

Giulietti, Corrado; Tonin, Mirco; Vlassopoulos, Michael

Working Paper

Racial Discrimination in Local Public Services: A Field Experiment in the US

IZA Discussion Papers, No. 9290

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Giulietti, Corrado; Tonin, Mirco; Vlassopoulos, Michael (2015) : Racial Discrimination in Local Public Services: A Field Experiment in the US, IZA Discussion Papers, No. 9290, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/120940>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 9290

**Racial Discrimination in Local Public Services:
A Field Experiment in the US**

Corrado Giulietti
Mirco Tonin
Michael Vlassopoulos

August 2015

Racial Discrimination in Local Public Services: A Field Experiment in the US

Corrado Giulietti

IZA

Mirco Tonin

*University of Southampton,
Central European University and IZA*

Michael Vlassopoulos

*University of Southampton
and IZA*

Discussion Paper No. 9290
August 2015

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Racial Discrimination in Local Public Services: A Field Experiment in the US*

Discrimination in access to public services can act as a major obstacle towards addressing racial inequality. We examine whether racial discrimination exists in access to a wide spectrum of public services in the US. We carry out an email correspondence study in which we pose simple queries to more than 19,000 local public service providers. We find that emails are less likely to receive a response if signed by a black-sounding name compared to a white-sounding name. Given a response rate of 72% for white senders, emails from putatively black senders are almost 4 percentage points less likely to receive an answer. We also find that responses to queries coming from black names are less likely to have a cordial tone. Further tests demonstrate that the differential in the likelihood of answering is due to animus towards blacks rather than inferring socioeconomic status from race.

JEL Classification: D73, H41, J15

Keywords: discrimination, public services provision, school districts, libraries, sheriffs, field experiment, correspondence study

Corresponding author:

Mirco Tonin
Economics Department
School of Social Sciences
University of Southampton
Southampton SO17 1BJ
United Kingdom
E-mail: m.tonin@soton.ac.uk

* We would like to thank Raj Chetty, Rajeev Dehejia, Guillermina Jasso, Gary King, Deniele Paserman, Uwe Sunde and Adam Szeidl for useful discussions. We also thank participants of seminars at IZA, Central European University, Temple, Köln and participants of the 2015 Society for Government Economists Conference for their comments. This project has received ethical approval from the Institute for the Study of Labor (IZA) in Bonn, Germany, and from the University of Southampton.

1 Introduction

Blacks have a disadvantaged position in American society in terms of labor market outcomes, educational achievements, incarceration, health and life expectancy.¹ Discrimination is commonly proposed as one of the possible causes of this predicament and has been documented in several realms, including the labor market, the judicial system, housing and product markets.² In his review of racial inequality, Fryer (2011) underlines that “the new challenge is to understand the obstacles undermining the achievement of black and Hispanic children in primary and secondary school” (page 857). Local public service providers like school districts and libraries have a major role to play in this regard; thus, discrimination in access to these services represents an important obstacle towards addressing racial inequality. More generally, public institutions at the local level – so-called street-level bureaucracies (Lipsky, 1980) – are at the front-line of service delivery and thus play a key role in the policy-implementation process, exerting great influence on how policies are actually carried out and experienced by citizens. It is hence important to examine their attitudes and behavior vis-à-vis discrimination.

A central tenet of U.S. law is the prohibition of racial discrimination by the government, with racial discrimination by public authorities prohibited and the principle of non-discrimination central to governmental policy throughout the country (US Government, 2013). For instance, the Civil Rights Act of 1964 bans discrimination based on race by government agencies that receive federal funding. This is supplemented by several other provisions in federal and state law. For instance, under the Minnesota Human Rights Act (363A.12), “It is an unfair discriminatory practice to discriminate against any person in the access to, admission to, full utilization of or benefit from any public service because of race [...]”. Thus, discrimination by providers of public services not only has a potentially detrimental impact on the economic and social lives of those affected, but is also illegal. Furthermore, taste-based discrimination à la Becker (1957) is predicted to fade with intensified market competition and lower barriers to entry. While deregulation and globalization may have increased competition in the US economy, thus placing pressure on discriminatory attitudes in the private sector (see Levine et al., 2014, for relevant evidence from the financial industry), this has certainly been much less the case for the public sector.

In this study, we investigate racial discrimination across a wide range of public services. In particular, we compile all available emails of the targeted local public service providers, which gives us more than 19,000 cases, corresponding to roughly half of the total number of providers. Targeted services include school districts, local libraries, sheriff offices, county clerks, county treasurers and

¹Altonji and Blank (1999) provide an overview of race differential in the labor market. Fryer (2011) focuses on the racial achievement gap in education. Sabol et al. (2009) provide figures about incarceration by race. CDC (2011) reports on race disparities in mortality and morbidity.

²Charles and Guryan (2011) discuss research on discrimination against blacks in labor market outcomes. Alesina and La Ferrara (2014) show evidence of racial bias in capital sentencing. Ewens et al. (2014) is a recent contribution on discrimination in housing. Product market discrimination is studied, for instance, by Doleac and Stein (2013).

job centers. Some of these services play important roles in relevant economic domains as they directly relate to employment (job centers) or education (school districts). Libraries also perform an important role by both promoting literacy and providing access to information and computer technology, thus facilitating activities like job-searching.³ The other services that we study involve typical government functions, like law enforcement (sheriffs), taxation (county treasurers) and general public administration (county clerks).

To identify in a credible way whether there is racial discrimination in access to local public services, we conduct an email correspondence study. In particular, we solicit information relevant to access a public service (the office opening hours or some specific information, e.g., the documentation needed for school enrollment) from 19,079 local public offices and observe whether or not we receive a reply depending on whether the request was signed with a distinctively white or black name. This methodology has been used to investigate discrimination in a variety of settings, including employment (Bertrand and Mullainathan, 2004), housing (Ewens et al., 2014), product markets (Gneezy et al., 2012; Doleac and Stein, 2013), financial markets (Bayer et al., 2014) and along different dimensions, including race, ethnicity, gender, age, disability, sexual orientation, obesity, caste and religion.⁴

Failing to provide information about a service is not equivalent to denying access to a service. However, there is growing evidence showing that the provision of information has an important impact on decisions and take-up rates. For instance, regarding the Earned Income Tax Credit, Bhargava and Manoli (2015) show that “the mere receipt of a second notice just months after the receipt of an initial IRS notice causes 0.22 of the non-respondents to take-up. Hoxby et al. (2013) show that providing information on the application process and colleges’ net costs has an effect on college enrollment. Hastings and Weinstein (2008) find that providing low-income families with direct information about school-level academic performances has an impact on parents’ school choice. Duflo and Saez (2003) show the effects of information on the decision to enroll in a tax deferred account retirement plan. Daponte et al. (1999) document that ignorance about the food stamp program contributes to nonparticipation. Thus, making it more difficult for a citizen to obtain information about a service is not merely a nuisance, but can also have an important impact on whether and how the citizen engages with the service. Moreover, experiencing even relatively small episodes of discrimination in a specific domain may erode the trust that an individual has

³A nationally representative survey by the Pew Research Center (2013) finds that over half of Americans aged 16 and older have used a public library in some way in the previous 12 months, with many using facilities provided by libraries for purposes related to education (42 %, e.g., taking online classes or working on assignments and schoolwork), employment (40 %, e.g., search for job opportunities, submission of online job applications or work on resumes), and health (37 %, e.g., learning about medical conditions, finding health care providers, and assessing health insurance options). Many also report using a library computer to download government forms or find out about a government program or service. Interestingly, the study shows that library services are particularly important to “[w]omen, African-Americans and Hispanics, adults who live in lower-income households, and adults with lower levels of educational attainment”.

⁴See Riach and Rich (2002), Guryan and Charles (2013) and Rich (2014), for earlier and more recent reviews of the literature.

in government institutions in general, potentially leading to the development of an “oppositional culture” with negative consequences, for instance, in terms of educational achievement (Fryer and Torelli, 2010). Furthermore, the medical and psychological literature provides ample evidence of the negative effect of discrimination on physical and mental health (Harrell et al., 2003), including so-called racial micro-aggressions, i.e., subtle everyday experiences of racism (Wong et al., 2014). Finally, it seems implausible that a discriminatory attitude would only express itself in a very specific element of the service delivery, without having a more general impact. In other words, a librarian not replying to requests for information coming from blacks is also likely to treat blacks differently in other aspects of the service, e.g., by being less forthcoming when asked about a certain library resource. Thus, our measure of discrimination is informative about the general attitude permeating local public services.

Our results show that emails signed with a distinctively black name are less likely to receive a reply than identical emails signed with a distinctively white name, thus indicating the presence of discrimination in access to public services. In terms of magnitude, given a response rate of almost 72% for white senders, emails from putatively black senders are almost 4 percentage points less likely to receive an answer. This differential response is particularly strong among sheriff offices, but is also present in libraries and school districts. We also find that responses to inquiries coming from black names are less likely to have a cordial tone. Our analysis points to the fact that the differential in the likelihood of answering is due to animus towards blacks, rather than a form of statistical discrimination arising from assigning low socioeconomic status to a sender with a distinctively black name. In particular, we deploy two approaches, whereby the first entails predicting the race of the recipient and checking whether black recipients are more likely to respond to emails signed by black senders, as one would expect if taste-based discrimination by white recipients were at play. In the second approach, we attempt to directly fix the socioeconomic status of the sender by signaling his occupation in the email. Both approaches indicate that taste-based discrimination is the primary driver of the race gap in our response rate.

These results are consistent with other studies that uncover evidence of discrimination in public services, mostly involving various aspects of law enforcement. For instance, Alesina and La Ferrara (2014) find evidence consistent with the presence of racial prejudice in capital sentencing, driven exclusively by Southern states. Glaeser and Sacerdote (2003) look at vehicular homicides and find that drivers who kill blacks receive significantly shorter sentences. Abrams et al. (2012) find support for the hypothesis that some judges treat defendants differently based upon their race. A recent study by political scientists regarding discrimination in the electoral process (White et al., 2015) finds that emails about voting requirements sent to over 7,000 local U.S. election officials from Latino aliases are significantly less likely to receive responses and, if granted, to receive responses of lower quality than those sent from non-Latino white aliases. Thee related study by Butler and Broockman (2011) – albeit focusing on lawmakers rather than bureaucrats – involved

sending emails asking for help with registering to vote to almost 5,000 U.S. state legislators. They find that putatively black requests receive fewer replies than requests coming from white aliases, even when the email signaled the sender’s partisan preference. Furthermore, Distelhorst and Hou (2014) find discriminatory behavior against ethnic Muslims by unelected public officials in China. However, some studies do not find any bias against blacks, such as Coviello and Persico (2013) on NYPD’s “Stop and Frisk Program”. Our study is the first to explore discrimination in a variety of local public services that perform important functions and constitute the majority of interactions between government institutions and citizens. The fact that we find evidence of discrimination has important implications for public policy, which we will discuss in the conclusions, after presenting the experimental set up in the next section and the results in sections 3 and 4.

2 The Field Experiment

The field experiment – conducted in March/April 2015 – entailed us sending email queries to over 19,000 local public offices, signing the emails with names that strongly evoke the race of the sender (white or black). In what follows, we describe the procedures surrounding the selection of public services, email queries and names of sender, as well as the experimental design.

2.1 Type of Public Services, Emails and Names of Senders

The first step is to select which public services to target. We chose public services according to two criteria: (i) the provision of the service is at the county or sub-county level to ensure a large number of observations and broad geographic coverage; and (ii) email addresses are publicly available or a directory of email addresses could be obtained. We came up with six types of public services that span a variety of local public services: sheriff offices, local libraries, school districts, county clerks, county treasurers and job centers. We were able to obtain over 21,000 email addresses and finally use 19,079 valid ones (the sources of email addresses are listed in Table B1 in the Appendix).⁵ This constitutes our target population. Table 1 presents the breakdown of numbers and shares of emails by type of public service. The three most numerous public services are school districts, libraries and sheriff offices, which jointly account for almost 90% of the emails sent. The emails used in the field experiment account for nearly 50% of all existing potential recipients (Table A1 reports the detailed number of existing recipients and of emails in the sample).

Figure 1 illustrates the geographic coverage and dispersion of our field experiment. It is evident that more populated counties and regions – which hence have a larger number of available recipients (such as the North-East) – receive a relatively high number of emails.


⁵A small random sample was used for testing; about 2,000 emails were eliminated because they were either undelivered or caught by anti-spam software. Indeed, we checked and corroborated that the probability that an email is eliminated does not correlate with our key variables.

Table 1: Emails by Type of Service

	Number	Percentage
School D.	9,873	51.75
Library	4,894	25.65
Sheriff	1,836	9.62
Treasurer	1,129	5.92
Job Center V.R.	731	3.83
County Clerk	616	3.23
Total	19,079	100

Source: own computations.

Figure 1: Location of recipients




Notes: Each pin represents the location (e.g., city) where recipients are located and might represent more than one observation.

Figure 2 plots the share of emails against the share of recipients across states. As can be seen, most observations are clustered around the 45-degree line, suggesting that the number of email addresses is proportional to the number of available recipients. We will account for any discrepancies in one of our robustness checks.

Our emails contain simple queries that were chosen not to impose significant effort/investment on the recipients' part. Specifically, we use two types of email: simple and complex. Simple emails ask about the opening hours of the office, while complex emails ask for some additional yet basic information that an ordinary citizen might need to know to carry out a transaction with the office. Every email has the following format:

Figure 2: Sample representativeness



Share of emails corresponds to the number of emails in a state over the total number of emails in the US; share of recipients corresponds to the number of recipients in the state over the total number of US recipients.

From: [Black/White Name]
Subject: [Opening Hours] or [Inquiry]

Hi,
My name is [Black/White Name] and I live in [City Name].
[Simple Query/Complex Query]

Thank you,
[Black/White Name]

To make the name of the sender as noticeable and salient as possible, we show it three times: in the sender field, the main body and the signature. The complete set of questions are presented in Table B2 in the Appendix.

Names of senders were chosen to evoke race as much as possible. We use two distinctively white names (Jake Mueller and Greg Walsh) and two distinctively black names (DeShawn Jackson and Tyrone Washington). Both the first names and surnames of our chosen names are among the most recognizable black and white names and have been previously used in similar correspondence studies (Bertrand and Mullainathan, 2004; Butler and Broockman, 2011; Broockman, 2013; White et al., 2015).

We created four email addresses, with the local part comprising two letters and six numbers and the domain part being gmail.com, corresponding with the four chosen names. In each case, the display name of the email sender was the sender's full name.

2.2 Experimental Design

We sent the emails over a period of two weeks due to limits in the number of emails that can be sent daily. Emails are differentiated by the race of the person who signs it (white or black) and the type of query that it contains (simple or complex). This gives rise to a 2x2 research design with four treatments that correspond to the four possible pairs of race/email complexity. In the analysis, we pool the two black and two white names. We randomize the treatments at the state level and for each type of public service separately.

Table 2 shows summary statistics of various county characteristics of recipients broken down by whether the email that they received was signed by a distinctively white or black name. As can be seen, our sample is balanced across all of these characteristics (the data sources of these characteristics are presented in Table B3 in the Appendix).

Six weeks after the first set of emails were sent to all recipients, we sent a second wave of emails.

Table 2: County characteristics of email recipients

	Black	White	t-test (p-value)
% of black among employed	6.1 (10)	6.1 (10.2)	0.83
% of black in public sector	10.9 (10.6)	10.4 (10.3)	0.05
Unemployment rate (%)	6.1 (2.1)	6.1 (2.1)	0.82
% of hispanic	10.1 (13.9)	9.8 (13.8)	0.27
Average labor income (USD)	794 (239)	789 (237)	0.17
Crime rate (%)	2.5 (1.3)	2.4 (1.3)	0.47
% of Dem votes	43.0 (15.1)	43.0 (15)	0.67
Urbanization	71.8 (45)	70.7 (45.5)	0.09

Source: own computations. Standard deviations in parentheses.

The structure of the email was the same as in the first wave, aside from modifying the signature as illustrated below:

```

*****
[Black/White Name]
Real Estate Agent
Buy - Sell - Rent
*****

```

The purpose of this is to fix the emails recipients' perceptions about the socioeconomic background of the sender. In the second wave, we randomized the race of the sender and changed the email type, whereby those recipients who received a simple email in the first wave were sent a complex one in the second wave and vice-versa. To avoid any suspicion that may arise from receiving two emails from the same person, we used a different name within the same race for those cases where a recipient was randomized to receive an email signed by the same race in both waves.

Our main outcome is whether the email is answered. As a secondary outcome, we also track the time elapsed between when the query was sent and when the response was received. Furthermore, we also investigate a measure of the quality of response by focusing on the number of responses received, the length of the response (number of words) and how cordial the response is.

3 Results

3.1 Descriptive Statistics

Overall, about 70 percent of the 19,079 emails that we sent received a response (see Table A2 for detailed statistics). This indicates that public service providers are generally quite responsive to queries coming from the public, despite a non-negligible share of them going unanswered. The response rate was 68.9% for simple emails and 70.8% for complex emails, which is surprising given that complex emails seemingly require more effort from the recipient. A possible explanation could be that responders may consider the information solicited by simple emails (i.e., opening times) to be easily available from various sources and thus they feel less compelled to provide an answer to such a query.

Emails signed by white-sounding names (we will refer to them as “white emails” hereafter) receive a response in 71.67% of the cases, while those signed by black-sounding names (henceforth: “black emails”) in only 67.96% of the cases, with the difference of 3.7 percentage points being strongly statistically significant (z-stat p-value <0.000). The response rate to emails coming from the two white-sounding names is almost identical (71.76% for Jake Mueller and 71.57% for Greg Walsh; z-stat, p-value 0.84), while there is a difference between the two black-sounding names (69.05% for DeShawn Jackson and 66.91% for Tyrone Washington; z-stat, p-value 0.03). Given that both first names are among the most recognizable black names, according to Fryer and Levitt (2004), it is possible that this difference emerges because one of the last names has a stronger association with black people than the other. Indeed, among those persons who are called Jackson, 53.02% are black and 41.93% are white, while for Washington the figures are 89.87% and 5.16%, respectively. In both cases, the response rate is significantly lower than that for white emails, with p-values <0.00. Hence, hereafter we will consider the difference between white and black emails without distinguishing between the two names within each category.

The descriptive evidence above indicates considerable racial differences in the response rate to the emails.

3.2 Main Results

Next, we examine whether there are racial differences in response rates in a regression framework, which allows us to control for various factors such as the type of public service, state fixed effects and several county characteristics. Specifically, we estimate linear probability regressions of the form:

$$Response_i = \beta' \sum ServiceType_i + \gamma ComplexEmail_i + \delta Black_i + X_i' \theta + s + d + \epsilon_i, \quad (1)$$

where *Response* is a binary variable indicating whether a response to the email was provided, *ServiceType* indicates the type of public service to which the email was sent (six types), *ComplexEmail* is a binary variable indicating whether the email was simple or complex, *Black* is a binary variable indicating whether the email was signed by a distinctively black name, X is a vector of county level characteristics that we use as controls, s represents state fixed effects and d are indicators for the calendar days when emails were sent out. Standard errors are clustered at the state/public service type level.

The main coefficient of interest in these regressions is δ , which tells us whether there is a differential response according to the racial identity of the sender.

Table 3 summarizes the main regression results. Column I includes state, public service type and sending day fixed effects. The estimated racial gap in response rates – at 3.8 percentage points – does not substantially differ from that emerging from the raw comparison reported in the previous sub-section (3.7 percentage points). Column II adds a dummy variable that takes the value of one if the email question is complex. In line with the raw comparison, complex emails are more likely (1.8 percentage points) to receive a response than simple emails. In column III, we examine whether the differential response rate between white and black emails varies according to complexity by adding an interaction term between black name and the complex email dummy. The estimate of the interaction term proves to be small and statistically insignificant, indicating that the differential in the response rate is not specific to the nature of the query. In column IV, which represents our baseline specification, we include various county level characteristics (unemployment rate, average wage, share of hispanic population, crime rate, share of democratic votes, rural/urban dummy). Unsurprisingly, since the emails are randomly assigned, we find that the inclusion of these controls does not change the racial difference in response rate estimated in column I. Finally, in column V we exploit the second wave of emails and particularly the fact that half of the recipients receive emails from senders with different races across the two waves. We can subsequently estimate a model with email fixed effects (and calendar day fixed effects). The within-recipient variation in the responsiveness to white and black emails is similar to that estimated in column I (3.2 percentage points).

As mentioned in the experimental design section, the share of emails for each state does not perfectly match the share of potential recipients in each state. Despite generally not being very large, this discrepancy makes some states under-represented and others over-represented. To correct for this, in unreported regressions, we have reestimated the model in column IV of Table 3 by weighting observations by the ratio of the number of recipients in each state to the number of emails sent in the state. The estimate (-0.036, s.e. 0.007) is remarkably close to that of the baseline specification. Furthermore, we checked the sensitivity of our results to the clustering of the standard errors. Clustering at a level other than state/public service type does not affect the precision of our

Table 3: Difference in response rates

	I	II	III	IV	V
Black	-0.038*** (0.006)	-0.038*** (0.006)	-0.038*** (0.009)	-0.038*** (0.006)	-0.032*** (0.005)
Complex		0.018*** (0.007)	0.019** (0.008)	0.018*** (0.007)	0.029*** (0.004)
Black \times Complex			-0.002 (0.013)		
R^2	0.045	0.045	0.045	0.050	0.023
N	19,079	19,079	19,079	19,079	38,168
State/Service/Date F.E.	Y	Y	Y	Y	Y
County controls	N	N	N	Y	N
Recipient F.E.	N	N	N	N	Y

Robust standard errors in parentheses clustered at the state/public service type level. Dependent variable is a binary variable indicating whether a response to the email was provided (linear probability model). Average response rate is 0.698.

County controls are: unemployment rate, average wage, share of hispanic in the population, crime rate, share of votes to democrats in presidential elections, and rural/urban counties.

Recipient fixed-effects refers to a regression which uses data from the two waves. In the model in column V, R^2 represents the within R^2 .

estimates. In particular, if we were to cluster at the state/public service type/sender name level, the standard error of the benchmark model would be 0.008, whereas if we were to cluster at the county/public service type level, the standard error would be 0.006.⁶

3.3 Type of Public Service

It is important to recall that our sample comprises six different public services with different sizes in the sample due to a combination of differences in how many of them are present in the country and in email availability. One could worry that our results might be driven by one particular type of public service. Hence, we analyze the results by type of service (for reference, Table A3 in the Appendix contains detailed summary statistics). When considering the response rates, the pattern of higher response rates for white emails holds in all cases, except for job centers, for which the racial difference is not statistically significant (p-value=0.73). For school districts, libraries and sheriff offices (the public services with the largest number of observations), the difference in response rates is statistically significant. The magnitude of the racial gap ranges between 3.4 and 7.0 percentage points. For instance, this means that a white sender has almost a 15 percent higher likelihood of receiving a reply to an email sent to a sheriff office than a black sender. For treasurers, the difference is marginal (p-value=0.11), while for county clerks (the smallest group) the difference is not statistically significant.

⁶A final check that we perform concerns the functional form. We estimated the benchmark specification using probit model. The marginal effects (-0.039, se 0.006) are remarkably similar to the estimates of the linear probability model, reassuring us that estimates are not sensitive to the chosen functional form.

In Table 4, we estimate the econometric model in Eq. (1) for each type of public service. The results essentially confirm the patterns of the descriptive statistics, with estimates only being statistically significant for school districts, libraries and sheriff offices and the largest racial difference found in the latter group.

In additional unreported analysis, we estimated a regression model that attributes equal importance to each service. We achieve this by weighting observations by the ratio of the total number of emails sent to the number of emails of each type of public service. The estimated coefficient (-0.032, s.e. 0.012) is not too dissimilar from that of the unweighted regressions, suggesting that the differential treatment between black and white emails is robust to giving equal weight to each of the six services.

Table 4: Type of public service

	School D.	Library	Sheriff	Treasurer	Job Center	County Clerk
Black	-0.035*** (0.009)	-0.041*** (0.010)	-0.074*** (0.021)	-0.039 (0.029)	0.008 (0.031)	-0.014 (0.048)
\bar{Y}	0.748	0.670	0.498	0.718	0.725	0.649
R^2	0.041	0.047	0.113	0.085	0.164	0.093
N	9,873	4,894	1,836	1,129	731	616

Robust standard errors in parenthesis clustered at the state level.

Dependent variable is a binary variable indicating whether a response to the email was provided (linear probability model).

All regressions include controls of col IV of Table 3.

3.4 Geographic Heterogeneity

Racial disparities might not be equally distributed across the U.S. For example, recent evidence from Stephens-Davidowitz (2014) shows that Google search queries with racially charged language are particularly intense in Southern states. We therefore explore whether there is geographic heterogeneity in the racial difference in the response rate. For this purpose, we split our sample into the four regions defined by the Census Bureau (North-East, Mid-West, South and West) and estimate our baseline specification on each subsample.⁷ Table 5 summarizes the results. We find a significant racial gap in all four regions, with the estimate ranging from 2.6% in the North-East to 4.9% in the Mid-West. To further probe this pattern of geographical variation, in columns V and VI, we classify counties into urban and rural and split the sample along this dimension.⁸ This gives

⁷The state composition of each region is the following: North-East includes Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont; Mid-West includes Illinois, Indiana, Michigan, Ohio, and Wisconsin, Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, and South Dakota; South includes Delaware, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia, Washington D.C., West Virginia, Alabama, Kentucky, Mississippi, Tennessee, Arkansas, Louisiana, Oklahoma, and Texas; West includes Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, Wyoming, Alaska, California, Hawaii, Oregon, and Washington.

⁸We apply the six-level classification developed by the National Center for Health Statistics (NCHS). The six categories are: Large central metro, Large fringe metro, Medium metro, Small metro, Micropolitan and Noncore. We

rise to 1,312 rural counties and 1,780 urban counties. The results indicate that the racial gap in response rate is substantially larger in rural areas, namely almost double than urban areas. This is consistent with the fact that we find a larger racial gap in the Mid-West, where the incidence of rural counties is highest.

Evidence that the differential treatment in response vis-à-vis blacks is not worse in Southern states might appear as striking given the relatively higher density of black population. However, an important consideration is that the percentage of blacks employed in public services is also higher in such regions. As we will document in section 4, the race of the recipient plays an important role in determining the magnitude of the racial disparity.

Table 5: Heterogeneity by geographical areas

	Regions				Counties	
	North-East	Mid-West	South	West	Rural	Urban
Black	-0.026** (0.010)	-0.049*** (0.011)	-0.031** (0.013)	-0.039** (0.017)	-0.058*** (0.011)	-0.030*** (0.007)
R^2	0.036	0.050	0.076	0.037	0.088	0.038
N	3,666	7,346	4,975	3,092	5,488	13,591

Robust standard errors in parenthesis clustered at the state/public service type level. Dependent variable is a binary variable indicating whether a response to the email was provided (linear probability model).

All regressions include controls of col IV of Table 3.

3.5 Additional Results: Other Outcomes

The outcome analyzed so far is whether an inquiry receives a reply. In this sub-section we investigate some additional outcomes. In particular, we explore whether there are differences in the quality of the reply, as measured by the number of replies sent by the receiver and the length of the email (number of words). We also use a measure of cordiality of the response: a binary variable concerning whether the respondent addresses the sender by name or with a salutation.⁹ Finally, we consider the intensive margin of replies, measuring how long it takes for the recipient to reply (number of hours).

Table A4 shows descriptive statistics related to these outcomes. Most respondents send just one reply, although a few also send some follow-up emails. The average length of emails is just above 170 words, and it takes on average 25 hours to receive a reply. For these three outcomes, a raw comparison suggests no difference between black and white senders. There appears to be a difference in the measure of cordiality, with 72% of responses to white emails being classified as cordial as opposed to 66% of responses to black emails. This is confirmed by the regression analysis in Table 6, whereby cordiality represents the only significant difference between black and whites.

designate the last category as rural.

⁹For salutations, we search the text for the following keywords "Hi", "Mr", "Dear", "Hello", "Good", "Thank".

Therefore, it appears that black emails are not only less likely to receive a response, but also that – conditional on receiving a response – it is less likely to have a cordial tone. This result seems consistent with evidence of prejudice rather than statistical discrimination. Even if – for instance – dealing with citizens of low socioeconomic background is more costly in terms of time or effort and recipients use race to infer the socioeconomic background of the sender, once a response is sent it seems unjustified to use a less cordial tone towards black senders. We explore the interpretation of our overall results in depth in the next section.

Table 6: Other outcomes

	Number	Length	Cordial	Delay
Black	-0.001 (0.003)	-0.030 (1.854)	-0.064*** (0.007)	-1.110 (1.404)
R^2	0.015	0.055	0.109	0.024
N	13,322	13,322	13,322	13,322

Robust standard errors in parenthesis clustered at the state/public service type level.

Dependent variables are, respectively: number of replies obtained, length of replies (n. of words), whether the sender was addressed by name or with salutations, and delay in obtaining a reply.

All regressions include controls of col IV of Table 3.

4 Interpretation: Taste-based vs Statistical Discrimination

Thus far, our results indicate a statistically and economically significant difference in the response to emails signed by white and black names. One possible interpretation is that this represents taste or prejudice-based discrimination, whereby responders may have an aversion to interacting with citizens with black-sounding names due to racially prejudicial attitudes or they may consider such citizens less worthy of their effort and attention. Another possibility is that the lower response to black emails represents a form of statistical discrimination, whereby the distinctively black names might signal some other personal trait, besides race, such as a certain socioeconomic background (Fryer and Levitt, 2004). In the labor market context, it has been argued that employers may use race to infer unobserved characteristics that are relevant for productivity. Thus, profit maximizing employers may statistically discriminate against some groups even if they are unprejudiced. In the context of public service provision by agencies not maximizing profit, it is conceivable to think about some objective function (e.g., effort minimization) that may give rise to something similar. For instance, dealing with citizens of low socioeconomic background could be more costly in terms of the time and effort of public service workers. In what follows, we explore which of these two


possible explanations prevails in our study by using two approaches: the first uses the inferred race of the respondent, while the second uses the socioeconomic background signaled by the sender.

4.1 Race of the recipient

In the first approach, we consider that if statistical discrimination were the primary driver of the difference, we would expect the recipient’s race not to be an important predictor of a response to a black email. Accordingly, white and black recipients should have a similar propensity to respond to names conveying low socioeconomic background, i.e., black names.

As a first attempt at assessing this view, consider Figure 3, which plots the gap in the response rate to black and white emails against the share of black population in the public sector, both at the state level. The relationship between these two variables – weighted by the number of emails sent in each state – appears to be negative.¹⁰

Figure 3: Difference in response rates and density of black population in employment



Black/white gap in response rate is obtained by pooling the data of the two waves and aggregating the data at the state level.

Observations are weighted by the number of emails sent in each state.

N=50 (Washington D.C. is excluded).

Since we do not have exact information about the race of the recipient, we try to proxy for

¹⁰Note that the graph is constructed using information from the emails of both waves. We exclude from the graph Washington D.C.. Only 44 emails were sent there, 23 in wave 1 and 21 in wave 2. At the same time, the response rate is 37% higher for blacks. Note that the black population is particularly over-represented in Washington D.C.. Our data show that the share of blacks in employment is 38% (vis-à-vis 8% at the national level) and that the share of blacks in public employment is 34% (while only 11% for the whole U.S.). As discussed in the next sections, the racial composition of the respondents play a key role in determining our results. The presence of the Washington D.C. “outlier” does not change the slope of line in the figure, which is weighted by number of emails in each state.

it by using two different methodologies. First, in Table 7, we introduce a proxy for the racial identity of the recipient, captured by the share of blacks among all employed individuals in the county (columns I and II) or the share of blacks among all employed in the public sector in the county (columns III and IV).¹¹ We interpret this as a proxy for the probability of the person who receives the email being black. We observe that the share of blacks in the county (column I) or the public sector (column III) is associated with a significant reduction in the probability of receiving a response.¹² When we interact the share of blacks with the black dummy variable (columns II and IV), we obtain a positive and statistically significant coefficient, indicating that the higher the probability of the recipient being black, the higher the probability of responding to a black email. In additional tests, we estimated models interacting the quartiles of the “share of blacks” variables with the black dummy variable, finding similar results.¹³ To facilitate an interpretation of these estimates, Figure 4 shows the predicted probability of response using the estimates in column II of Table 7, by deciles of the distribution of the share of blacks in employment. Predictions are calculated by varying the values of the share of blacks in employment (represented by the mid-point of each decile) and averaging over the remaining covariates. The figure shows that there is a statistically significant difference in the predicted probability of response across races where the likelihood of the recipient to be black is less than 10 percent (bottom eight deciles). In the top two deciles, where the probability that a recipient is black becomes more substantial, the predicted response rates for the two races become indistinguishable.

The evidence presented thus far suggests that black recipients are less likely to ignore black emails, which supports the interpretation that the estimated gap reflects actual taste-based racial discrimination.

In the second approach, we attempt to identify the race of the recipient more directly by inferring it from the surname associated with each email address. Given that each email address in our database is associated with the name of the recipient, we have this information for both respondents and non-respondents. For each surname, we compute two indices for the “probable race” of recipient, one for black names and one for white names, corresponding to the frequencies of surnames by race and ethnicity as reported in the 2000 Census. The idea is to proxy for the

¹¹The correlation between these two measures is 0.91. Note that in the case of the share of blacks in public employment, we have a smaller number of observations. This is because such information, obtained from the pooled 2006-2010 American Community Survey (ACS), is available only for some counties. The fact that the estimate in column III is lower than our benchmark is not surprising, given that the available counties are all classified as urban areas, where we know that the racial gap is relatively less pronounced.

¹²For example, moving from the 1st to the 3rd quartile of the share of blacks in the county (column I) implies a reduction in the response rate of nearly 2 percentage points (from 71.32% to 69.48%).

¹³In particular, in the case of the model with the share of blacks in employment, the estimate for the baseline category (first quartile) is -0.073 (s.e. 0.015) and the estimates for the interactions between the 2nd, 3rd and 4th quartiles and the black dummy are: 0.029 (s.e. 0.017), 0.045 (s.e. 0.020) and 0.063 (s.e. 0.020). For the model using the share of blacks among employed in the public sector, the estimate for the baseline category is -0.040 (s.e. 0.020), while the estimates for the interactions between the 2nd, 3rd and 4th quartiles and the black dummy are: 0.021 (s.e. 0.031), -0.004 (s.e. 0.025) and 0.069 (s.e. 0.027).

Table 7: Mechanisms: Race of recipient

	Share of blacks in employment		Share of blacks in public sector	
Race	-0.038*** (0.006)	-0.051*** (0.008)	-0.019** (0.009)	-0.047*** (0.013)
% Black in county	-0.243*** (0.072)	-0.344*** (0.078)	-0.121 (0.153)	-0.261 (0.167)
Black × % Black in county		0.208*** (0.070)		0.272*** (0.090)
R^2	0.050	0.051	0.028	0.029
N	19,079	19,079	7,406	7,406


Robust standard errors in parenthesis clustered at the state/public service type level.

Dependent variable is a binary variable indicating whether a response to the email was provided (linear probability model).

% of black refers to the county share of black population among the employed population (columns I and II) and among the employed population in the public sector (columns III and IV).

All regressions include controls as in col IV of Table 3.

Figure 4: Race of recipient



Predictions from col II of Table 7.

The x-axis represents deciles of the share of black population employed in the public sector. The estimates are calculated at the values of the 5%, 15%, ... , 95% percentiles. These values are: 0, 0.002, 0.004, 0.008, 0.013, 0.023, 0.042, 0.072, 0.123, 0.272.

probability that a certain name is white or black. We can then order surnames in our database according to the confidence by which we can associate them with a certain race, thereby obtaining a distribution for the “probable race” index. Subsequently, we set several thresholds corresponding to fixed percentiles of this distribution. For example, a threshold of 1% means that we select the top 1% values of the distribution. In the case of blacks, this threshold includes values of the probable race index which range from 48.38% (e.g., the surname Mack, with the census showing that nearly half of the people holding this surname are black) to 94.39% (e.g., the surname Ravenell, for which blacks represent the great majority). In the case of whites, the 1% threshold includes values ranging from 99% (e.g., the surname Kobylski) to 99.82% (e.g., the surname Sickie). Lower thresholds include surnames that are less frequent, e.g., for blacks Nicholson (with a value of 18.74%) and for whites Kline (with a value of 95.38%). We can subsequently select samples of recipients for which we are increasingly confident about their association with a specific race and estimate the corresponding racial gap in response. We present these estimates in Figure 5, which shows that: (i) samples where the name of the recipient is identified as being black are associated with a smaller race gap in response than those identified as being white, and (ii) the more accurately (e.g., a threshold of 5% or 1%) we can designate the race of the recipient as being black (white), the smaller (larger) the estimated adverse treatment experienced by blacks, although estimates become more imprecisely estimated. Again, these results are consistent with the taste-based discrimination interpretation of the differential response rate that we find.


4.2 Fixing the socioeconomic background

We next turn attention to the second wave of emails, which – as mentioned in the description of the experiment – includes a signature indicating the sender’s occupation (real estate agent). Note that, according to data from the Bureau of Labor Statistics for 2014, the annual mean wage of real estate agents (\$55,530) is above the annual mean wage for all occupations (\$47,230).¹⁴ Hence, this occupation should act as a signal to the recipient that the sender is a middle-class person or at least that he does not belong to a particularly low socioeconomic group. This feature of the design is meant to assess whether the racial difference found in wave 1 is primarily attributable to socioeconomic background.

The overall response rate to the second wave’s emails was slightly lower than the first wave (63.25%). In column I of Table 8, we report the racial difference in response rate estimated from our baseline linear probability regression using only data from the second wave. We find the racial gap to be 3.6 percentage points, namely almost identical to what we estimated in the first wave where there was no occupation signal. This provides evidence that the differential response to white versus black names is not specific to black names that are associated with low socioeconomic background, but is also present when we compare emails sent by individuals who belong to

¹⁴http://www.bls.gov/oes/current/oes_nat.htm.

Figure 5: Discrimination conditional on probable race of recipient



Points represent regression coefficients estimated on subsamples of data. Subsamples are defined first by constructing an index of the probability of belonging to the white or black race (matching surname from our data to the 2000 Census). Thresholds indicate the upper portion of the frequency distribution of the index and are used to define the various subsamples.

a middle income occupational group. In column II, we pool observations of the two waves, finding remarkably similar results. In column III, we test whether there is a difference in the treatment between black and white emails across the two waves, finding that the interaction term between the black dummy and the dummy for the second wave is very small and not statistically significant. In columns IV to VI, we estimate the same regressions using cordial reply as the outcome variable. The estimates are remarkably similar to those obtained in Table 6, indicating that the differential in the likelihood of receiving a reply with a cordial tone is not attributable to black names that might evoke low socioeconomic background. Together with the results presented in the previous sub-section, this evidence points to prejudice-based discrimination being behind our finding.

5 Conclusions

We carry out an email correspondence study that aims to identify whether racial discrimination exists in the provision of information regarding public services offered by local offices in the U.S. (school districts, libraries, sheriff offices, treasurers, job centers and county clerks). Overall, we find that requests of information coming from a person with a distinctively black name are less likely

Table 8: Fixed socioeconomic background

	Reply			Cordial reply		
	Wave II	Pooled waves		Wave II	Pooled waves	
Black	-0.037*** (0.007)	-0.037*** (0.005)	-0.041*** (0.013)	-0.064*** (0.009)	-0.064*** (0.006)	-0.066*** (0.015)
Black \times Wave II			0.002 (0.008)			0.002 (0.010)
R^2	0.052	0.052	0.052	0.099	0.102	0.102
N	19,089	38,168	38,168	12,073	25,395	25,395

Robust standard errors in parenthesis clustered at the state/public service type level.

Dependent variable in columns I-III is a binary variable indicating whether a response to the email was provided (linear probability model) and in columns IV-VI is a binary variable indicating whether the sender was addressed by name or with salutations.

Wave II refers to the follow-up wave where occupation is signaled in the email.

Pooled refers to the pooling of Wave I and Wave II.

All regressions include controls as in col IV of Table 3.

to receive a reply than those from a person with a distinctively white name. This holds even if we signal that the sender is not of low socioeconomic status, which indicates that the differential treatment of black emails is unlikely to be attributable to statistical discrimination.

Besides being illegal, discrimination by public service providers is particularly startling, since governments could be major players in the effort to eradicate discrimination in American society. For instance, school districts and libraries can play an important role in closing the educational achievement gap of black children. Indeed, our interest in the local level government relates to the fact that street-level bureaucrats are responsible for the implementation of policy enacted at both the federal and state level.

One common criticism of correspondence studies in the labor market is that these analyses may not measure labor market discrimination that blacks experience in equilibrium. The explanation is that blacks may respond to the presence of discrimination by sorting themselves across firms (e.g., minimizing their contact with the most discriminatory ones) or adopting different job-search strategies than whites (e.g., sending more resumés, see Charles and Guryan, 2011). This is less of an issue in the case of local public services, since providers are local monopolies in many cases. Thus, residents of a given locality cannot usually choose with which school district or sheriff office to interact. Of course it is true that black citizens may respond to the differential treatment that we have uncovered by becoming more vocal in asking public officials to fulfil their duties (for instance, by sending more “reminders” to unresponsive offices). However, this entails a cost, both psychologically and in terms of time. Moreover, besides “voice”, there is the alternative option of “exit” (Hirschman, 1970), whereby black citizens who feel discriminated by public offices may reduce their interaction with them as much as possible, with potentially high costs in terms of the foregone consumption of public services. In our settings, we cannot investigate which type of reaction prevails. However, the fact that – for instance – African Americans account for 12.5% of

the citizen voting age population, but only occupy 5.7% of city council seats and 3% of local offices (Joint Center for Political and Economic Studies, 2015) suggests a certain disengagement with local offices.

Overcoming discriminatory practices in local public services is a complex issue. The persistence of such practices despite their illegality suggests that they will not be eradicated through a quick legislative fix. Possible interventions include hiring policies aimed at increasing diversity among the workforce or promoting racial matching between employees and the communities they serve (Lang, 2015). What this paper shows is that discriminatory practices are present in terms of access to public services and that policy makers should consider such interventions.

References

- Abrams, D., M. Bertrand, and S. Mullainathan (2012). Do Judges Vary in Their Treatment of Race? *Journal of Legal Studies* 41(2), 347–383.
- Alesina, A. and E. La Ferrara (2014). A Test of Racial Bias in Capital Sentencing. *American Economic Review* 104(11), 3397–3433.
- Altonji, J. and R. Blank (1999). Race and Gender in the Labor Market. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics* 3C, Chapter 48, pp. 3143–3259. New York: Elsevier.
- Bayer, P., F. Ferreira, and S. L. Ross (2014). Race, Ethnicity and High-Cost Mortgage Lending. *NBER Working Paper 20762*.
- Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press, Chicago.
- Bertrand, M. and S. 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.
- Bhargava, S. and D. Manoli (2015). Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment. *American Economic Review* (forthcoming).
- Broockman, D. E. (2013). Black Politicians Are More Intrinsically Motivated to Advance Blacks’ Interests: A Field Experiment Manipulating Political Incentives. *American Journal of Political Science* 57(3), 521–536.
- Butler, D. M. and D. E. Broockman (2011). Do Politicians Racially Discriminate against Constituents? A Field Experiment on State Legislators. *American Journal of Political Science* 55(3), 463–477.
- CDC, C. (2011). Health Disparities and Inequalities Report. *Supplement Vol 60*, <http://www.cdc.gov/mmwr/pdf/other/su6001.pdf> (last accessed on July 27, 2015).
- Charles, K. K. and J. Guryan (2011). Studying Discrimination: Fundamental Challenges and Recent Progress. *Annual Review of Economics* 3(1), 479–511.
- Coviello, D. and N. Persico (2013). An Economic Analysis of Black-White Disparities in NYPD’s Stop and Frisk Program. *NBER Working Paper 11195*.
- Daponte, B. O., S. Sanders, and L. Taylor (1999). Why do Low-Income Households not use Food Stamps? Evidence from an Experiment. *Journal of Human Resources* 34(3), 612–628.

- Distelhorst, G. and Y. Hou (2014). Ingroup Bias in Official Behavior: A National Field Experiment in China. *Quarterly Journal of Political Science* 9(2), 203–230.
- Doleac, J. L. and L. C. Stein (2013). The Visible Hand: Race and Online Market Outcomes. *Economic Journal* 123(572), F469–F492.
- Duflo, E. and E. Saez (2003). The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment. *Quarterly Journal of Economics* 118(3), 815–842.
- Ewens, M., B. Tomlin, and L. C. Wang (2014). Statistical Discrimination or Prejudice? A Large Sample Field Experiment. *Review of Economics and Statistics* 96(1), 119–134.
- Fryer, R. G. (2011). Racial Inequality in the 21st Century: The Declining Significance of Discrimination. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics 4B*, Chapter 10, pp. 855–971. New York: Elsevier.
- Fryer, R. G. and S. D. Levitt (2004). The Causes and Consequences of Distinctively Black Names. *Quarterly Journal of Economics* 119(3), 767–805.
- Fryer, R. G. and P. Torelli (2010). An Empirical Analysis of Acting White. *Journal of Public Economics* 94(5), 380–396.
- Glaeser, E. L. and B. Sacerdote (2003). Sentencing in Homicide Cases and the Role of Vengeance. *Journal of Legal Studies* 32(2), 363–382.
- Gneezy, U., J. List, and M. K. Price (2012). Toward an Understanding of why People Discriminate: Evidence from a Series of Natural Field Experiments. *NBER Working Paper 17855*.
- Guryan, J. and K. K. Charles (2013). Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots. *Economic Journal* 123(572), F417–F432.
- Harrell, J. P., S. Hall, and J. Taliaferro (2003). Physiological Responses to Racism and Discrimination: an Assessment of the Evidence. *American Journal of Public Health* 93(2), 243–248.
- Hastings, J. S. and J. M. Weinstein (2008). Information, School Choice, and Academic Achievement: Evidence from two Experiments. *Quarterly Journal of Economics* 123(4), 1373–1414.
- Hirschman, A. O. (1970). *Exit, Voice and Loyalty*. Harvard University Press, Cambridge/Massachusetts.
- Hoxby, C., S. Turner, et al. (2013). Expanding College Opportunities for High-Achieving, Low Income Students. *Stanford Institute for Economic Policy Research Discussion Paper 12-014*.

- Joint Center for Political and Economic Studies (2015). 50 Years of the Voting Rights Act. <http://jointcenter.org/sites/default/files/> (last accessed on July 27, 2015).
- Lang, K. (2015). Racial Realism: A Review Essay on John Skrentny's After Civil Rights. *Journal of Economic Literature* 53(2), 351–359.
- Levine, R., Y. Rubinstein, and A. Levkov (2014). Bank Deregulation and Racial Inequality in America. *Critical Finance Review* 3(1), 1–48.
- Lipsky, M. (1980). *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services*. New York, Russel Sage Foundation.
- Pew Research Center (2013). *How Americans Value Public Libraries in Their Community*.
- Riach, P. A. and J. Rich (2002). Field Experiments of Discrimination in the Market Place. *Economic Journal* 112(483), F480–F518.
- Rich, J. (2014). What do Field Experiments of Discrimination in Markets Tell us? A Meta Analysis of Studies Conducted since 2000. *IZA Discussion Paper 8584*.
- Sabol, W. J., H. C. West, and M. Cooper (2009). Prisoners in 2008. *Bureau of Justice Statistics Bulletin* (NCJ 228417), 1–45.
- Stephens-Davidowitz, S. (2014). The Cost of Racial Animus on a Black Candidate: Evidence Using Google Search Data. *Journal of Public Economics* 118, 26–40.
- US Government (2013). United States Periodic Report to the Committee on Elimination of Racial Discrimination. <http://www.state.gov/documents/organization/210817.pdf> (last accessed on July 27, 2015).
- White, A. R., N. L. Nathan, and J. K. Faller (2015). What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials. *American Political Science Review* 109(1), 1–14.
- Wong, G., A. O. Derthick, E. David, A. Saw, and S. Okazaki (2014). The What, the Why, and the How: A Review of Racial Microaggressions Research in Psychology. *Race and Social Problems* 6(2), 181–200.

Appendix

Table A1: Details of emails

Recipient	N. recipient	Sent emails	Undelivered/testing	Final sample size	Emails in sample / N. recipients
Wave I					
School D.	13,567	10,882	1,009	9,873	73%
Library	14,638	5,350	456	4,894	33%
Sheriff	3,080	2,087	251	1,836	60%
Treasurer	3,143	1,252	123	1,129	36%
Job Center V.R.	3,146	890	159	731	23%
County Clerk	3,143	691	75	616	20%
Total	40,717	21,152	2,073	19,079	47%
Wave II					
School D.	13,567	10,882	1,029	9,853	73%
Library	14,638	5,350	420	4,930	34%
Sheriff	3,080	2,087	247	1,840	60%
Treasurer	3,143	1,252	118	1,134	36%
Job Center V.R.	3,146	890	169	721	23%
County Clerk	3,143	691	80	611	19%
Total	40,717	21,152	2,063	19,089	47%

Source: own computations.

Table A2: Response rates - by sending name

	DeShawn Jackson	Tyrone Washington	Total Black	Greg Walsh	Jake Mueller	Total White
Response rate	69.05 (46.23)	66.91 (47.06)	67.96 (46.67)	71.57 (45.11)	71.76 (45.02)	71.67 (45.06)
N	4,637	4,835	9,472	4,918	4,689	9,607
Difference within race (abs)	2.15		0.19			
z-stat (p-val)	0.025		0.836			
Difference B-W	-2.62	-4.76				
z-stat (p-val)	0.001	0.000				

Source: own computations. Entries are response rates multiplied by 100, standard deviations in parentheses.

Table A3: Response rates

	School D.	Library	Sheriff	Treasurer	Job Center	County Clerk	Total
White	76.53 (42.39)	69.08 (46.23)	53.23 (49.92)	73.90 (43.96)	71.93 (44.99)	65.72 (47.54)	71.67 (45.06)
Black	73.10 (44.35)	64.96 (47.72)	46.26 (49.89)	69.57 (46.05)	73.08 (44.42)	64.09 (48.05)	67.96 (46.67)
Difference B-W	-3.43	-4.12	-6.98	-4.33	1.14	-1.63	-3.71
z-stat (p-val)	0.000	0.002	0.003	0.107	0.730	0.672	0.000

Source: own computations. Entries are response rates multiplied by 100, standard deviations in parentheses.

Table A4: Other outcomes, summary statistics

	Black	White	t-test (p-val)
Number of replies	1.03 (0.18)	1.03 (0.19)	0.78
Length of reply (# words)	171.4 (99.2)	170.6 (101.3)	0.66
Cordial reply	66.3 (47.3)	72.2 (44.8)	0
Delay in reply (hours)	24.5 (73.9)	25.7 (97.2)	0.43

Source: own computations. Standard deviations in parentheses.

Table B1: Data sources of email addresses

Recipient	Source of email addresses	Accessed/obtained
School Districts	http://schoolinformation.com/	November 3, 2014
Libraries	http://www.americanlibrarydirectory.com	October 7, 2014
Sheriffs	http://www.sheriffs.org	October 7, 2014
Treasurers	http://www.uscounties.org	October 8, 2014
Job Centers	http://www.servicelocator.org	November 18, 2014
County Clerks	http://www.uscounties.org	October 8, 2014

Table B2: Email Queries by Recipient

Recipient	Simple Query	Complex Query
School District	I would like to enroll my son in a school in this district and I have some questions. Could you please tell me what your opening hours are?	I would like to enroll my son in a school in this district. Could you please let me know the documents I would need to do this? Do I also need an immunization record?
Library	I would like to become a member of the library. Could you please tell me what your opening hours are?	I would like to become a member of the library. Could you please explain what I need to do for this? Do I need proof of address?
Sheriff	I am performing a background check on a local individual. Could you please tell me what your opening hours are?	I am performing a background check on a local individual. Could you please tell me what the procedure is for a criminal record search and how much it would cost?
Treasurer	I am about to purchase a house and I have some questions about property taxes. Could you please tell me what your opening hours are?	I am about to purchase a house. Could you please explain how I can check whether there are unpaid taxes on the house? If there are unpaid taxes, who would be liable for them?
Job Center	I am recently unemployed and have some questions about benefits. Could you please tell me what your opening hours are?	I am recently unemployed. Could you tell me what conditions I need to meet to be eligible for benefits and how would I apply to receive them?
County Clerk	My partner and I would like a marriage license. Could you please tell me what your opening hours are?	My partner and I would like a marriage license. Could you please let me know the procedure for applying for one? Also would such a license only be valid in this county, or would it be recognized elsewhere?

Table B3: Data sources of ancillary variables

Variable	Description	Source	Year	Weblink
% of black	Total black population over total population in each county	Census: Profile of general population	2010	http://factfinder.census.gov/
% of black among employed	Total black population over total population employed in each county	American Community Survey	2006-2010	http://www.census.gov/people/eoetabulation/data/eoetables20062010.html
% of black in public sector	Total black population employed in public sector over total population employed in public sector in each county	American Community Survey	2006-2010	http://www.census.gov/people/eoetabulation/data/eoetables20062010.html
% of hispanic population	Total hispanic population over total population in each county	Census: Profile of general population	2010	http://factfinder.census.gov/
Unemployment rate	Total population who is unemployed over total population in labor force in each county	Local Area Unemployment Statistics	2014	http://www.bls.gov/lau/#cntyaa
Average labor income	Average weekly wage based on the 12-monthly employment levels in each county	Quarterly Census of Employment and Wages	2014	http://www.bls.gov/cew/datatoc.htm
Crime rate	Total number of violent and propriety crime over total population in county	Uniform Crime Reporting Program Data	2012	http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/35019?searchSource=revise&q=Uniform+Crime+Reporting+Program+Data+county+2012&paging.rows=100&sortBy=0
% of Dem votes	Shares of votes given to Barack Obama over total votes in each county	The guardian	2012	http://www.theguardian.com/news/datablog/2012/nov/07/us-2012-election-county-results-download http://www.elections.alaska.gov/results/12GENR/data/results.htm
Urbanization	Rural corresponds to the category "Non-core", Urban to the categories "Large central metro, Large fringe metro, Medium metro, Small metro and Metropolitan"	National Center for Health Statistics	2013	http://www.cdc.gov/nchs/data_access/urban_rural.htm