. Census Genealogy Records. The plot shows Democratic Hispanic candidates are, on average, more moderate in comparison with white Democratic candidates. It is worth noting that this is at odds with studies in the political science literature that finds voters stereotype minority candidates as having more extreme ideological views. While there are only 4 Republican Hispanics, these candidates appear to be more conservative than the average Republican.

Table 9 formalizes these observations in a regression framework. The dependent variable is Bonica (2014) measure of the candidate's ideology. Column (1) regresses this measure on an indicator for whether or not the candidate is Hispanic. The estimates imply that the average Democratic Hispanic candidate locates 0.280 units towards the right in comparison with the average white Democratic candidate whose score is -0.673. While this difference is very imprecisely estimated, the magnitude is sizable given that the distance between the average white Democrat and Republican is 0.5. In column (2), we relate the candidate's ideal point with the fraction Hispanic of the congressional district. The point estimates are consistent with the notion that Democratic Hispanic candidates are more moderate in less favorable districts.

Columns (3) and (4) show analogous results for Republican candidates. It is interesting to speculate as to why these estimates diverge from those for Democratic candidates. In particular, Republican Hispanic candidates are more conservative than the average Republican candidate even though our

Table 9: Selection Effects of the Candidate's Race

Dep Var: Estimate of Candidate's Ideal Point

	Political Party:			
	Democrats		Republicans	
Indicators for Whether:	(1)	(2)	(3)	(4)
Candidate is Hispanic	0.280		0.245	
	(0.227)		(0.211)	
District is:				
50 to 75% Hispanic		0.097		0.113
		(0.210)		(0.171)
75% or more Hispanic		-0.142		-0.000
		(0.284)		(0.171)
Constant	-0.673***	-0.626***	1.173***	1.176***
	(0.104)	(0.129)	(0.052)	(0.058)
Observations	43	43	49	49
R-squared	0.036	0.016	0.028	0.010

Note: The dependent variable is the candidate's ideology score computed by Bonica (2014). We restrict the data to U.S. House of Representative candidates from the state of Texas. Candidates are categorized as Hispanic using the Census Genealogy records.

model predicts the opposite. One possible explanation may be linked to the fact that our theoretical framework abstracts away from primary elections. If the median Republican voter is extremist, then intra-party competition between white and Hispanic Republican candidates could lead to greater policy moderation towards the party median by Republican Hispanic candidates. Thus, a model that allows the distribution of policy preferences (e.g. $f(x_i^*)$) to vary by political party and incorporates primary elections may lead to richer predictions that better align with these data.

7 Interpretation of Results

How should the reader interpret our estimates - are they local or average treatment effects? With respect to voting behavior, our results are not so different from existing literature that estimate the effects of race in high stakes election. For example, Stephens-Davidowitz (2012) estimates that President Obama lost roughly 4 percentage points in vote share due to racial animus. This is comparable to our findings that in statewide low level elections, Hispanic candidates lose roughly 5 to 6 percentage points in vote share. The consistency across studies implies that the estimates may be valid across high and low level elections within non-midterm elections. At the same time, we find disparities in the effects of race across midterm versus non-midterm elections as well as across statewide versus local low level elections. Thus, our estimates do not generalize everywhere and should not be interpreted as average treatment effects.

Instead, it seems that our estimates would be relevant for a social planner who considers various ways to structure democratic elections. For example, there is a current and sometimes contentious debate on whether we should remove party labels in judicial elections. While non-partisan elections would arguably insulate judges better from political pressure, the absence of party labels could lead voters to place more weight on racial cues. Alternatively, in light of Bertrand and Mullainathan (2004), a number of firms have changed their human resource policy to redact names from applicant resumes. One could imagine a social planner considering similar changes in which ballots

only list initials rather than full names in statewide "down-ticket" races.[24] Our results would arguably enrich these types of policy discussions by providing quantitative benchmarks as to how these changes would differentially impact minority candidates.

While our results on candidate selection are largely descriptive and statistically imprecise, they are important. Existing literature focuses on how voter bias affects voter behavior even though selection effects on candidates could have consequential impact on policy outcomes. This is especially relevant given the existing discourse on race and politics. Consider, for example, the following quote from 2015 MacArthur Grant recipient Ta-Nehisi Coates:[25]

> *But as our first black president, he has avoided mention of race almost entirely. In having to be "twice as good" and "half as black," Obama reveals the false promise and double standard of integration.*

Our results comport with but also build on the general sentiment in this statement. While it may be true that Obama has failed to meet certain expectations, we would add that the expectations may have been misplaced from the start. In equilibrium, voter bias should affect the types of minority candidates that are selected into office.

Finally, while our paper obviously has close ties to Bertrand and Mullainathan (2004), our findings further sharpen the implications of their own. In their paper, there is ambiguity as to whether or not taste-based or statistical

[24] Note that we are not taking a stance on the merits of this policy. We only note the possibility of its proposal in light of similar changes that have taken place in recent years.

[25] From an article in the September 2012 issue of *The Atlantic* entitled "Fear of a Black President".

discrimination drives their results. They conclude that, perhaps, a different model of racial animus could better fit the data. They write:

> *This discussion suggests that perhaps other models may do a better job at explaining our findings. One simple alternative model is lexicographic search by employers. Employers receive so many resumes that they may use quick heuristics in reading these resumes. One such heuristic could be to simply read no further when they see an African-American name.*

Unlike resumes, ballots provide no additional candidate-specific information which precludes voters from reading beyond names. While we leave it to future research to precisely formalize the differences across models, our findings provide further justification for the delineation of lexicographic search.

8 Conclusion

In recent survey pieces on the economics of racial discrimination (Charles and Guryan (2011), Lang and Lehmann (2012)), there is no mention of studies on voter discrimination in democratic elections. To be clear, this is *not* a criticism.[26] It does, however, reflect the fact that voter discrimination receives far less academic attention even though there is evidence that voters respond to the candidate's race (Stephens-Davidowitz (2012), Washington (2006)) and

[26] We emphasize that this is not a criticism. The papers cited in these surveys have made enormous contributions to the literature. The surveys themselves are open about their intent to review studies on labor market discrimination. Our only point is to highlight the relative scarcity of research on voter discrimination.

that diversity can substantially impact policy outcomes in ways that favor minority groups. For example, gender diversity in the federal judiciary (Boyd et al. (2010)), racial diversity among police (Anwar and Fang (2006)), and on juries (Anwar et al. (2012)) have been found to influence the adjudication of sexual harrassment cases, highway searches, and criminal conviction rates, respectively. It seems fair to say that voter discrimination probably deserves more attention than it currently receives.

We find evidence that voters discriminate against minority candidates. Using low level statewide elections, we find that both political parties lose 5 to 6 percentage points in vote share when their candidate is Hispanic. There are multiple pieces of evidence that imply these results are not driven by policy-related interests of well-informed voters. The relatively low demand for information, absence of voter roll-off, incumbency advantage, and turnout effects in "down-ticket" statewide elections support this view. In addition, we find no differential effects across counties that we might expect to have more at stake in highly specialized elections (e.g. Railroad Commissioner). Finally, we find robust evidence that the minority disadvantage is stronger in counties that are predicted to have higher levels of racial animus.

We use local elections to examine how candidates respond to voter bias. Democratic Hispanic candidates are much less likely to enter local elections in which Hispanics constitute a minority of the population. The exposure to Hispanic candidates increases sharply in majority Hispanic voting blocs. Given that racial composition strongly predict racial animus, these patterns are consistent with Democratic Hispanic candidates sorting on the basis of

racial considerations. Using measures of ideology from Bonica (2014), we also find some evidence that Democratic Hispanic candidates are more moderate, on average. There are two caveats to keep in mind in relation to these results. The estimates associated with policy moderation are statistically imprecise. In addition, these behaviors do not extend to Republican Hispanic candidates. In future research, it may be fruitful to develop richer theoretical models in order to better characterize the asymmetric candidate behavior across political parties.

Finally, it is worth addressing what appears to be a disconnect between the generality of the paper's title "Voter Discrimination" and the narrowness of the empirical results. While all of the results compare across Hispanic and white candidates, it seems plausible that the same mechanisms that affect Hispanic candidates would generalize to others. While this is purely anecdotal, the fact that two of Donald Trump's delegates named "Nabi Fakroddin" and "Raja Sadiq" lost sizable vote share in the 2016 Republican Illinois primary in comparison with other Trump delegates with more familiar names points to the possibility that the behaviors documented in this paper may extend to non-Hispanic minority candidates as well. We imagine that future work will apply this research design more broadly to study the effects on other minority groups, in other geographies, and along other salient political outcomes.

References

Ansolabehere, S. and J. M. Snyder Jr (2002). The incumbency advantage in us elections: An analysis of state and federal offices, 1942-2000. *Election Law Journal 1*(3), 315–338.

Anwar, S., P. Bayer, and R. Hjalmarsson (2012). The impact of jury race in criminal trials. *The Quarterly Journal of Economics 127*(2), 1017–1055.

Anwar, S. and H. Fang (2006). An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *The American Economic Review*.

Arrow, K. (1972). *Some Mathematical Models of Race in the Labor Market.* Lexington Books.

Arrow, K. (1973). The theory of discrimination. *Discrimination in labor markets 3*(10).

Ashworth, S. and E. B. De Mesquita (2008). Electoral selection, strategic challenger entry, and the incumbency advantage. *The Journal of Politics 70*(04), 1006–1025.

Ashworth, S. and E. B. De Mesquita (2014). Is voter competence good for voters?: Information, rationality, and democratic performance. *American Political Science Review 108*(3), 565–587.

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

Becker, G. (1971). The economics of discrimination.

Bertrand, M. and S. Mullainathan (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *The American Economic Review 94*(4), 991–1013.

Besley, T. and S. Coate (1997). An economic model of representative democracy. *The Quarterly Journal of Economics*, 85–114.

Bonica, A. (2014). Mapping the ideological marketplace. *American Journal of Political Science 58*(2), 367–386.

Borjas, G. J. and S. G. Bronars (1989). Consumer discrimination and self-employment. *Journal of Political Economy 97*(31).

Boyd, C. L., L. Epstein, and A. D. Martin (2010). Untangling the causal effects of sex on judging. *American journal of political science 54*(2), 389–411.

Charles, K. K. and J. Guryan (2008). Prejudice and wages: an empirical assessment of beckers the economics of discrimination. *Journal of Political Economy 116*(5), 773–809.

Charles, K. K. and J. Guryan (2011). Studying discrimination: Fundamental challenges and recent progress. *Annu. Rev. Econ. 3*(1), 479–511.

Chattopadhyay, R. and E. Duflo (2004). Women as policy makers: Evidence from a randomized policy experiment in india. *Econometrica 72*(5), 1409–1443.

Chetty, R., N. Hendren, and L. F. Katz (2015). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. Technical report, National Bureau of Economic Research.

Coate, S. and G. C. Loury (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 1220–1240.

Ferraz, C. and F. Finan (2008). Exposing corrupt politicians: The effects of brazil's publicly released audits on electoral outcomes. *The Quarterly Journal of Economics 123*(2), 703–745.

Fryer Jr, R. G. and S. D. Levitt (2004). The causes and consequences of distinctively black names. *The Quarterly Journal of Economics*, 767–805.

Kartik, N. and R. P. McAfee (2007). Signaling character in electoral competition. *The American economic review*, 852–870.

Knowles, J., N. Persico, and P. Todd (2001). Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy 109*(1).

Lang, K. and J.-Y. K. Lehmann (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature 50*(4), 959–1006.

Lundberg, S. and R. Startz (1998). On the persistence of racial inequality. *Journal of Labor Economics 16*(2), 292–323.

Matson, M. and T. S. Fine (2006). Gender, ethnicity, and ballot information: Ballot cues in low-information elections. *State Politics & Policy Quarterly 6*(1), 49–72.

McDermott, M. L. (1997). Voting cues in low-information elections: Candidate gender as a social information variable in contemporary united states elections. *American Journal of Political Science*, 270–283.

Neal, D. (2006). Why has black–white skill convergence stopped? *Handbook of the Economics of Education 1*, 511–576.

Pande, R. (2003). Can mandated political representation increase policy influence for disadvantaged minorities? theory and evidence from india. *The American Economic Review 93*(4), 1132–1151.

Prato, C. and S. Wolton (2015). The voters' curses: Why we need goldilocks voters. *American Journal of Political Science*.

Rogoff, K. (1990). Equilibrium political budget cycles. *The American Economic Review*, 21–36.

Snyder Jr, J. M. and D. Strömberg (2010). Press coverage and political accountability. *Journal of Political Economy 118*(2), 355–408.

Stephens-Davidowitz, S. (2012). The effects of racial animus on a black presidential candidate: Using google search data to uncover what traditional surveys miss.

Stone, W. J. and E. N. Simas (2010). Candidate valence and ideological positions in us house elections. *American Journal of Political Science 54*(2), 371–388.

Washington, E. (2006). How black candidates affect voter turnout. *The Quarterly Journal of Economics 121*(3), 973–998.