

What Do I Need to Vote?

Bureaucratic Discretion and Discrimination

by Local Election Officials*

Ariel R. White, Noah L. Nathan and Julie K. Faller[†]

This version: 11 July 2014

Abstract

Do street-level bureaucrats discriminate in the services they provide to constituents? We use a field experiment to measure differential information provision about voting by local election administrators in the United States. We contact over 7,000 election officials in 48 states who are responsible for providing information to voters and implementing voter ID laws. We find that officials provide different information to potential voters of different putative ethnicities. Emails sent from Latino aliases are significantly less likely to receive any response from local election officials than non-Latino white aliases and receive responses of lower quality. This raises concerns about the effect of voter ID laws on access to the franchise and about bias in the provision of services by local bureaucrats more generally.

*Thanks to David Broockman, David Doherty, Ryan Enos, Claudine Gay, Nahomi Ichino, Dustin Tingley, Kris-Stella Trump, the anonymous reviewers, and participants in Government 2008 and the American Politics Research Workshop at Harvard for helpful comments. An earlier version of this paper was presented at the 2013 APSA annual meeting in Chicago. David Kimball and Martha Kropf generously shared data on the partisanship of election officials, Matthew Disler, Tian Kisch, and Yena Oh provided research assistance, and Ryan Enos and Dustin Tingley provided the funding for this project. All errors are our own. This research was approved by Harvard University's Committee on the Use of Human Subjects in Research. Replication data will be made available upon publication at: dvn.iq.harvard.edu.

[†]Ph.D. Candidates, Department of Government, Harvard University. Email: jfaller@fas.harvard.edu, nlnathan@fas.harvard.edu, arwhite@fas.harvard.edu.

1 Introduction

On October 2, 2012, with the general election weeks away, a Pennsylvania judge suspended the state’s new voter identification law, which would have required Pennsylvanians to show ID at the polls. Despite this, television ads highlighting the then-defunct requirement were still broadcast after the ruling and a large Spanish-language billboard stating that ID would be required remained on a prominent Philadelphia highway for weeks. This raised concerns that a lack of information would create confusion among voters.¹ Similar issues extend beyond Pennsylvania; since 2002, twenty other US states adopted additional ID requirements for voting. Many voters – even in states without these new requirements – held inaccurate beliefs about whether they needed ID in the 2012 elections.² Amid heated debate over possible discriminatory effects of these laws, there has been concern that changes to voting requirements raise informational costs to voting, differentially reducing turnout among those with less access to information, such as the Spanish-speaking voters targeted by the Philadelphia billboard (Vercellotti and Andersen 2009).

At the front lines of the effort to inform voters about new voting requirements stand nearly 8,000 “street-level bureaucrats,” the county and municipal administrators who implement the nation’s electoral laws (Kimball and Kropf 2006). Like other street-level bureaucrats, local election administrators often have considerable discretion in how they manage the election system, operate with little direct oversight from state officials, and are frequently time and resource constrained (Burden et al. 2012, 2013).³ Scholars of bureaucracy have argued that when local bureaucrats have

¹Romy Varghese, “Confusion Sown in Pennsylvania by Lingering Voter-ID Ads,” *Bloomberg News*, 12 October 2012, <http://www.bloomberg.com/news/2012-10-12/confusion-sown-in-pennsylvania-by-lingering-voter-id-ads.html>. An earlier 2012 survey of Pennsylvanians found that 34.3% of registered voters did not know of the ID requirement and that 12% incorrectly believed themselves to have the required ID (Barreto et al. 2012).

²A 2012 Pew survey found that only 82% of registered voters (and 76% of registered Latinos) living in states that required photo IDs knew about this requirement and that 52% of registered voters (and 59% of registered Latinos) living in non-voter ID states erroneously thought they needed ID (Lopez and Gonzalez-Barrera 2012).

³Lipsky (1980) defines street-level bureaucrats as those bureaucrats who “interact directly with citizens” and “have substantial discretion in the execution” of policies that they do not make (3). Although some local election officials

discretion in allocating effort and resources, racial, ethnic, or partisan bias can occur unless clear procedures are in place to prevent bias (Jones et al. 1977, Lipsky 1980, Brodtkin 1997). Discretionary bias from bureaucrats is of particular concern in the electoral system because differential treatment of potential voters has implications for the composition of the electorate and the nature of democratic political representation.

As a wave of new voting requirements is implemented across the US, do the street-level bureaucrats managing the electoral system discriminate by race or ethnicity in interactions with constituents about voting rules? We employ a large-scale, email-based field experiment, similar in design to Butler and Broockman (2011) and McClendon (2012), to explore whether minorities face higher costs than other voters to learning about voting requirements from election officials. In doing so, we conduct the first field experiment of this scale on street-level bureaucrats in the US. We test whether the local officials directly responsible for election administration provide different information about these requirements to Latino versus non-Latino white voters. We randomly assign essentially all county or municipal election officials in 48 US states (over 7,000 officials) to receive emails from fictitious constituents with either putatively Latino or non-Latino white names asking two questions about voting before the 2012 election: a question about voter ID or a control question measuring baseline differences in responsiveness.⁴

Analyzing over 5,300 replies, we find clear, causally identified evidence of bias against Latinos in the responsiveness of local election officials. Emails from Latino names are roughly five percentage points less likely to receive a reply to a question about voter ID requirements than those from non-Latino whites. Replies that Latino emailers do receive are less likely to convey accurate information about ID requirements. There is some evidence that this bias is greater in responses

are themselves elected, they meet this classic definition and existing literature treats them as street-level bureaucrats (Burden et al. 2012, Kimball and Kropf 2006). We find no differences in behavior of elected versus appointed officials (see Section 4.4).

⁴Much debate about voter ID laws has focused on whether these laws prevent voting by undocumented immigrants or instead deter eligible Latino voters. Adding an African-American treatment condition would have substantially reduced statistical power, so we focus on Latino names for this study.

to questions about voter ID compared to a baseline question, but this is not statistically significant across model specifications.

Our results cannot be explained by differences in the resource constraints or professionalization of bureaucrats and there is little support for plausible “statistical discrimination” mechanisms, such as if the bias were driven by the partisanship of local officials. We do, however, find evidence consistent with the claim that those working in institutions with mechanisms in place aimed at preventing discrimination are less discriminatory (Pager and Shepherd 2008); despite bias against Latinos nationally across all types of officials, we find no bias in counties and municipalities covered under Sections 5 and 203 of the federal Voting Rights Act.⁵

Overall, our results likely provide a lower bound for the bias real Latinos may face in interactions with these local election officials. Lipsky (1980) argues that bias will be highest for actions that are complex and costly to the bureaucrat; replying to the short emails in our experiment is particularly low cost.⁶ Yet these emails stand in for a larger class of interactions between election officials and potential voters – from providing directions to polling places to determining voters’ registration eligibility – that can require more complicated actions and for which bias may be greater.⁷ In addition, real bias can result from officials using a bundle of characteristics to decide who is deserving of their effort; characteristics such as ethnicity, race, class, accent, and literacy can all influence discrimination (Segrest Purkiss et al. 2006, Hancock 2007, Hosoda et al. 2012). But while these characteristics are correlated within the American population, including for Latinos, our experiment only primes ethnicity.⁸ As a result, we likely provide a conservative estimate

⁵Section 5 was still fully in place during our experiment.

⁶Many accurate responses to our email about voter ID involved simply pasting a link to an informational website, while accurate responses to the control email often contained only the word “No.”

⁷We might be concerned that the impersonal and anonymous nature of email allows bureaucrats to discriminate more than they would in face-to-face interactions with prospective voters. But in the extensive private sector literature on studies of racial discrimination in hiring, Pager (2007) notes that “correspondence” studies in which resumés are sent by mail to employers often find lower estimates of discrimination than audit studies in which actors apply for jobs in person. This suggests that our manner of communication is unlikely to inflate our estimates of discrimination.

⁸In a pre-test of our experimental primes, described below, we did not find large differences in the extent to which

for the bias that could be faced by a poor, non-white, and/or non-English-speaking Latino voter in comparison to a wealthier, white, and English-speaking non-Latino voter.

We believe our experiment is strengthened by the realism of the design. Survey evidence (see footnote 2) indicates that many voters do not have accurate knowledge about ID requirements. While no existing data specifically measures rates at which voters email local officials about voter ID laws, Americans (including Latinos) are increasingly likely to contact officials like those in this study online.⁹ Most importantly, we directly study the relevant population through real individual interactions, rather than observing bureaucratic behavior in the face of numerous confounders.

Ultimately, this research makes two main contributions. To our knowledge, this is the first large-scale field experiment on bureaucratic behavior in the US offering causal evidence of discrimination on the basis of ethnicity or race. This brings field experimental evidence to bear on long-standing debates about the conditions under which street-level bureaucrats are biased in responsiveness to constituents (e.g., Jones et al. 1977, Lipsky 1980, Mladenka 1981, Goodsell 1983).¹⁰ Our experiment expands recent observational research on bias in implementation of social and welfare services (Keiser et al. 2004, Fording et al. 2007, Davis et al. 2011) to the study of election administration and more directly isolates whether individual-level patterns of bias are best explained by “taste-based,” strategic, or resource-driven factors than existing literature.

Second, our results have clear implications for states with new voter ID laws and for minority voting more generally. While evidence of aggregate turnout effects of these laws is mixed

our Latino and non-Latino names cued education level, age, or political interest.

⁹A 2009 poll found that of those Americans who reported contacting government within the last year, 60% had done so online or via email at least once (Smith et al. 2009) and the proportion of white, black, and Hispanic internet-users reporting online or email-based government interactions is similar (Smith 2010). Further, citizens regularly contact local officials, like the election administrators we email in this study. A 2007 poll found that 58% of Americans reported contacting a government office in the past year; of these, 27% said they contacted their local government, and 12% multiple levels of government, which could include local officials (Estabrook et al. 2007).

¹⁰Grose (2014) highlights the lack of existing field experimental work on bureaucratic behavior in the US. Previous scholars of election administrators have noted the difficulty of identifying the underlying factors driving the discretionary behavior of these officials through observational studies (see Burden et al. 2013, 907).

(Alvarez et al. 2008, Ansolabehere 2009, Erickson and Minnite 2009, Mycoff et al. 2009), they are often thought to have discriminatory effects because poor and minority voters are less likely to have required ID (Barreto et al. 2007, 2009) or because poll workers may challenge poor or minority voters to show ID more frequently (Atkeson et al. 2010, Cobb et al. 2012). We instead explore a previously unexamined mechanism – the provision of information before election day – through which these laws may differentially affect voters. Observational evidence indicates that negative turnout effects of voter ID laws are largest when states have recently implemented them, especially among Latino voters and other groups that face steeper “learning curves” about new requirements (Vercellotti and Andersen 2009). But more broadly, our finding of bias against Latinos even in states without voter ID suggests that minorities may face discrimination throughout their interactions with the electoral system.

This paper proceeds as follows. In Section 2, we review existing literature on bias by street-level bureaucrats and its implications for the implementation of voter ID laws. In Section 3, we introduce the hypotheses and design of our experiment, then present results and examine possible mechanisms in Section 4. Section 5 concludes.

2 Existing Literature

The behavior of street-level bureaucrats directly affects the realized forms of public policies (Brodkin 1997, Lipsky 1980). Although a policy might be written as universal, resource-constrained bureaucrats frequently have discretion over which services are delivered and to whom; discretionary decisions made by local bureaucrats often become *de facto* public policy. This behavior may also shape the long-run political context in which bureaucrats operate. Existing literature indicates that interactions with local bureaucracy – such as enrolling children in Head Start or receiving welfare benefits – can create feedbacks and influence rates of political participation (Bruch et al. 2010, Mettler and Soss 2004). If interactions with election officials have similar effects, bias could influence

the long-term composition of the electorate by affecting who is a voter. Given their critical role, it is important to understand why bureaucrats may treat constituents differently in discretionary interactions.

Existing research has advanced two broad classes of explanations for observed differences in the treatment of constituents. First, some studies argue that decisions by bureaucrats are driven largely by organizational rules, such as the technical-rational criteria in place for responding to complaints, and available resources, such as the manpower available in different jurisdictions (Jones et al. 1977, Mladenka 1981). This predicts that bureaucrats respond differently to requests because of variation in procedures and resources, not because of individual-level bias against groups of constituents.

A second set of studies argues instead that differences in treatment of constituents do stem from biases. These works focus on how individual characteristics – including ethnicity and partisanship – of the bureaucrat, the citizen, or the interaction of the two affect bureaucratic behavior. This literature proposes a combination of taste-based and more strategic, or political, mechanisms for this bias. Much of this comes from studies of the welfare state. Focusing on administration of the Temporary Assistance for Needy Families (TANF) program, Keiser et al. (2004) finds that non-white recipients' benefits are sanctioned by case workers for violating program rules more than white recipients within each jurisdiction. Similarly, Ernst et al. (2013) finds that a white applicant received higher quality information about benefits in an audit of welfare offices, arguing that white applicants may be seen by caseworkers as more deserving of benefits from the state. Fording et al. (2007) adds a political dimension, finding that more conservative localities have higher levels of TANF sanctions, and proposes that this is because local bureaucrats share their jurisdiction's majority political ideology about welfare programs. Davis et al. (2011) finds that where local TANF program managers have closer ties to black political networks in areas with more black

political power, black constituents are more likely to be given discretionary access to vocational training programs.¹¹

Existing studies of local election officials indicate that they may be similarly influenced by partisanship or ethnicity. The interaction between the partisanship of these officials and their district is associated with aggregate turnout (Burden et al. 2013), the number of provisional ballots cast and counted (Kimball et al. 2006), and local election officials' support for provisional ballot policies (Kropf et al. 2013), all suggestive that partisanship may affect how vigorously election rules are enforced. There is also evidence that citizens of different ethnicities have different experiences on election day, with blacks and Latinos more likely to be asked for ID by poll workers (Ansolabehere 2009, Atkeson et al. 2010, Cobb et al. 2012) and reporting lower quality experiences at the polls (Hall et al. 2009). As with the literature on discretionary bias in the welfare state, however, these studies cannot clearly distinguish whether these patterns emerge from underlying biases of local administrators and poll workers or if they are instead driven by more political and strategic considerations.

Recent changes to voter ID laws provide an important case in which to test theories of bureaucratic behavior. Since the passage of the federal Help America Vote Act (HAVA) in 2002, twenty-one states have adopted new ID requirements. Thirty states currently ask for some type of ID on election day (*National Conference of State Legislatures* 2012). These laws fall on a spectrum from requiring a photo ID to those that allow voters to present non-photo IDs and those which request some form of ID, but allow voters to cast ballots without it under certain conditions. Even if these new laws are designed to simplify the voting process, they increase complexity in the short

¹¹Outside the bureaucratic context, a large body of experimental research suggests that individuals – even those who openly espouse equality and fairness – often display discriminatory behavior when making discretionary choices. Experiments in areas as disparate as medical diagnosis and treatment, housing searches, and job applications find that otherwise identical whites and non-whites are treated differently (Ross et al. 2002, Green et al. 2007, Bertrand and Mullainathan 2004, Pager 2007, Pager and Shepherd 2008, Pager et al. 2009). Studying politicians rather than bureaucrats, Butler and Broockman (2011) finds that state legislators are more likely to respond to emails about voter registration from putatively white constituents than from putatively black constituents.

term by changing rules that voters and officials had become familiar with. This is especially true when laws are challenged in court and it is unclear whether they will be enforced. In states with new or legally contested voting laws, more voters can be expected to seek new information on requirements, and any underlying ethnic or racial differences in information provision may have a more pronounced effect on voting.

It may seem *a priori* unlikely that small differences in access to information about ID requirements would influence turnout. Recent scholarship, however, has shown that small changes in the costs of obtaining information about an action can significantly influence voluntary behaviors. For example, simplifying information and providing assistance in completing forms can lead to significantly higher rates of college enrollment among students from low-income families (Bettinger et al. 2012), more applications to selective colleges from high-achieving students who are eligible for scholarships (Hoxby and Turner 2013), and higher uptake of valuable tax credits among eligible tax payers (Bhargava and Manoli 2013, Saez 2009). Importantly, these experimental studies show that small information costs deter action even when substantial individual benefits are at stake. Similar information costs may also affect voting, which lacks such obvious benefits. Indeed, existing research has also shown that small changes to voting procedures, such as changing the location of polling places by several city blocks, decrease turnout by increasing information costs to voting (Brady and McNulty 2011, McNulty et al. 2009).

3 Hypotheses and Research Design

We conduct a field experiment measuring whether local election officials in 48 US states respond to email queries from constituents of different ethnicities at different rates or with information of differing quality.

3.1 Hypotheses

H1 – Bias in responses against Latino emailers: Consistent with literature showing that discretionary bureaucratic behavior is influenced by constituents’ individual characteristics, we hypothesize that election officials respond to emails from putatively Latino constituents at lower rates than to those from putatively non-Latino white constituents. This bias could also manifest in less accurate, informative, and friendly responses.

H2 – Greater bias against Latino emailers in emails about voter ID laws: We expect particularly strong bias in responses to questions about voter ID laws compared to other questions. This could be for two reasons. First, the focus on undocumented immigrants in public discourse on voter ID laws may have caused these laws to become “racialized,” priming underlying biases.¹² Alternatively, bias may be greater for tasks that are more complex for bureaucrats (Lipsky 1980), with greater differences across ethnic groups in responses to more complicated queries.

H3 – Greater bias against Latino names in states with stricter voter ID laws: Finally, we hypothesize that officials in states with stricter voter ID laws display greater bias in responsiveness against Latinos. We expect concerns about voting by non-citizens to be more salient in states with stricter laws. This may be because such concerns caused these states to adopt strict voter ID laws, or because exposure to more debate about ID requirements has made these laws more “racialized” in these states.

3.2 Experimental Set-up

To test these hypotheses, we contacted every local official or election commission for each county or municipality at which elections are administered in 48 states. We sent emails to publicly avail-

¹²Existing research suggests that media coverage linking a policy to race or ethnicity causes people to associate the policy with their pre-existing biases, even when that policy does not have explicit racial or ethnic content (Gilens 2004).

able addresses through which these officials already field constituent requests.¹³ Results are analyzed from 46 of these states ($N=6825$); two states were dropped due to independence assumption violations.¹⁴ Table 1 lists sample sizes and voter ID law types by state.¹⁵

[Table 1 about here]

Each local official or commission received an email on September 4, 2012 with a randomly assigned text and name. The “voter ID email” asked a question about voter ID laws, while the “control email” asked a simpler and less-politicized question about voting in primary versus general elections. Having two questions allows us to distinguish between baseline discrimination and discrimination specific to voter ID. We also varied the sender of the emails among four names, two that are putatively Latino (“José Martinez” and “Luis Rodriguez”) and two that are putatively non-Latino white (“Greg Walsh” and “Jake Mueller”).¹⁶ To send many emails in a short time window, we registered domain names and created email addresses such as `lrodriguez@ajnmail.net`. The text of the voter ID email was:

¹³In most states, elections are administered at the county level, and the unit of analysis is the county. In six states elections are administered by town or city, and the unit of analysis is the municipality. In two states we analyze – Maryland and Illinois – elections are generally administered by county, but by municipality in large cities. Officials in Maine and Alaska were not contacted. In Maine, email addresses for local election officials were unavailable. In Alaska, electoral districts do not correspond directly to municipal or county boundaries, making it difficult to merge data on covariates.

¹⁴In Virginia, separate emails forwarded to the same officials in some cases, meaning they received multiple treatment conditions simultaneously, violating the *SUTVA* (stable unit treatment value) assumption of no spillover between units needed to identify effects (Rosenbaum and Rubin 1983). There were similar concerns about Minnesota, where officials indicated that they had communicated with each other about the emails in their first set of replies. We drop Virginia and Minnesota from all analysis.

¹⁵There may be concern that our results are being driven primarily by the states with many jurisdictions, such as Michigan. We re-estimate each of the models reported in Section 4.1 and 4.2, sequentially dropping each of the seven states with the largest sample sizes. All results are robust to dropping each of these states (see Supporting Information).

¹⁶We use two names per ethnicity to ensure that results are not due to idiosyncratic characteristics of any specific name. The non-Latino white names are taken directly from Butler and Broockman (2011) and Bertrand and Mullainathan (2004). We follow Butler and Broockman (2011) and other audit studies in using only male names to conserve power and avoid proliferation of treatment conditions. To test for consistency violations, we conducted a preliminary survey using Amazon’s *Mechanical Turk* (Berinsky et al. 2012). These names overwhelmingly (> 96%) cued the correct ethnicities. Within the pairs of Latino or non-Latino names, the two names each cued very similar characteristics on other dimensions, such as perceived level of education, income, race, age, and political engagement. Results of this pre-test are in the Supporting Information.

Hello,

I've been hearing a lot about voter ID laws on the news. What do I need to do to vote?

Thank you,

(Jose Martinez, Jake Mueller, Luis Rodriguez, or Greg Walsh)

The control email was:

Hello,

I've been wondering about this. Do you have to vote in the primary election to be allowed to vote in the general elections?

Thank you,

(Jose Martinez, Jake Mueller, Luis Rodriguez, or Greg Walsh)

In sum, there were eight treatment conditions (4 names x 2 email texts). We analyze the pairs of Latino or non-Latino names together, collapsing the study to four conditions.¹⁷ These treatments were assigned within each state after pair matching localities, following Imai et al. (2009). We used coarsened exact matching (Iacus et al. 2012) to create pairs of counties or municipalities within states that were most similar in terms of population density, income per capita, proportion black, proportion Latino, and Democratic vote share in the 2008 presidential election.¹⁸ This improves balance on observable covariates prior to randomization. Importantly, by matching on factors that should predict bureaucratic capacity (population density and income per capita), and more generally by randomizing the treatment, we ensure that any ethnic disparities observed in the

¹⁷There are not substantively different outcomes within the pairs of Latino or non-Latino names.

¹⁸We merged demographic data from the 2010 US Census and 2005 American Community Survey (ACS) about each county or municipality to the list of official email addresses. We match on population density and income per capita as predictors of capacity and professionalization of local officials. We match on the proportion of a county's population that is African-American or Latino as predictors of the racial and ethnic diversity of an area. Finally, we match on Democratic vote share in the last presidential election to account for partisan responses to potential Latino voters.

experiment cannot be driven by differences in bureaucratic procedures or other jurisdiction-level factors.

The ethnicities of the email senders (Latino, non-Latino) and the two email texts were randomized within each pair with equal probabilities, such that one of the units received an email from a Latino name and the other did not.¹⁹ After matching, balance for both treatments (text and ethnicity of emailer) remained imperfect, because of the small number of observations within some states (see Supporting Information).²⁰ For this reason, we control for these covariates in the analysis below rather than looking only at differences in means across treatment conditions.

3.3 Coding of Responses

We received over 5300 replies to these emails from at least 4557 local election officials.²¹ Using a dataset of voter ID requirements in each state,²² we used double-blind human coding to assess the accuracy of each of the 2549 replies to the voter ID email. Specifically, the names of our fictitious emailers were removed and research assistants were not told the research questions. Research assistants then determined whether each email was “absolutely accurate,” meaning the reply contained full and accurate information about voter ID requirements in the state, or fell into a range of other categories.²³ Each reply was initially coded by two research assistants; replies for which the

¹⁹In states with odd numbers of units, the least similar unit was randomly assigned to treatment on its own.

²⁰For example, in many states the largest city or urban county, such as Chicago in Illinois, did not have any close match on the covariates.

²¹Multiple emails per official usually involved an automatic “out of the office” message followed by a substantive reply. In such cases, we collapse the emails into a single observation. Although some officials replied from different addresses than we had emailed, we can match all but 19 replies to the officials we originally emailed. The 19 unidentified emails are dropped from all analyses.

²²Voter ID requirements in some states changed in the final months before the 2012 election due to litigation and court rulings. We code accuracy with respect to the laws in each state as of the date our email was sent (Sept. 4). The requirements used in our coding process are available on request.

²³The full coding rules and examples of each type of email are provided in the Supporting Information, as is a set of additional analyses demonstrating that collapsing our coding categories in different ways does not substantially change the results presented here.

two coders disagreed were sent to a third.²⁴

In contrast, the control email was designed to examine baseline differences in response rates and had a single, simple answer in all states; no state requires voting in primaries to vote in general elections and this email was answered accurately essentially everywhere. Therefore, we analyze differences between the two email texts in response rate, not in accuracy.

4 Experimental Results

Table 2 contains summary statistics, broken out by treatment condition. Overall, officials responded to about 71% of emails, a higher response rate than in a similar study of elected state legislators (Butler and Broockman 2011). In addition, officials were more likely to respond to the control email than the voter ID email. This is consistent with the control email being marginally easier to answer; yet despite the low cost, officials responded to fewer than 74% of these emails. Only 46% of emails asking a question about voter ID got some type of accurate response. Furthermore, we find that local officials in states without voter ID laws were more likely to respond accurately to the voter ID email than those in states with additional requirements.²⁵ Finally, the summary statistics provide evidence of bias; Latino emailers receive responses at lower rates and are less likely to receive absolutely accurate responses than non-Latino emailers.

[Table 2 about here]

²⁴On the first pass, the two research assistants disagreed on the coding for accuracy in 23% (558) of the emails. Of these, 46% (270) were disagreements over the distinction between the “Absolutely accurate” category versus two other categories for responses that were either too general or too narrow but did not contain explicitly inaccurate information. Emails for which the first two coders disagreed on accuracy or other variables, such as friendliness, were sent to a third coder. If all three coders disagreed about at least one variable, the authors (single-) blind coded the responses, settling coding for 126. This left 90 emails for which we could not determine the accuracy, coded as “Ambiguous.” All results below are robust to dropping ambiguous and author-coded observations (see Supporting Information).

²⁵We do not see significant differences across states with different types of ID laws, however; see below.

Because randomization did not achieve perfect balance, we prefer parametric approaches that control for locality characteristics that may affect response rates.²⁶ We use two different modeling approaches. First, we run multi-level models in which coefficients for the effect of the Latino name and the intercept vary by state, while coefficients for the control variables are pooled across all states. This estimates different effects for every state, including those with small sample sizes, but borrows information across states (Gelman and Hill 2007). We also run a single fully pooled regression model with state fixed effects. This allows intercepts for each state to vary, but estimates a single treatment effect for the entire country, driven more heavily by states with larger sample sizes. We report results from both models for comparison.

4.1 Response Rate

We find strong evidence supporting our first hypothesis that the response rate to Latino names is significantly lower than the response rate to non-Latino names ($p < 0.001$ in both models). Because logistic coefficients are difficult to interpret, Panel (A) of Figure 1 plots the simulated first difference in the predicted probability of a response to a Latino versus non-Latino name for both email texts and 95% confidence intervals around these point estimates.²⁷ Responses to Latino names are three-and-a-half to four percentage points less likely than to non-Latino white names in both fully and partially pooled models.

[Figure 1 about here]

²⁶As shown in Table 2, 240 emails bounced back because of an incorrect email address, which could also affect balance. Dropping the other half of each pair in which an observation was not treated due to a bounce-back and re-estimating the results presented in Sections 4.1 and 4.2 does not change these findings (see Supporting Information).

²⁷Models include controls for each jurisdiction's 2008 Democratic presidential vote, percent non-Latino, percent Black, logged population density, and logged per capita income. Following Hanmer and Kalkan (2013), for all first differences in the paper we simulate the change in predicted probability of a response after changing the Latino name treatment from 0 to 1 for all observations in the data, keeping covariates fixed to the observed values for each observation, with 1000 simulations each.

We find limited support for our second hypothesis that the gap between response rates to non-Latinos and Latinos is larger for questions about voter ID laws than a general question. Panel (B) of Figure 1 shows the simulated first differences to Latino and non-Latino white names for each email text separately. The estimated bias is 5 percentage points (95% CI: -0.08,-0.02) for the voter ID email and 2 percentage points (95% CI: -0.05, 0.00) for the control email. Bias against Latino emailers in responses to the voter ID email text is thus estimated to be approximately three percentage points larger than in the control email, but the difference between them is not statistically significant ($p > 0.17$ in both specifications).

We find no evidence supporting our third hypothesis that states with voter ID laws exhibit greater differences in response to Latinos versus non-Latinos.²⁸ Testing this final hypothesis involves a non-experimental comparison of treatment effect sizes across states with different laws that were chosen endogenously. Furthermore, due to small sample sizes by law type (e.g., only 11 states with photo ID laws), we should only find evidence confirming our third hypothesis if the difference between states with different law types is very large. In two additional specifications, we include indicators for each state's type of law, and interact these with the Latino name variable. We find no interactions between the type of voter ID law and the Latino name variable on response rate (the smallest p -value on these interactions is 0.19; see Supporting Information).

²⁸We follow the *National Conference of State Legislatures's* coding of laws on two dimensions. First, laws require different types of ID (photo, non-photo, or none). Second, under "strict" laws, a ballot cast by a voter without ID will not be counted without additional steps by the voter to prove her identity after Election Day, while under laws that are "not strict" at least some voters lacking ID have options to cast ballots that will still be counted on Election Day, such as by signing an affidavit when voting attesting to their own identity. We find no interactions between law type and bias on either of these dimensions, or their interaction. These null results also hold when including states in which laws were enacted but not yet enforced.

4.2 Characteristics of Responses

We also investigate whether, conditional on receiving a reply, the nature of responses varies between emails to Latino and non-Latino aliases.²⁹ We focus only on responses to the voter ID email because these varied in accuracy.

First, Latino names are less likely to receive absolutely accurate responses ($p < 0.05$ in both specifications).³⁰ A key reason Latino emailers received fewer absolutely accurate responses is that rather than answering their question directly, local officials were more likely to respond to Latino emails “non-informatively,” side-stepping our direct question by, for example, asking whether the emailer wished to vote absentee or requesting that the emailer call the office to discuss the issue ($p < 0.05$ in both specifications).³¹ Overall, simulations based on partially pooled models estimate Latino emailers to be 4 percentage points (95% CI: -0.08,0.00) less likely to receive absolutely accurate responses and nearly 5 percentage points (95% CI: 0.01, 0.08) more likely to receive a non-informative response compared to non-Latino emailers.

In addition to the content of responses, we examine whether the tone or style of responses to Latino and non-Latino emailers differs. Two-thirds of all responses to the voter ID question were coded as “friendly.”³² Latino emailers are estimated to be 4 percentage points less likely to receive a friendly response, but this result falls short of statistical significance in both model specifications

²⁹This section focuses only on responses received; we have no way of inferring how accurate responses from localities that did not respond at all might have been. That we still find differences in the responses to Latino and non-Latino names after conditioning on receiving any response is perhaps even more striking. If we look instead at whether each email received an accurate response or did not, combining cases where we received no response at all with cases that received a response not coded as accurate, the negative effect of the Latino name appears even larger because the non-response bias against Latino names compounds with the quality-of-response bias reported in this section.

³⁰Regression tables for all analysis discussed in this section appear in the Supporting Information.

³¹Within the non-informative responses, Latino emailers were not more likely than non-Latino emailers to be asked whether they were registered to vote.

³²Emails marked as “friendly” contained “explicitly friendly language, such as use of the senders’ name in the salutation or sign-off.” Examples included “Dear (name),” “Let us know if you have any more questions” and “Have a great day.”

($p=0.13$). Taken together, these results suggest differential treatment of Latino and non-Latino emailers, even after conditioning on response rate.

4.3 The Voting Rights Act: Bias under Monitoring

Within the clear aggregate pattern of bias in responsiveness against Latinos, we also find suggestive evidence that bias from local election officials varies with the degree of monitoring they face.³³ Research on racial and ethnic discrimination in employment, housing, financial, and consumer markets suggests that private sector employees are less likely to be discriminatory in discretionary decisions about clients when organizational procedures are in place to make them aware of the risk for discrimination and to explicitly monitor against it (Pager and Shepherd 2008, 14-17).³⁴ If similar dynamics hold for election officials, biased behavior should be less likely when institutional mechanisms are in place to prevent it.

We test this by analyzing response rates from jurisdictions that are and are not covered by two federal anti-discrimination policies explicitly aimed at preventing local-level discrimination in election administration. Under Section 5 of the federal Voting Rights Act (VRA) – which was in effect during our study – 706 localities in our data were subject to “pre-clearance.” To prevent discrimination, this coverage placed localities under additional federal oversight and prohibited them from changing election procedures without federal approval. Under Section 203 of the VRA, 186 jurisdictions from our study are specifically required by the federal government to make all voting information and ballots accessible to Spanish-language voters because they have large Latino citizen populations, a high proportion of Latino citizens, or many eligible voters who are Spanish language-dominant.³⁵

³³The previous analyses followed a pre-analysis plan, available on request. The remaining analyses are exploratory.

³⁴Similarly, research on policing has found less racial profiling when street-level officers are monitored and given less discretion in enforcement (Gottfredson and Gottfredson 1988, Harris 1997).

³⁵There is some overlap between these: 93 observations were covered by both Sections 5 and 203 in 2012.

Assignment to increased scrutiny under these provisions is not random and we cannot estimate causal effects of the VRA. Nonetheless, adding indicator variables for whether an observation is covered under either Section 5 or Section 203 to the models from Section 4.1 and interacting these indicators with the Latino name treatment shows that bias against Latinos is significantly lower in these subsets. Places covered by the VRA under either Section 5 or Section 203 show no response bias ($p=0.30$), while non-covered places show significant bias in response rates ($p = 0.0001$), with an estimated 5 percentage points fewer responses to Latino emailers (95% CI: -0.07,-0.02).³⁶ Plotting this interaction, Panel (A) of Figure 2 shows that bias against Latinos is 7.5 percentage points lower (95% CI: 0.8, 13.3) in covered localities than in non-covered localities. Because areas covered by the VRA likely differ from non-covered places on many other dimensions, we also use nearest neighbor propensity score matching to prune the set of non-covered localities to those that most resemble covered localities on the covariates listed above. Re-estimating the model on this restricted dataset shows a similar pattern, with 6.6 percentage points (95% CI: -1.8, 15.4) less discrimination in covered than non-covered jurisdictions, although this is only significant at $p = 0.12$, with the reduced sample size.

[Figure 2 about here]

We also analyzed Section 5 and Section 203 coverage separately because the determinants and nature of coverage for each differ. Panel (B) of Figure 2 shows that re-estimating the models for Section 5 alone returns a similar result, with estimated 6.1 percentage points less discrimination in covered than non-covered jurisdictions ($p = 0.08$, 95% CI: -0.8, 13.2).³⁷ In the unmatched data, Section 5 jurisdictions are estimated not to discriminate against Latinos in responsiveness ($p = 0.49$), while the un-covered jurisdictions do discriminate ($p = 0.0002$), with 4.4 percentage

³⁶The p-value on the interaction between VRA coverage and the Latino treatment indicator is 0.02 (see Supporting Information for corresponding table).

³⁷The p-value on the interaction in the smaller matched dataset is 0.17, however.

points estimated bias against Latinos (95% CI: -0.07,-0.02).³⁸ We also find significant differences across Section 203 coverage. Panel (C) of Figure 2 plots the interaction between the Latino name treatment and an indicator for Section 203 coverage, showing significantly less discrimination in covered than non-covered localities.³⁹ In the unmatched data, officials in the 186 observations covered by Section 203 for Spanish-speaking populations are signed as more likely to respond to the Latino than the non-Latino names, while officials in localities not covered are still 4 percentage points less likely to reply to Latinos than non-Latinos ($p = 0.0001$, 95% CI: -0.06, -0.02).

An additional factor that could produce variation in discrimination across localities is the ethnicity of the officials themselves. Studies of representative bureaucracy argue that the interests of groups will be better represented by a bureaucratic workforce that includes members of those groups (Kingsley 1944, Meier 1975, Krislov 2012). Local officials who are themselves Latino/a may be less biased, and local election officials are more likely to be Latino in localities with larger Latino populations, including those with Section 203 coverage. We are unable to identify the ethnicity of most officials in our dataset. But for those states where ethnicity data is available, it appears that the overwhelming majority of local election officials are non-Latino.⁴⁰ We cannot fully reject the possibility that some of the difference between Section 203 covered places and other jurisdictions is driven by these localities also being more likely to have Latino/a officials.⁴¹

³⁸Because Section 5 coverage is highly correlated with being in the South, however, the result for Section 5 jurisdictions cannot be distinguished from a more general regional difference between southern states and the rest of the country (see Supporting Information). But Section 203 localities are spread widely across the US and also appear to be less biased against Latinos.

³⁹The p-value on the interaction between Section 203 coverage and the Latino treatment is 0.01 in the unmatched data and 0.12 in the matched data, with the smaller sample size of $N = 372$.

⁴⁰In many cases, we do not know the specific official's name (unless they happened to sign their email response), and even when we do, there is little available information to code ethnicity. We attempted to identify ethnicity of these officials using a database of Latino local government officials maintained by the National Association of Latino Election and Appointed Officials (NALEO). In many states, NALEO does not collect data on the relevant office in our study. But in the 17 states where most of the offices we contacted are included in the dataset, we find that 69 out of 1290 officials contacted, or only 5%, were Latino/a. This is too few to estimate bias among Latino officials separately to compare with the non-Latino officials.

⁴¹We attempt to estimate a separate effect of Section 203 coverage from the effect of the Latino population by using total Latino population as a matching variable in addition to percent Latino population (which is already included in

We can show, however, that the reduction in bias in localities covered by Section 5 in the VRA exists separately from any differences in Latino population because areas covered under Section 5 range widely in local Latino population. Removing all localities that are more than 10% Latino and re-estimating the models for Section 5 coverage shows that these results are the same without including localities with high Latino populations (see Supporting Information). There is thus still evidence that monitoring may change bureaucratic behavior even in places that are unlikely to have Latino bureaucrats.

4.4 Potential Mechanisms

The bias we observe in areas without VRA coverage could result from several mechanisms. Although the appearance of bureaucratic bias elsewhere might result from differences in organizational rules or resources (Mladenka 1981), randomization ensures that these factors do not drive the bias we observe in our experiment. We explore two additional mechanisms. First, local bureaucrats may engage in taste-based discrimination when trying to decide which of many tasks to prioritize. Alternatively, public officials may not respond as often to Latinos because they believe that they are less likely to support these officials politically, to actually live in their districts, or to be citizens. This second set of reasons could be seen as forms of “statistical discrimination,” with bias emerging from potentially rational beliefs about differences between Latino and non-Latino constituents (Altonji and Blank 1999). While our design does not allow us to causally identify these mechanisms, a series of additional tests suggest that the most plausible statistical discrimination mechanisms are insufficient.

Partisan elected officials may have political incentives to reply at different rates, or with differing accuracy, if ethnicity serves as a cue for partisanship (Butler and Broockman 2011). For

the matching model). Results for this additional matched dataset are similar and are provided in the Supporting Information.

example, if Latino voters are perceived as more likely to be Democrats, Republican officials could reply less to discourage turnout of perceived opponents. Similar incentives would not hold for non-partisan civil servants, who have no political advantage at stake in these emails. We update data from Kimball and Kropf (2006) to identify the partisanship and appointment method of the local election officials and test several arguments about political incentives.⁴²

We find no evidence that political incentives affect bias. Elected and appointed officials respond and are accurate at similar rates. Appointment method does not interact with the Latino name – both types of officials are similarly biased overall.⁴³ Similarly, while there is some evidence that partisan officials reply less overall than non-partisan officials, partisan and non-partisan officials show no significant difference in bias towards Latinos. We also find no difference in bias between Democratic versus Republican election officials.⁴⁴ Finally, the partisanship of each official’s jurisdiction, measured as Obama vote share in 2008, does not interact with the Latino name or with each officials’ partisanship – i.e., Republican officials are not particularly likely to be biased in Democratic jurisdictions, or vice versa.

Beyond political incentives, bias in response rates could result from officials not seeing Latinos as plausible constituents. Because election administration is organized by county or locality in the US and the Latino population is not evenly distributed across the country, most local election

⁴²This was updated by consulting local elections results, where available, and otherwise individually searching for officials by office and email address. We are unable to reproduce this data for Michigan and Wisconsin, where Kimball and Kropf (2006) do not code municipal-level officials and heterogeneous local selection procedures would have to be individually researched for thousands of municipalities. Of the remaining 3433 officials for which we have updated data, 50.1% serve in officially partisan positions and 63.6% were elected as of 2012. These dimensions are cross-cutting: 35% of non-partisan officials are elected and 8% of partisan officials are appointees. For those offices (excluding Michigan and Wisconsin) where officials are formally partisan, we are able to identify the specific partisanship as of 2012 for 89%. We code officials who serve in formally non-partisan capacities as non-partisan, following Kimball and Kropf (2006).

⁴³See Supporting Information for regression tables.

⁴⁴When comparing Democrats and Republicans, we use a restricted sample of 1522 localities for which officials are partisan and partisanship is known. These 1522 localities are slightly less Latino and less dense than the national sample, while being more Republican in past election results, more African-American, and wealthier. Importantly, the overall results of the experiment continue to hold within this subset, with bias against Latinos in response rate and accuracy.

administrators serve in jurisdictions with low Latino populations. There are relatively few jurisdictions in the data with large Latino populations; the median observation is only 2.1% Latino. Interacting the Latino name treatment with the Latino population share in each locality returns a statistically significant interaction suggesting that bias against Latinos in responsiveness only becomes significant in areas with few Latinos.⁴⁵ But this does not account for VRA coverage. A disproportionate share of the localities with large Latino populations are covered under Sections 203 or 5 of the VRA.⁴⁶ Removing all localities where the above analysis suggests we should not find significant bias because they are covered by the VRA, there is no evidence of an interaction between the Latino name treatment and the Latino population share of each locality ($p = 0.80$ on this interaction term, see Supporting Information). Ultimately, because of the prevalence of VRA coverage in localities with high Latino populations, we cannot observe how much discrimination would exist in the most Latino localities in the US if the VRA were not in effect.⁴⁷

A related statistical discrimination mechanism would hold if officials could disproportionately tell that the Latino emailers were not their constituents in comparison to the non-Latino emailers. An extreme example would be if our findings were driven by officials in very small towns or counties who know most their constituents personally and can easily tell that our aliases are fake. We re-estimate the results in Sections 4.1 and 4.2, dropping localities in the bottom 10th percentile of total population. This has no substantive effect and all treatment effect estimates remain as before (see Supporting Information).

Another version of this mechanism would occur if local officials were less likely to find our Latino aliases than our non-Latino aliases in official lists of registered voters or locality residents,

⁴⁵ $p = 0.02$ on the interaction term, see Supporting Information for the full regression table.

⁴⁶Of the 725 localities that are more than 10% Latino, 351, or 48%, are covered by the VRA. In contrast, only 8% of the remaining 5717 observations in the data with populations less than 10% Latino are covered by the VRA.

⁴⁷As discussed in Section 4.3, we *can* show that the difference in response rate between VRA and non-VRA localities is separate from differences in Latino population because of the extensive variation in local Latino population within the Section 5 jurisdictions.

and then did not reply (or do so accurately) as a result. To test for this, our research assistants coded indicators for the officials being suspicious about our aliases: whether officials asked if the emailer lived in that locality or was registered to vote, as well as if officials indicated that they looked for the emailer in their voter registration file. The Latino names were no more likely to receive such questions about residency or registration than the non-Latino names. Instead, 12% of the officials who replied explicitly stated that they had no evidence that the emailer (of either ethnicity) lived in their locality. Moreover, our non-Latino white surnames, Mueller and Walsh, are actually much less common in the US population than Martinez and Rodriguez⁴⁸ and our specific Latino first and last name combinations are also orders of magnitude more prevalent than our non-Latino names.⁴⁹ Though the prevalence of our Latino aliases is not uniform across the US, in the clear majority of localities in our data the non-Latino aliases will not appear in the voter registration file either. Nonetheless, over 70% of officials replied to our non-Latino emailers, making it implausible that finding names in the voter registration file accounts for whether or not officials responded.

⁴⁸As of the 2000 Census, Rodriguez was the ninth most common surname in the entire country; Martinez was eleventh. Walsh and Mueller ranked 265th and 467th, respectively. See <http://www.census.gov/genealogy/www/data/2000surnames/> (accessed 21 April 2013).

⁴⁹We searched for each of our aliases in <http://names.whitepages.com/>, a database that estimates the total instances of a specific name in the current population (accessed 7 July 2014) and which is increasingly used to construct samples for opinion polls (Guterbock et al. 2011). Although this undercounts each name, because of unlisted phone numbers and people with only mobile phones, the search shows a massive difference in name frequencies, returning 27,188 instances of “Jose Martinez” and 10,917 of “Luis Rodriguez,” compared to 442 of “Greg” or “Gregory Walsh” and 183 of “Jake” or “Jacob Mueller.” Greg/Gregory Walsh and Jake/Jacob Mueller appear in 42 and 30 states, respectively, while the Latino names appear in every state. Moreover, nationally representative survey data from 2006 indicates that Hispanic households are 15.2 percentage points *less likely* to have a listed phone number than non-Hispanic households, suggesting that our Latino aliases may be even more prevalent than the non-Latino aliases than indicated by phonebook data (National Center for Health Statistics and Prevention 2007). If anything, many officials will be more likely to have been exposed to people in their localities with the same names as the Latino, rather than Anglo, aliases.

5 Conclusions

We bring field experimental evidence to debates about the degree to which constituents' individual characteristics drive bureaucratic behavior and causally identify the effect of constituent ethnicity on the responsiveness of local election officials. This experiment allows us to directly study the public officials tasked with providing election information to citizens and to isolate the effect of ethnicity on their real-world performance. We find strong evidence that election officials were less responsive to questions from putatively Latino constituents. In addition, the responses that Latino emailers received to voter ID questions were less likely to be accurate and were more likely to be non-informative.

While we cannot causally identify mechanisms, our results suggests that bias from street-level bureaucrats can occur even when there are not clear strategic reasons for officials to discriminate on the basis of ethnicity. Randomization ensures that observed bias cannot be explained by differences in organizational capacity (e.g., Jones et al. 1977) and we find no evidence that whether local officials are elected or appointed, partisan or nonpartisan, or members of specific political parties influences the extent of bias. This differs from earlier observational work on partisan patterns in the behavior of election officials in other aspects of election administration (Burden et al. 2013, Kimball et al. 2006, Kropf et al. 2013) as well as observational evidence of partisan bias by other types of street-level bureaucrats (Fording et al. 2007). But consistent with the claim that enhanced monitoring reduces discrimination, we find no evidence of bias against Latinos in jurisdictions subject to federal regulation under the VRA.

These findings have important implications for debates about voter ID laws, and indeed for any changes to voting requirements or election administration. Our results indicate that changes to existing voting regulations are likely to differentially increase information costs for Latino voters because public officials are less responsive to their requests for information. That we find some evidence that officials respond at lower rates to Latinos even when asked a question requiring a

single word answer (“No”) suggests that public officials can be biased even in exceptionally low-cost interactions. If minority voters are less able to acquire information about ID requirements and more likely to be asked for ID at the polls (Ansolabehere 2009), this could manifest in lower voting rates. This may be greatest where officials are not monitored to prevent discrimination. The recent Supreme Court decision (*Shelby County v. Holder* (2013)) striking down portions of the VRA makes it all the more imperative that future reforms to electoral rules include serious efforts to disseminate information to all voters to offset potential bias at an individual level.

More broadly, our experiment confirms that bias by street-level bureaucrats based on citizens’ ethnicity is not always driven by partisan incentives, other plausible forms of “statistical discrimination,” or organizational characteristics. In addition to elections, street-level bureaucrats have discretionary control over the provision of many goods and services controlled by local governments in the US – from trash collection and snow plowing to the management of welfare offices. Similar ethnic or racial biases may affect the quality of services delivered in these other arenas. Future research can expand the use of experimental methods to examine the presence of bias in service delivery in many other aspects of local administration in the US.

Table 1: Number of Email Addresses Used by State

State	No.	Unit	Law Type	State	No.	Unit	Law Type
Alabama	67	County	Non-photo	Montana	56	County	Non-photo
Alaska	–	–	–	Nebraska	93	County	HAVA
Arizona	15	County	Non-photo	Nevada	17	County	HAVA
Arkansas	75	County	Non-photo	New Hampshire	233	Municipal	Photo
California	54	County	HAVA	New Jersey	17	County	HAVA
Colorado	64	County	Non-photo	New Mexico	33	County	HAVA
Connecticut	168	Municipal	Non-photo	New York	62	County	HAVA
Delaware	3	County	Non-photo	North Carolina	100	County	HAVA
Florida	67	County	Photo	North Dakota	53	County	Non-photo
Georgia	159	County	Photo	Ohio	87	County	Non-photo
Hawaii	3	County	Photo	Oklahoma	77	County	Non-photo
Idaho	42	County	Photo	Oregon	36	County	HAVA
Illinois	110	Mixed	HAVA	Pennsylvania	63	County	HAVA
Indiana	91	County	Photo	Rhode Island	39	Municipal	Non-photo
Iowa	99	County	HAVA	South Carolina	46	County	Non-photo
Kansas	105	County	Photo	South Dakota	64	County	Photo
Kentucky	120	County	Non-photo	Tennessee	95	County	Photo
Louisiana	64	County	Photo	Texas	254	County	Non-photo
Maine	–	–	–	Utah	25	County	Non-photo
Maryland	24	Mixed	HAVA	Vermont	245	Municipal	HAVA
Massachusetts [†]	315	Municipal	HAVA	Virginia*	–	–	–
Michigan [†]	1284	Municipal	Photo	Washington	39	County	Non-photo
Minnesota*	–	–	–	West Virginia	54	County	HAVA
Mississippi	81	County	HAVA	Wisconsin [†]	1756	Municipal	HAVA
Missouri	115	County	Non-photo	Wyoming	23	County	HAVA

Total = 6825. †: 20 units from Massachusetts, Michigan, and Wisconsin were used in a pilot and dropped from the full experiment. *: Minnesota and Virginia were emailed, but not analyzed due to concerns about contamination. Law types are those *in effect* in September 2012, when our emails were sent, based on definitions from *National Conference of State Legislatures* (2012). The “photo” category includes any laws that request voters present photo ID; the “non-photo” category includes any laws that require any identification; the “HAVA” category includes states without regulations beyond those required by federal law.

Table 2: Summary Statistics, All States

	<i>Voter ID Email</i>		<i>Control Email</i>		<i>Both Emails</i>	
	Latino	Non-Latino	Latino	Non-Latino	Latino	Non-Latino
States	46	46	46	46	46	46
Emails sent	1662	1664	1683	1683	3345	3347
Incorrect address / bounce	64	67	48	61	112	128
Response rate	1036 (64.8%)	1126 (70.5%)	1184 (72.4%)	1211 (74.7%)	2220 (68.7%)	2337 (72.6%)
Accurate, absolutely	523 (50.5%)	619 (55%)	1184 (100%)	1211 (100%)	1707 (76.9%)	1830 (78.3%)
Accurate, general or narrow	190 (18.3%)	203 (18%)	0 (0%)	0 (0%)	190 (8.6%)	203 (8.7%)
Non-informative	272 (26.3%)	248 (22%)	0 (0%)	0 (0%)	272 (12.3%)	248 (10.6%)
Inaccurate	13 (1.3%)	10 (0.9%)	0 (0%)	0 (0%)	13 (0.6%)	10 (0.4%)
Ambiguous	38 (3.7%)	45 (4%)	0 (0%)	0 (0%)	38 (1.7%)	45 (1.9%)
Friendly	665 (64.2%)	762 (67.7%)	0 (0%)	0 (0%)	665 (30%)	762 (32.6%)

“Incorrect address” refers to emails that either immediately were “returned to sender” or to which we received a reply informing us that the email was addressed to an incorrect recipient. Refer to text for other definitions.

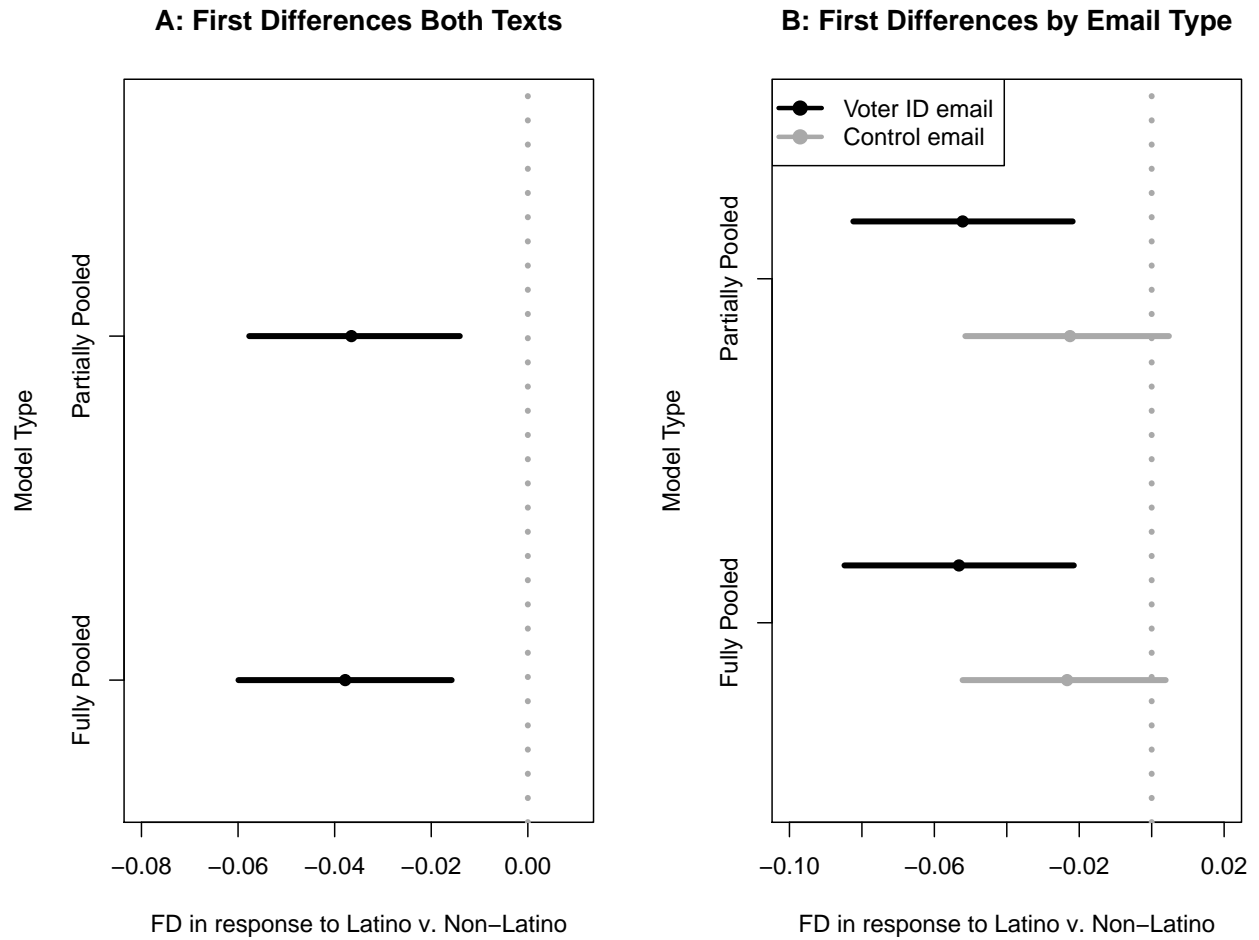


Figure 1: *First Differences in Overall Response Rates*: Panel (A) shows simulated differences between the non-Latino and Latino names in the overall predicted probability a local official responds to either email text, with 95% confidence intervals. Panel (B) shows first differences in response rate between non-Latino and Latino names broken down by each email text. Regression tables are in the Supporting Information

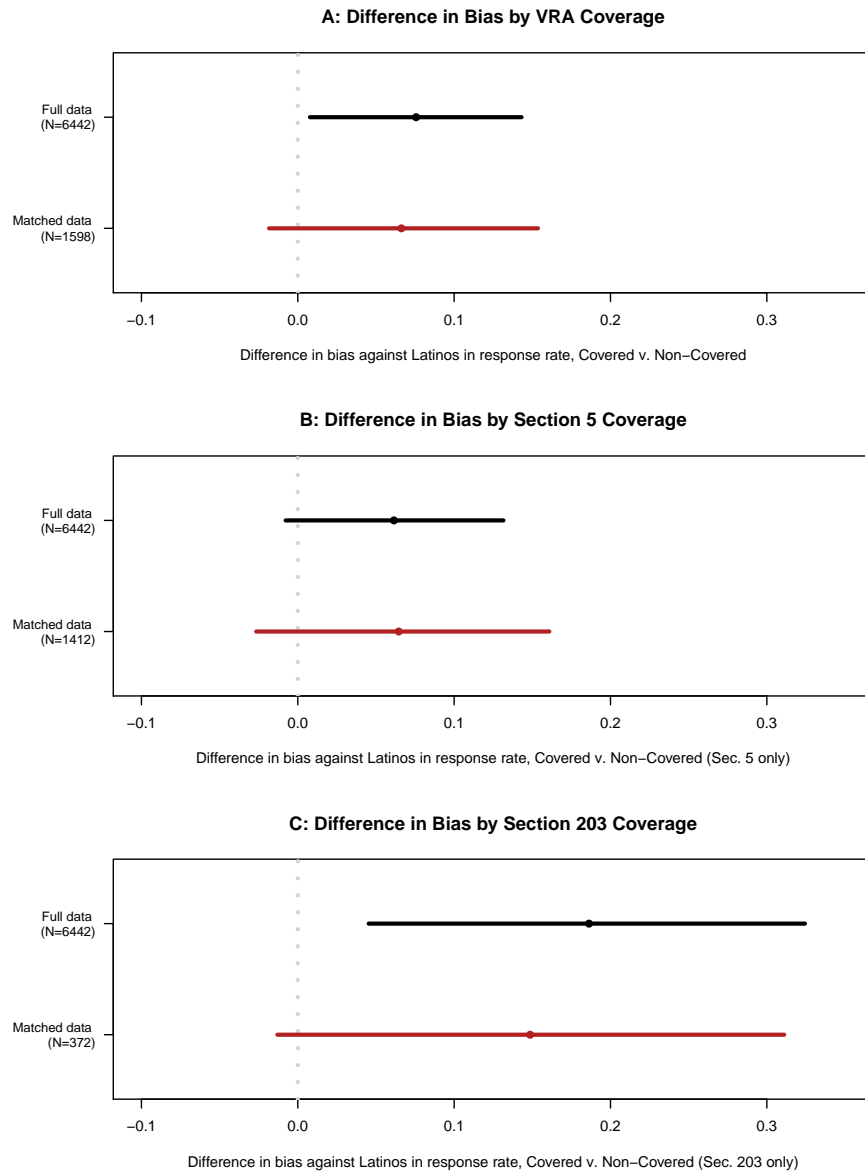


Figure 2: *First Differences in Latino Treatment Effect by VRA Status*: Simulated differences in the size of the Latino name treatment effect after a change in VRA status, with 95% confidence intervals. Positive values indicate a smaller Latino treatment effect, and thus *less* bias. The matched data is from propensity score matching on VRA status as the treatment, as described in the text. Corresponding regression tables are in the Supporting Information.

References

- Altonji, Joseph G. and Rebecca Blank. 1999. Race and Gender in the Labor Market. In *Handbook of Labor Economics Volume 3c*, ed. Orley C. Ashenfelter and David Card. Elsevier Science.
- Alvarez, R. Michael, Delia Bailey and Jonathan N. Katz. 2008. “The Effect of Voter Identification Laws on Turnout.” California Institute of Technology Social Science Working Paper 1267R.
- Ansolabehere, Stephen. 2009. “Effects of Identification Requirements on Voting: Evidence from the Experiences of Voters on Election Day.” *PS: Political Science & Politics* 42(1):127–130.
- Atkeson, Lonna Rae, Lisa A. Bryant, Thad E. Hall, Kyle L. Saunders and R. Michael Alvarez. 2010. “A new barrier to participation: Heterogenous application of voter identification policies.” *Electoral Studies* 29(1):66–73.
- Barreto, Matt A., Gabriel R. Sanchez and Hannah Walker. 2012. “Rates of Possession of Valid Photo Identification, and Public Knowledge of Voter ID Law in Pennsylvania.” Expert Report Submitted on Behalf of Plaintiffs in Applewhite, et al. v. Commonwealth of Pennsylvania et al., No. 330 MD 2012.
- Barreto, Matt A., Stephen A. Nuño and Gabriel R. Sanchez. 2007. “Voter ID Requirements and the Disenfranchisement of Latino, Black and Asian Voters.” Prepared for presentation at: 2007 American Political Science Association Annual Conference.
- Barreto, Matt A., Stephen A. Nuño and Gabriel R. Sanchez. 2009. “The Disproportionate Impact of Voter-ID Requirements on the Electorate - New Evidence from Indiana.” *PS: Political Science & Politics* 42(1):111–116.
- Berinsky, Adam J., Gregory A. Huber and Gabriel S. Lenz. 2012. “Evaluating Online Labor Markets for Experimental Research: Amazon.com’s Mechanical Turk.” *Political Analysis* 20(3):351–368.
- Bertrand, Marianne and Sendhil Mullainathan. 2004. “Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination.” *American Economic Review* 94(4):991–1013.
- Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos and Lisa Sanbonmatsu. 2012. “The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment.” *The Quarterly Journal of Economics* 127(3):1205–1242.
- Bhargava, Saurabh and Dayanand Manoli. 2013. “Why Are Benefits Left on the Table? Assessing the Role of Information, Complexity, and Stigma on Take-up with an IRS Field Experiment.” Working Paper.
- Brady, Henry E. and John E. McNulty. 2011. “Turning out to Vote: The Costs of Finding and Getting to the Polling Place.” *American Political Science Review* 105(1):115–134.

- Brodkin, Evelyn Z. 1997. "Inside the Welfare Contract: Discretion and Accountability in State Welfare Administration." *Social Service Review* 71(1):1–33.
- Bruch, Sarah K., Myra Marx Ferree and Joe Soss. 2010. "From Policy to Polity: Democracy, Paternalism, and the Incorporation of Disadvantaged Citizens." *American Sociological Review* 75(2):205–226.
- Burden, Barry C., David T. Canon, Kenneth R. Mayer and Donald P. Moynihan. 2012. "The Effect of Administrative Burden on Bureaucratic Perception of Policies: Evidence from Election Administration." *Public Administration Review* 72(5):741–751.
- Burden, Barry C., David T. Canon, Stephane Lavertu, Kenneth R. Mayer and Donal P. Moynihan. 2013. "Selection Method, Partisanship, and the Administration of Elections:." *American Politics Research* 41(6):903–936.
- Butler, Daniel M. and David E. Broockman. 2011. "Do Politicians Racially Discriminate Against Constituents? A Field Experiment on State Legislators." *American Journal of Political Science* 55(3):463–477.
- Cobb, Rachael V., D. James Greiner and Kevin M. Quinn. 2012. "Can Voter ID Laws Be Administered in a Race-Neutral Manner? Evidence from the City of Boston in 2008." *Quarterly Journal of Political Science* 7:1–33.
- Davis, Belina Creel, Michelle Livermore and Younghee Lim. 2011. "The Extended Reach of Minority Political Power: The Interaction of Descriptive Representation, Managerial Networking, and Race." *Journal of Politics* 73(2):494–507.
- Erikson, Robert S. and Lorraine C. Minnite. 2009. "Modeling Problems in the Voter Identification – Voter Turnout Debate." *Election Law Journal* 8(2):85–101.
- Ernst, Rose, Linda Nguyen and Kamilah C. Taylor. 2013. "Citizen Control: Race at the Welfare Office." *Social Science Quarterly* 94(5):1283–1307.
- Estabrook, Leigh, Evans Witt and Lee Rainie. 2007. "Information Searches that Solve Problems: How People Use the Internet, Libraries, and Government Agencies When They Need Help." Pew Internet & American Life Project: <http://bit.ly/19Ivj7g>.
- Fording, Richard C., Joe Soss and Sanford F. Schram. 2007. "Devolution, Discretion, and the Effect of Local Political Values on TANF Sanctioning." *Social Service Review* 81(2):285–316.
- Gelman, Andrew and Jennifer Hill. 2007. *Data Analysis Using Regression and Multi-level/Hierarchical Models*. New York: Cambridge University Press.
- Gilens, Martin. 2004. Poor people in the news: Images from the journalistic subconscious. In *Class and News*, ed. Don Heider. Oxford, UK: Rowman and Littlefield pp. 44–60.

- Goodsell, Charles T. 1983. *The Case for Bureaucracy: A Public Administration Polemic*. Chatham, NJ: Chatham House Publishers.
- Gottfredson, Michael R. and Don M. Gottfredson. 1988. *Decision Making in Criminal Justice*. Springer.
- Green, Alexander R., Dana R. Carney, Daniel J. Pallin, Ngo H. Long, Kiristal L. Raymond, Lisa I. Iezzoni and Mahzarin R. Banaji. 2007. "Implicit bias among physicians and its prediction of thrombolysis decisions for black and white patients." *Journal of General Internal Medicine* 22(9):1231–1238.
- Grose, Christian R. 2014. "Field Experimental Work on Political Institutions." *Annual Review of Political Science* 17:335–370.
- Guterbock, Thomas M., Abdoulaye Diop, James M. Ellis, John Lee Holmes and Kien Trung Leigh. 2011. "Who Needs RDD? Combining Directory Listings with Cell Phone Exchanges for an Alternative Telephone Sampling Frame." *Social Science Research* 40(3):860–872.
- Hall, Thad E., J. Quin Monson and Kelly D. Patterson. 2009. "The Human Dimension of Elections: How Poll Workers Shape Public Confidence in Elections." *Political Research Quarterly* 62(3):507–522.
- Hancock, Ange-Marie. 2007. "When Multiplication Doesn't Equal Quick Addition: Examining Intersectionality as a Research Paradigm." *Perspectives on Politics* 5(1):63–79.
- Hanmer, Michael J. and Kerem Ozan Kalkan. 2013. "Behind the Curve: Clarifying the Best Approach to Calculating Predicted Probabilities and Marginal Effects from Limited Dependent Variable Models." *American Journal of Political Science* 57(1):263–277.
- Harris, David A. 1997. "'Driving While Black' and All Other Traffic Offenses: The Supreme Court and Pretextual Traffic Stops." *Journal of Criminal Law and Criminology* pp. 544–582.
- Hood, III, M.V. and Charles S. Bullock, III. 2012. "Much Ado About Nothing? An Empirical Assessment of the Georgia Voter Identification Statute." *State Politics & Policy Quarterly* 12(4):394–414.
- Hosoda, Megumi, Lam T. Nguyen and Eugene F. Stone-Romero. 2012. "The effect of Hispanic accents on employment decisions." *Journal of Managerial Psychology* 27(4):347–364.
- Hoxby, Caroline and Sarah Turner. 2013. "Expanding College Opportunities for High-Achieving, Low Income Students." SIEPR Discussion Paper No. 12-014.
- Iacus, Stefano M., Gary King and Giuseppe Porro. 2012. "Causal Inference without Balance Checking: Coarsened Exact Matching." *Political Analysis* 20(1):1–24.
- Imai, Kosuke, Gary King and Clayton Nall. 2009. "The essential role of pair matching in cluster-randomized experiments, with application to the Mexican universal health insurance evaluation." *Statistical Science* 24(1):29–53.

- Jones, Bryan D., Saadia R. Greenberg, Clifford Kaufman and Joseph Drew. 1977. "Bureaucratic Response to Citizen-Initiated Contacts: Environmental Enforcement in Detroit." *American Political Science Review* 71(1):148–165.
- Keiser, Lael R., Peter R. Mueser and Seung-Whan Choi. 2004. "Race, Bureaucratic Discretion, and the Implementation of Welfare Reform." *American Journal of Political Science* 48(2):314–327.
- Kimball, David C. and Martha Kropf. 2006. "The Street-Level Bureaucrats of Elections: Selection Methods for Local Election Officials." *Review of Policy Research* 23(6):1257–1268.
- Kimball, David C., Martha Kropf, and Lindsay Battles. 2006. "Helping America Vote? Election Administration, Partisanship, and Provisional Voting in the 2004 Election." *Election Law Journal* 5(4):447–461.
- Kingsley, John Donald. 1944. *Representative Bureaucracy: An Interpretation of the British Civil Service*. Antioch Press.
- Krislov, Samuel. 2012. *Representative Bureaucracy*. Quid Pro Books.
- Kropf, Martha, David C. Kimball and Timothy Vercellotti. 2013. "Representative Bureaucracy and Partisanship: The Implementation of Election Laws." *Public Administration Review* 73(2):242–252.
- Lipsky, Michael. 1980. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services*. New York: Russell Sage Foundation.
- Lopez, Mark Hugo and Ana Gonzalez-Barrera. 2012. "Latino Voters Support Obama by 3-1 Ratio, But Are Less Certain than Others about Voting." Pew Hispanic Center: <http://bit.ly/15IOviw>.
- McClendon, Gwyneth H. 2012. "Race, Responsiveness, and Electoral Strategy: A Field Experiment with South African Politicians." Working Paper, Princeton University.
- McNulty, John E., Conor M. Downing and Margaret H. Ariotti. 2009. "Driving Saints to Sin: How Increasing the Difficulty of Voting Dissuades Even the Most Motivated Voters." *Political Analysis* 17:435–455.
- Meier, Kenneth J. 1975. "Representative Bureaucracy: An Empirical Analysis." *American Political Science Review* pp. 526–542.
- Mettler, Suzanne and Joe Soss. 2004. "The Consequences of Public Policy for Democratic Citizenship: Bridging Policy Studies and Mass Politics." *Perspectives on Politics* 2(1):55–73.
- Mladenka, Kenneth R. 1981. "Citizen Demands and Urban Services: The Distribution of Bureaucratic Response in Chicago and Houston." *American Journal of Political Science* 25(4):693–714.

- Mycoff, Jason D., Michael W. Wagner and David C. Wilson. 2009. "The Empirical Effects of Voter-ID Laws: Present or Absent?" *PS: Political Science & Politics* 42(1):121–126.
- National Center for Health Statistics, Centers for Disease Control and Prevention. 2007. "National Health Interview Survey (NHIS), 2006 Data Release Data and Documentation."
- Pager, Devah. 2007. "The use of field experiments for studies of employment discrimination: Contributions, critiques, and directions for the future." *The Annals of the American Academy of Political and Social Science* 609(1):104–133.
- Pager, Devah, Bruce Western and Bart Bonikowski. 2009. "Discrimination in a Low-Wage Labor Market: A Field Experiment." *American Sociological Review* 74(5):777–799.
- Pager, Devah and Hana Shepherd. 2008. "The Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit, and Consumer Markets." *Annual Review of Sociology* 34(1):181–209.
- Purkiss, Sharon L. Segrest, Pamela L. Perrewé, Treena L. Gillespie, Bronston T. Mayes and Gerald R. Ferris. 2006. "Implicit sources of bias in employment interview judgments and decisions." *Organizational Behavior and Human Decision Processes* 101(2):152–167.
- Rosenbaum, Paul R. and Donald B. Rubin. 1983. "The central role of the propensity score in observational studies for causal effects." *Biometrika* 70(1):41–55.
- Ross, Stephen L., George C. Galster, Erin B. Godfrey, Seon Joo Lee and Bo Zhao. 2002. *Discrimination in Metropolitan Housing Markets: National Results from Phase I HDS 2000*. Washington, D.C.: US Department of Housing and Urban Development.
- Saez, Emmanuel. 2009. "Details Matter: The Impact of Presentation and Information on the Take-up of Financial Incentives for Retirement Saving." *American Economic Journal: Economic Policy* 1(1):204–228.
- Smith, Aaron W. 2010. "Government Online: The Internet Gives Citizens New Paths to Government Services and Information." Pew Internet & American Life Project: <http://bit.ly/18VF5Ej>.
- Smith, Aaron W., Kay L. Schlozman, Sidney Verba and David H. Brady. 2009. "The Internet and Civic Engagement." Pew Internet & American Life Project: <http://bit.ly/ZtbpzV>.
- Vercellotti, Timothy and David Andersen. 2009. "Voter-Identification Requirements and the Learning Curve." *PS: Political Science & Politics* 42(1):117–120.
- Voter Identification Requirements*. 2012. National Conference of State Legislatures: <http://bit.ly/17oyiCz> (Accessed May 2012).