Persistent Bias Among Local Election Officials *

D. Alex Hughes[1], Micah Gell-Redman[2], Charles Crabtree[3], Natarajan Krishnaswami[1], Diana Rodenberger[1] and Guillermo Monge[1]

[1]School of Information, University of California, Berkeley
[2]Department of International Affairs and Department of Health Policy & Management, University of Georgia
[3]Department of Political Science, University of Michigan

April 11, 2019

Abstract

Results of an audit study conducted during the 2016 election cycle demonstrate that bias toward Latinos observed during the 2012 election has persisted. In addition to replicating previous results, we show that Arab/Muslim Americans face an even greater barrier to communicating with local election officials, but we find no evidence of bias toward blacks. An innovation of our design allows us to measure whether emails were opened by recipients, which we argue provides a direct test of implicit discrimination. We find evidence of implicit bias toward Arab/Muslim senders only.

*The data, code, and compute environment required to replicate all analyses in this article are available at the Journal of Experimental Political Science Dataverse within the Harvard Dataverse Network, at: https://doi.org/10.7910/DVN/8E1IIM (Hughes et al., 2019). The authors are aware of no conflicts of interest regarding this research.

Racial bias that limits access to the ballot threatens basic principles of democratic equality. One potential source of bias that has received little attention are the street level bureaucrats who administer elections in the U.S. (Lipsky, 1980). An audit study conducted during the 2012 U.S. election cycle showed these local election officials responded at significantly lower rates to inquiries from voters with putatively Latino, as opposed to white, surnames (White, Nathan and Faller, 2015). In this paper we report the results of a similar audit study performed during the 2016 election cycle. We find that the previously observed bias against Latinos is persistent. We also extend the previous study by testing the effects of two racial primes other than Latino. Voters with Arab/Muslim names received responses at significantly lower rates (11 percentage points) than whites, while black voters did not.

The two primary motivations for this study are to determine whether the previous finding of bias toward Latinos stands up to replication, and to examine whether this bias extends to blacks and Arab/Muslim Americans. In spite of the ample evidence of racial disparities in political participation (Hajnal and Lee, 2011; Abrajano and Alvarez, 2010; Hajnal and Abrajano, 2015; García-Bedolla and Michelson, 2012) and in every-day life (Bertrand and Mullainathan, 2004), relatively little empirical work demonstrates the role of race in limiting access to the ballot in contemporary America (McNulty, Dowling and Ariotti, 2009), and some claims in this area have aroused skepticism (Hajnal, Lajevardi and Nielson, 2017; Grimmer et al., 2018). The pervasive discrimination that blacks face in various arenas of American politics (Butler, 2014) suggests that this group could be at risk of bias in interacting with local election officials. While there is also ample evidence of discrimination toward Arab and Muslim Americans (Gaddis and Ghoshal, 2015), this group has received comparatively less attention from scholars (Jamal and Naber, 2007; Panagopoulos, 2006). In an era of political rhetoric increasingly characterized by appeals to group identity, it is particularly important to understand how racially-motivated bias impacts the day-to-day mechanics of elections for a range of racial/ethnic groups.

To seek evidence of bias, we focus on the thousands of local-level administrators charged with conducting elections in the United States. These bureaucrats are generally capable of exercising discretion in carrying out their job duties, which include responding to inquiries about the mechanics of voting and eligibility to participate in elections. Our core contention is that in exercising such discretion, street-level bureaucrats may be consciously or unconsciously influenced by the characteristics (e.g., race or partisanship) of individuals seeking public services (Lipsky, 1980; White, Nathan and Faller, 2015).

1 Experiment Design

To determine the extent to which previously documented bias is persistent and extends to other racial groups, we conduct an email audit study of local election officials (Pager, 2003).[1] Our intended sample comprises all such officials with publicly available email addresses and the analytic sample includes 6,439 local election officials from 44 states (Figure A1).

The experimental stimulus consists of a single email sent to each local election official. All emails follow the same structure, greeting the official by name, referencing voter identification laws, and asking about the requirements to vote in the state corresponding to the official. Our design closely parallels White, Nathan and Faller (2015), but differs in that we send only messages that mention voter ID laws. Additionally, to minimize possible spillover issues, we create 27 variants of this request (SI section A4 and section A6).

Our experimental treatment is the putative identity of the email sender. In line with convention we expose officials to four distinct group identities by manipulating senders' names (Bertrand and Mullainathan, 2004; Bertrand and Duflo, 2017; Butler and Homola, 2017). Because the identities signaled in our treatments have elements which could be described as racial, ethnic, or religious, we refer to these generically as group identity treatments. To mitigate possible name effects, each group identity condition is signaled by 100 unique names. We check that the chosen names reliably prime ethnicity by conducting a manipulation check on Amazons Mechanical Turk service in which workers read sets of names and ascribe probabilities that a name belongs to a particular racial or ethnic group.[2] In total, we send 4,900 unique experimental conditions which combine variants of the contact language with treatment identities.

1.1 Treatment assignment and implementation

We block treatment assignment on logged population density, two-party vote share in the 2012 presidential election, percent African American, percent Latino, percent of households with incomes below 150 percent of the federal poverty level, and a dummy variable indicating whether a county was previously covered by Section 5 of the Voting Rights Act. Further details are provided in SI section A8.

[1] We received Human Subjects approval from the University of California, Berkeley and University Michigan Human Subjects Committees. Both committees waived the requirement of informed consent. Additional implementation details are made available in the Supplemental Information. The study design, and pre-analysis plan were registered at Evidence in Governance and Politics. Data, code, and computing environments are available at ALEX INSERT.

[2] SI section A7 describes the procedure for choosing names, and section A17 provides the complete list of names.

Within each block we assign local election officials a racial condition and message version at random.

We sent 6,235 emails the morning of October 31, 2016, one email to each election official that was a part of the study.[3] Emails were sent from a purpose-built domain, `ez-webmail.com`. Sending addresses took the form of the senders' first initial, last name, and a two-digit string between twenty and forty. To mitigate the possibility that elections officials would be suspicious of our contact, we structured the email headers so that inboxes displayed the full name of the purported voter (see Figure A1). The variety in our treatments was intended to reduce the likelihood that different offices would receive emails from identical senders. In twenty-nine of the forty-three states in our analytic sample every official received a contact from a distinct name.

One key innovation in this experiment permits the identification of whether emails were received and opened by election officials. We include a 1x1 pixel image with a unique link – commonly referred to as a tracking pixel – in the email body so that upon opening the email, most email clients loaded the image from our server and provided a positive record that the email had been opened by a particular official. This measurement permits inference about differential open-rates, a test of implicit bias we examine in subsection 2.1.

An open question in correspondence studies concerns whether observed effects are merely an artifact of differential treatment of stimulus by the internet and email infrastructure, i.e., spam filters. Through pilot testing we are able to comment on this question. Before taking steps to develop positive server reputation, no messages reached any test inboxes. However, by carefully managing our digital authentication and consulting with individuals at a digital marketing company, in pilot testing we were able to place every message, from every attempted sender, into test inboxes (see SI section A2).

The choice to contact election officials eight days before the election is designed to make our study reflective of the real constraints on individuals seeking and providing information about voting requirements. To minimize the impact of our intervention on election officials' time, the specific request contained in the email is one that would require little effort to fulfill. Using data gathered via our mailing system, we estimate that the median time to compose and send a response to our email is three minutes, six seconds. We contend that any costs borne by public officials as a result of our intervention are counterbalanced by the benefits of uncovering persistent bias in electronic communications between constituents and local election officials.

[3] We also sent two waves of pilot email, 54 on October 26, 2016; and, 146 on October 28, 2016. For details, see SI section A12.

Table 1: Response Rates by Experimental Condition

Ethnic Cue	White	Minority	Latino	Black	Arab
Response Rate (%)	61.3	56.6	58.4	61.4	50.1
Standard Error	1.21	0.71	1.23	1.21	1.25
N	1,611	4,828	1,609	1,613	1,606

Notes: The *Minority* column includes all data from the *Latino*, *Black*, and *Arab* columns. Response rates and standard errors are reported in percentage terms.

Our pre-registered analysis uses a single outcome measure, GOTRESPONSE, coded 1 if an election official replied to our email prior to election day, and 0 otherwise.[4] We do not count auto-replies, away messages, or bounces as valid replies. We further report an exploratory analyses of a novel outcome measure made possible through our engineering: whether a local election official opened the message.

2 Results

Overall, 57.8 percent of the emails we sent received at least one reply from local election officials. While lower than the 67.7 percent response rate previously obtained from a similar sample (White, Nathan and Faller, 2015), this rate compares favorably with experiments on elected officials in the U.S., suggesting that our requests were taken at face value (Butler and Broockman, 2011).

Election officials respond at considerably lower rates when queries come from minority as opposed to white senders (difference in mean, $\Delta\mu = -4.70$ percentage points, *Wilcox Rank-Sum* $P < 2 \times 10^{-16}$). However, as we report in Table 1 responsiveness to minority senders is not uniformly lower. Nonparametric tests using white senders as the baseline find that a Latino name is sufficient to suppress the likelihood of a response by nearly 3 percentage points ($\Delta\mu = -2.97$, $P = 0.07$). Strikingly, an Arab/Muslim name lowers the likelihood of a response by greater than 11 percentage points ($\Delta\mu = -11.3$, $P < 1 \times 10^{-10}$). In contrast, black senders receive responses at a rate indistinguishable from white senders ($\Delta\mu = 0.11$, $P = 0.90$). Figure 1 (a) plots the Intent to Treat (ITT) causal effects of our treatments. Regression estimates with robust standard errors are reported in columns 1 and 2 of Table A6, and produce similar results.

Figure 1 (b) plots a precision weighted meta-analysis estimate (Gerber and

[4]Pre-analysis Plan Registered at EGAP (Hughes, Gell-Redman and Crabtree, 2016).

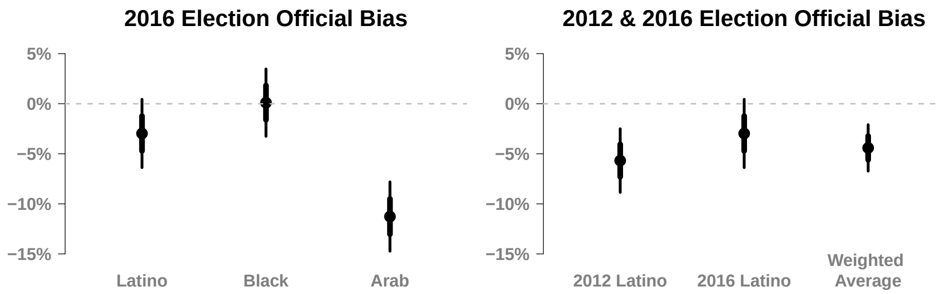

Figure 1: Points represent the ITT, the estimated difference in response rates to emails from the named identity, compared to the white response rate baseline. Thick bars report $ITT \pm SE$, thin bars report $ITT \pm 1.96 \times SE$. All estimates are difference in means, except the *Weighted Average* which estimates a precision weighted difference (Gerber and Green, 2012) utilizing 2012 (White, Nathan and Faller, 2015) and 2016 Latino evidence.

Green, 2012, p. 361) that combines the results of our intervention with those previously reported (White, Nathan and Faller, 2015). These data, gathered in independent audits conducted over two election cycles, show that Latinos receive replies from local election officials at a rate 4.4 percentage points lower than whites ($\Delta\mu = 4.4$, precision weighted $SE = 1.18$, $P < 0.0001$).

While the persistence of the treatment of Latino senders in the 2012 and 2016 elections is remarkable, perhaps more striking is the finding that Arab/Muslim names suffer a penalty more than two times greater than the one produced by a Latino stimulus. One potential concern is that the observed effect could be driven by the implausibility of the treatment, since many parts of the country do not have any appreciable population of Arab-Americans. To examine this possibility, we investigate whether treatment effects are smaller in the jurisdictions where Arab-Americans are more numerous. If treatment effects are driven by implausibility then they should be smaller in places where the presence of citizens with Arab names are more plausible. We do not find clear evidence that the proportion of Arab Americans moderates the treatment effect (Table A13, Model 3; Table A14; Table A15). Our most credible estimates find a 10.6 percentage point bias against Arab senders in counties with no Arab population ($\Delta\mu = -10.6$ percentage points, $SE = 2.5, P < 0.001$), but only a 2.6 percentage point improvement in the highest Arab population quartile of counties ($\delta\Delta\mu = +2.6$ percentage points, $SE = 4.4, P = 0.55$), although the distribution of Arab American settlement limits the strength of this robustness check.[5]

[5] In the highest Arab quartile, the mean Arab population is 1%.

2.1 Evidence of implicit discrimination

Local election officials who receive our intervention demonstrate bias insofar as they respond differentially based only on the signal of group identity delivered through our treatments. This observed response behavior is part of a chain of actions: the official must open, read, and then respond to the email. Standard analyses of audit experiments, which report an indicator of response or non-response as the dependent variable, focus only on the final result of this compound process. Innovations of our design allow us to consider the outcome at a prior step, the decision by the official to open the received email, conditional on the treatment delivered.

To respond to our experimental stimulus, an election official must identify our request from among the large number of other requests, categorize it mentally, and then open it. We argue that opening an email is a high-volume, low-attention task of the type scholars have associated with implicit, rather than explicit bias (Devine, 1989; Bertrand, Chugh and Mullainathan, 2005, p.96). The pattern of email opens suggests that, indeed, elections officials may be unintentionally or automatically screening requests from Arab/Muslim senders. There is no difference in open rates between white and latino names ($\Delta\mu = -0.74$ percentage points, $SE = 1.7$, $P = 0.68$) or white and black names ($\Delta\mu = -0.24$, $SE = 1.7$, $P = 0.90$). However, there is a pronounced gap in open rates for emails sent by senders with Arab/Muslim names, who have their emails opened at a rate 6.8 percentage points lower than white senders ($\Delta\mu = -6.8$, $SE = 1.8$, $P = 0.00013$).

2.2 Awareness of experiment

During the analysis phase of this project, it came to the researchers' attention another entity was pursuing a similar line of research using the same sending domain as White, Nathan and Faller (2015). As a result, some public officials became concerned that an audit study might be underway. News reports claim that these concerns prompted the National Association of Secretaries of State (NASS) to alert its state branches, who in turn had the opportunity to alert individual officials. In sum, some of our experimental subjects may have become aware of the presence of interventions.

Subjects' awareness of the intervention poses a general threat to audit studies, either by compromising independence between units, or by violating the exclusion restriction if minority names are more likely to raise suspicion than white names. Because subjects' awareness might prevent identification of causal effects, researchers should mitigate this risk by using many identities and a well-tuned sending architecture whenever feasible. When there is any observable information about the possibility of discovery, researchers can use this information to evaluate

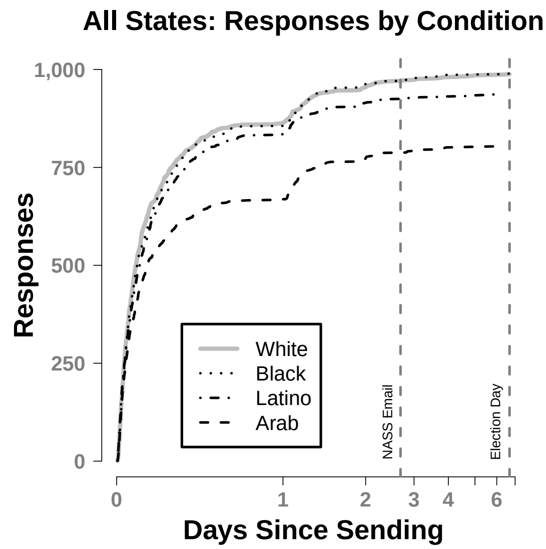

Figure 2: Rapidly slowing rates of response. The vertical axis plots the cumulative number of respones, split by group identity of sender; the horizontal axis plots time since sending. Election Day and NASS emails are noted with vertical dashed lines. Responses follow a clear diurnal rhythm, and patterns of bias appear rapidly.

whether apparent differences are likely the result of discovery.

Analysis of the timing of responses in this experiment does not suggest that discovery is leading to the observed results. First, as we present in Figure 2, the systematic pattern of unresponsiveness to minority names appears rapidly and well before the reported NASS broadcast. Second, as we report in Table A11 and Table A12, models that censor response data at the time of the NASS broadcast, and models that exclude states that witnessed interference between units both produce estimates very similar to our main results.

3 Conclusion

Previous experimental evidence showed local election officials were less responsive to inquiries from Latinos, raising concerns about bias in the electoral process. Using a similar experimental design, we demonstrate the firm basis for these concerns by replicating the initial finding. We also extend the results by testing for bias against other groups.

Our intervention showed Arab/Muslim Americans to be markedly disadvantaged in their interactions with local election officials. This finding is particularly salient given that it is not simply an artifact of Arab/Muslims being a relatively less numerous part of the electorate. We encountered no evidence of bias from local election officials toward African Americans, making ours at least the third recent study to produce a similarly unexpected null finding (Einstein and Glick, 2017; Gell-Redman et al., 2018). Rather than evidence of a lack of bias against African Americans, these null findings may be an artifact of the correspondence study method in which name alone, rather than other cues such as appearance, is used to signal identity.

Through this design, we also engage a challenge inherent to all audit studies, the risk that subjects become aware of the experiment. The relatively low technical sophistication required to conduct some forms of audit studies, mated with the potentially large sample size that is possible through email-based audits make these designs a potentially attractive way to identify discriminatory behavior. However, in an increasingly crowded field, researchers must face the possibility that experimental subjects become aware of the study, thereby damaging the inference. We determined that sending 4,900 distinct treatments on a custom-built server provided the best balance of a low possibility of discovery with the ability to identify a novel open rate outcome measure, and we would encourage future researchers to make a similar assessment.

References

Abrajano, Marisa A. and Michael M. Alvarez. 2010. *New Faces, New Voices: The Hispanic Electorate in America*. Princeton University Press.

Bertrand, M. and E. Duflo. 2017. Field Experiments on Discrimination. In *Handbook of Field Experiments*, ed. Abhijit Vinayak Banerjee and Esther Duflo. Vol. 1 North-Holland pp. 309 – 393.
URL: *http://www.sciencedirect.com/science/article/pii/S2214658X1630006X*

Bertrand, Marianne, Dolly Chugh and Sendhil Mullainathan. 2005. "Implicit discrimination." *American Economic Review* pp. 94–98.

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.

Butler, Daniel M. 2014. *Representing the Advantaged: How Politicians Reinforce Inequality*. Cambridge University Press.

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.

Butler, Daniel M and Jonathan Homola. 2017. "An Empirical Justification for the Use of Racially Distinctive Names to Signal Race in Experiments." *Political Analysis* 25(1):122–130.

Devine, Patricia G. 1989. "Stereotypes and prejudice: Their automatic and controlled components." *Journal of personality and social psychology* 56(1):5.

Einstein, Katherine Levine and David M. Glick. 2017. "Does Race Affect Access to Government Services? An Experiment Exploring Street-Level Bureaucrats and Access to Public Housing." *American Journal of Political Science* 61:100–116.

Gaddis, S Michael and Raj Ghoshal. 2015. "Arab American Housing Discrimination, Ethnic Competition, and the Contact Hypothesis." *The Annals of the American Academy of Political and Social Science* 660(1):282–299.

García-Bedolla, Lisa and Melissa R. Michelson. 2012. *Mobilizing Inclusion: Transforming the Electorate Through Get-out-the-Vote Campaigns*. Yale University Press.

Gell-Redman, Micah, Neil Visalvanich, Charles Crabtree and Christopher Fariss. 2018. "It's all about race: How state legislators respond to immigrant constituents." *Political Research Quarterly* .

Gerber, Alan S and Donald P Green. 2012. *Field experiments: Design, analysis, and interpretation*. WW Norton.

Grimmer, Justin, Eitan Hersh, Marc Meredith, Jonathan Mummolo and Clayton Nall. 2018. "Obstacles to estimating voter ID laws' effect on turnout." *Journal of Politics* 80(3).

Hajnal, Zoltan and Marisa Abrajano. 2015. *White Backlash: Immigration, Race, and American Politics*. Princeton University Press.

Hajnal, Zoltan, Nazita Lajevardi and Lindsay Nielson. 2017. "Voter Identification Laws and the Suppression of Minority Votes." *The Journal of Politics* 79(2):363–379.

Hajnal, Zoltan and T. Lee. 2011. *Why Americans Don't Join the Party: Race, Immigration, and the Failure (of Political Parties) to Engage the Electorate*. Princeton University Press.

Hughes, D. Alex, Micah Gell-Redman and Charles Crabtree. 2016. "Who gets to vote?" Evidence in Government and Politics, EGAP ID: 20161001AA.
URL: *http://egap.org/registration/2183*

Hughes, D. Alex, Micah Gell-Redman, Charles Crabtree, Natarajan Krishnaswami, Guillermo Monge and Diana Rodenberger. 2019. "Replication Data for: Persistent Bias Among Local Election Officials." *Harvard Dataverse* doi:10.7910/DVN/8E1IIM.

Jamal, Amaney and Nadine Naber. 2007. *Race and Arab Americans Before and After 9/11: From Invisible Citizens to Visible Subjects*. Syracuse University Press.

Lipsky, Michael. 1980. *Street-level Bureaucracy: Dilemmas of the Individual in Public Services*. Russell Sage.

McNulty, John E., Conor M. Dowling and Margaret H. Ariotti. 2009. "Driving Saints to Sin: How Increasing the Difficulty of Voting Dissuades Even the Most Motivated Voters." *Political Analysis* 17(4):435–455.

Pager, Devah. 2003. "The mark of a criminal record." *American journal of sociology* 108(5):937–975.

Panagopoulos, Costas. 2006. "The Polls-Trends: Arab and Muslim Americans and Islam in the Aftermath of 9/11." *Public Opinion Quarterly* 70(4):608–624.

White, Ariel R., Noah L. Nathan and Julie K. Faller. 2015. "What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials." *American Political Science Review* 109(1):129–142.