

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

Corrado Giulietti[†] Mirco Tonin[‡] Michael Vlassopoulos[§]

(forthcoming in the Journal of the European Economic Association)

Accepted version: August 2017

Abstract

We examine whether racial discrimination exists in access to public services in the U.S.. 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 from putatively black senders are almost 4 percentage points less likely to receive an answer compared to emails signed with a white-sounding name. Moreover, responses to queries coming from black names are less likely to have a cordial tone. Further tests suggest that the differential in the likelihood of answering is due to animus towards blacks rather than inferring socioeconomic status from race. Finally, we show that attitudes towards the government among blacks are more negative in states with higher discrimination.

*We would like to thank the Editor Daniele Paserman and three anonymous referees for very constructive comments. We also thank Matthew Blackwell, Raj Chetty, Rajeev Dehejia, Stefano DellaVigna, Uri Gneezy, Jonathan Guryan, Nathaniel Hilger, Guillermina Jasso, Gary King, Peter Kuhn, Uwe Sunde and Adam Szeidl for useful discussions. We also thank participants of seminars at IZA, Central European University, Temple, Köln, Newcastle, Surrey, Innsbruck, Verona, Rotterdam, and participants of the 2015 Society for Government Economists Conference, the Workshop on the Economics of Discrimination in Naples, the 2016 Royal Economic Society Conference, the 39th NBER Summer Institute, the 33rd European Association of Law and Economics for their comments. We thank Kerwin Kofi Charles and Jonathan Guryan for kindly providing us with their data. This project has received ethical approval from the Institute for the Study of Labor (IZA) in Bonn, Germany, and from the University of Southampton.

[†]University of Southampton. E-mail: C.Giulietti@soton.ac.uk

[‡]Free University of Bozen-Bolzano; CESifo; Dondena Centre for Research on Social Dynamics and Public Policy, Bocconi University; IZA. E-mail: Mirco.Tonin@unibz.it

[§]University of Southampton; IZA. E-mail: M.Vlassopoulos@soton.ac.uk

1 Introduction

African Americans have a disadvantaged position in American society in terms of economic 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 (U.S. 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 U.S. 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. Targeted services include school districts, local libraries, sheriff offices, county clerks, county treasurers and job center veteran representatives. In particular, we collect all available emails of the targeted local public service providers, which gives us more than 19,000 cases, corresponding to roughly half of

¹ Altonji 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.

² Charles 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).

the total number of providers. 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).⁴

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. Correspondence and audit studies have 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., 2017) and along different dimensions, including race, ethnicity, gender, age, disability, sexual orientation, obesity, caste and religion.⁵

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 led to substantial additional claiming. Hoxby and Turner (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

³ A 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., searching for job opportunities, submitting online job applications or working 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”.

⁴ County clerks are generally responsible for keeping records of deeds and marriage licenses and most other public records. They also issue many licenses and often have responsibilities for elections, including the preparation of ballots and the recruitment and training of poll workers. County treasurer offices' responsibilities include, for instance, the issuance of vehicle titles and registrations, the collection of vehicle fees, and the collection of property taxes for local governments. The responsibilities of school districts may include the selection of textbooks and other curriculum materials, the hiring and dismissal of staff, the monitoring of finances, ensuring compliance with relevant laws and the maintenance of school buildings. Sheriff offices have a range of duties that include criminal investigations, traffic enforcement, and operation of the jail. Their responsibilities may include serving warrants and evictions.

⁵ See Riach and Rich (2002), Guryan and Charles (2013), Rich (2014), Neumark (2016) and Bertrand and Duflo (2017), for earlier and more recent reviews of the literature.

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 non-participation. 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 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 may also 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 African-American 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 African-American names are less likely to have a cordial tone. Thus, despite a rising sentiment among whites that the so-called reverse discrimination is on the rise in the U.S. (Norton and Sommers, 2011) due, for instance, to affirmative actions, what we find is evidence of discrimination against African Americans by public service providers. Note that email is an important channel of communication between the public and government agencies in the U.S.. A recent survey (Smith, 2010) finds that contacting a government agency via email or a website is the second most preferred channel (28%) - telephone being the most preferred (35%) - and is a more preferred option than in-person visits (20%). Among internet users, online contact is the preferred option.

In the context of our study, the case for statistical discrimination is much weaker than in a labor market setting. Nevertheless, we also investigate whether the differential response rate is due to animus or to a form of statistical discrimination arising from assigning low socioeconomic status to a sender with a distinctively African-American 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 – in a second wave of emails that we sent – by signaling his occupation. A change in the racial difference in responses across waves due to this additional information would indicate the importance of statistical discrimination. Both approaches are suggestive of taste-based discrimination being an important driver of the race gap in the response rate.

Our results are consistent with other studies that uncover evidence of racial/ethnic 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 a response and, if granted, to receive responses of lower quality than those sent from non-Latino white aliases. A 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.

To the best of our knowledge, our study is the first to explore racial 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. In Section 5, we explore the association between our measure of discrimination and the racial gap in wages and attitudes towards government.

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: school districts, local libraries, sheriff offices, county treasurers, job center veteran representatives and county clerks. 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 A.1 in the Appendix).^{6,7} 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 B.1 in the Appendix reports the detailed number of existing recipients and of emails in the sample).

Table 1: Emails by Type of Service

Recipient	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	731	3.83
County Clerk	616	3.23
Total	19,079	100

Notes: Figures refer to the number of emails sent to each type of service.

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 large number of emails.

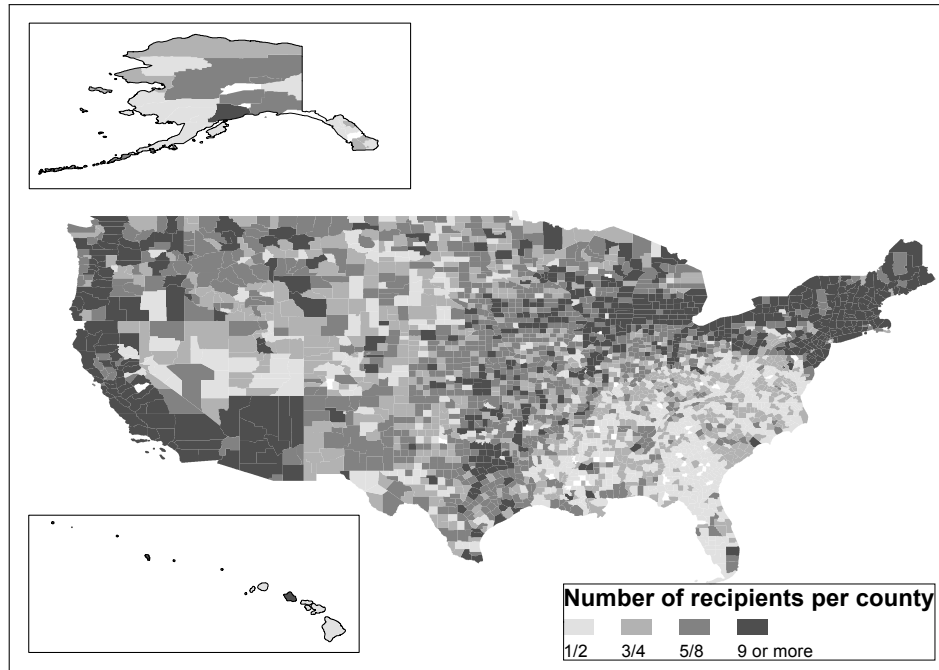
Figure 2 plots the share of emails against the share of recipients across states. As can be seen, most observations are clustered around the 45-degree line, suggesting that the number of email addresses is proportional to the number of potential recipients across states. We will account for

⁶ A small random sample was used for testing; about 2,000 emails were eliminated because they were undelivered. We checked and corroborated that the probability that an email is eliminated does not correlate with our key variables.

⁷ Most of our email addresses are of the “personal” type rather than “generic”, that is, they include the name or surname of the receiver and are not, for instance, of the type “info@...” or “office@...”. We do not know, however, who the actual reader is. It could well be the case that an email addressed to the sheriff is actually read by a deputy or a personal assistant.

any discrepancies in one of our robustness checks.

Figure 1: Location of Recipients



Notes: white fill indicates counties with no recipients.

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. For example, in the case of school districts, the simple email asks the following question: "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?". The complex email asks: "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?". The idea of having both types of emails is to check whether the degree of discrimination changes with the complexity (i.e., amount of effort) required by the task. Every email has the following format:


```

*****
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]
*****

```

Figure 2: Sample Representativeness



Share of emails corresponds to the number of emails in a state over the total number of emails in the U.S.; share of potential recipients corresponds to the number of potential recipients in the state over all potential recipients in the entire U.S..

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 A.2 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 African-American 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 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. The field “City Name” contains the name of the city where the public service provider is located.

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 most of the analysis, we pool the two black and two white names. We randomized the treatments at the state/public service type. This means, for example, that we randomized across school districts in California, and then across sheriff offices in California, etc. We also randomized the order in which emails were sent across the treatments, as well as across types of public services and states.

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 A.3 in the Appendix). As evident from Figure 1, the vast majority of counties (98.37%) are present in our sample with at least one recipient. Nevertheless, this is not the case for each type of recipient. Hence, in Table B.2 in the Appendix, we present summary statistics of the local characteristics of both counties that are covered in our sample and those that are not, by type of recipient.

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

Table 2: County Characteristics of Email Recipients

	Black	White	t-test (pval)
% of blacks in employment	6.10 (10.01)	6.07 (10.25)	0.83
Unemployment rate (%)	6.13 (2.14)	6.12 (2.11)	0.82
% of Hispanic	10.05 (13.91)	9.83 (13.81)	0.27
Average labor income (USD)	794 (239)	789 (237)	0.17
Crime rate (%)	2.45 (1.31)	2.44 (1.29)	0.47
% of Dem votes	43.05 (15.13)	42.96 (14.99)	0.67
Urbanization	71.80 (45)	70.68 (45.53)	0.09
N	9,472	9,607	

Notes: Standard deviations in parentheses. P-value refers to the t-test for the difference between two means.

– about half of the sample – 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. Furthermore, we also investigate measures of the quality of response by focusing on the number of responses received, the length of the response (number of words), the delay in the response and how cordial the response is.

3 Results

3.1 Descriptive Statistics

Overall, about 70% of the 19,079 emails that we sent received a response (see Table B.3 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 68.88% for simple emails and 70.84% 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 71.66% 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 (71.74% for Jake Mueller and 71.57% for Greg Walsh; z-stat, p-value 0.85), 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 African-American 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.⁸ Indeed, the 2000 Census shows that among the 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. Hence, hereafter we will consider the difference between white and black emails without distinguishing between the two names within each category.

To summarize, the descriptive evidence 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. This variable is coded as 1 if the response was received within 40 days and 0 otherwise. The variable *ServiceType* indicates the type of public service to which the email was sent, *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. Estimating regressions with state/public service fixed effects produces numerically similar results across all our specifications. 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 3.84 percentage points – does not substantially differ from that emerging from the raw comparison reported in the previous

⁸ In their analysis of the distribution of first names in California, Fryer and Levitt (2004) find that the first name Tyrone is slightly more represented among whites relative to DeShawn. In particular, they report that “[t]here are 463 children named DeShawn, 458 of whom are Black. The name Tyrone is given to 502 Black boys and only 17 Whites.” (Fryer and Levitt, 2004, page 770).

Subsection (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 (1.84 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 the black name and the complex email dummies. 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 hence estimate a model with office (recipient) fixed effects. The within-recipient variation in the responsiveness to white and black emails is similar to that estimated in column I (3.22 percentage points).

Table 3: Difference in Response Rates

	I	II	III	IV	V
Black	-0.038*** (0.006)	-0.038*** (0.006)	-0.037*** (0.009)	-0.038*** (0.006)	-0.032*** (0.005)
Complex		0.018*** (0.007)	0.019** (0.009)	0.018*** (0.007)	0.029*** (0.004)
Black \times Complex			-0.002 (0.013)		
\bar{Y}	0.698	0.698	0.698	0.698	0.665
R^2	0.045	0.045	0.045	0.049	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
Office 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).

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

Office 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 .

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 re-estimated 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. Next, 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 estimates. In particular, if we were to cluster at the state/public service type/sender name level, the standard error of the black dummy coefficient in the baseline model would be 0.008, whereas if we were to cluster at the county/public service type level, the standard error would be 0.007.⁹

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 might be interested in knowing whether our results might be driven by one particular type of public service. Hence, we analyze the results by type of service. When considering the response rates, the pattern of higher response rates for white emails holds in all cases, except for job centers. The raw differences in response rates are presented in Table B.4 in the Appendix. In Table 4, we estimate the econometric model in Equation (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. This is in line with the literature mentioned in the introduction finding evidence of discrimination in law enforcement. The finding of a lack of discrimination for more administrative-type jobs, like county clerks or job centers, goes against what found for U.S. election officials (White et al., 2015), but the roles are of course very different.¹⁰

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

⁹ Two further checks that we perform concern the functional form and the email structure. First, we estimated the baseline specification using a probit model. The marginal effect (for the baseline -0.040, s.e. 0.006) is remarkably similar to the estimate of the linear probability model, reassuring us that results are not sensitive to the chosen functional form. Second, we excluded from the sample emails with a generic structure (i.e., “office@...”, “admin@...”). When performing a regression in the resulting sample (15,851 observations), we find a coefficient estimate of -0.033 (s.e. 0.007).

¹⁰ In 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, showing 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)	-0.040 (0.029)	0.008 (0.031)	-0.016 (0.048)
\bar{Y}	0.748	0.670	0.498	0.718	0.725	0.649
R^2	0.040	0.047	0.108	0.082	0.164	0.092
N	9,873	4,894	1,836	1,129	731	616

Robust standard errors in parentheses clustered at the state level.

Dependent variable is a binary variable indicating whether a response to the email was provided (linear probability model). \bar{Y} refers to the average response rate.

All regressions include controls of col IV, Table 3.

into the four regions defined by the Census Bureau (North-East, Mid-West, South and West) and estimate our baseline specification on each subsample.¹¹ Table 5 summarizes the results. We find a significant racial gap in all four regions, with the estimate ranging from -0.026 in the North-East to -0.049 in the Mid-West.¹²

Table 5: Heterogeneity by Geographical Areas

	Regions				Counties	
	North-East	Mid-West	South	West	Rural	Urban
Black	-0.026** (0.010)	-0.049*** (0.010)	-0.031** (0.013)	-0.039** (0.017)	-0.058*** (0.011)	-0.030*** (0.007)
\bar{Y}	0.728	0.722	0.635	0.708	0.660	0.714
R^2	0.036	0.049	0.075	0.036	0.087	0.038
N	3,666	7,346	4,975	3,092	5,488	13,591

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). \bar{Y} refers to the average response rate.

All regressions include controls of col IV, Table 3.

To further document this pattern of geographical variation, in columns V and VI, we classify counties into urban and rural and split the sample along this dimension.¹³ This gives rise to 1,312 rural counties and 1,780 urban counties. The results indicate that the racial gap in response rate is

¹¹ The 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, 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.

¹² When we split the sample according to the 9 Census Divisions (see Figure C.1), all point estimates for the black dummy are negative albeit some of them are not statistically significant.

¹³ We apply the six-level classification developed by the National Center for Health Statistics: Large central metro, Large fringe metro, Medium metro, Small metro, Micropolitan and Noncore. We designate the last category as rural.

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 African Americans 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.

3.5 Additional Results: Other Outcomes

The outcome analyzed so far is whether an inquiry receives a reply. In this Subsection we investigate 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.¹⁴ Finally, we consider the intensive margin of replies, measuring the number of hours it takes for the recipient to reply.

Table 6: Other outcomes

	Number	Length	Cordial	Delay
Black	-0.001 (0.003)	-0.114 (1.847)	-0.062*** (0.007)	-0.918 (1.147)
\bar{Y}	1.029	171.041	0.708	24.684
R^2	0.015	0.055	0.114	0.031
N	13,321	13,321	13,321	13,321

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

Dependent variables are, respectively: number of replies obtained, length of replies (number of words), whether the sender was addressed by name or with salutations, and delay in obtaining a reply (number of hours). \bar{Y} refers to the average of the outcome variable reported in each column header. All regressions include controls of col IV, Table 3.

Table B.5 in the Appendix shows descriptive statistics related to these outcomes. Most respondents sent just one reply, although a few also send some follow-up emails. The average length of emails is just above 171 words, and it takes on average just over a day (25 hours) to receive a reply. For these three outcomes, a raw comparison suggests no difference between black and white

¹⁴ For salutations, we search the text for the following keywords “Hi”, “Mr”, “Dear”, “Hello”, “Good”, “Thank”.

senders. There appears to be a difference in the measure of cordiality, with 73.52% of responses to white emails being classified as cordial as opposed to 67.9% 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. Furthermore, Table C.1 in the Appendix shows that this difference is statistically significant in 5 out of the 6 types of public services (all but the job centers). 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.¹⁵

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 African-American senders. We explore the interpretation of our overall results in depth in the next Section.

4 Interpretation

Thus far, our results indicate a statistically and economically significant difference in the response to emails signed by white and African-American 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 African-American 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 our setting, characterized by a one-shot interaction between the sender and a public officer, the case for statistical discrimination is weaker and less likely to be of first order importance. Nevertheless, it is conceivable to think that public officers may attempt to minimize effort and perceive senders with a lower socio-economic status (SES) as requiring more help. In absence of direct information about SES, public officers may use race to infer it and, thus, be less likely to respond to queries coming from black-sounding names. Alternatively, officers may try to minimize the likelihood of being reprimanded, may consider senders with a low SES as less assertive and, thus, infer that senders with a black-sounding name are less likely to complain if their query is left unanswered. In what follows, we explore the relevance of taste-based vis-à-vis statistical discrimination adopting two approaches: the first uses the inferred race of the respondent and the

¹⁵ For robustness, we also constructed an alternative definition of cordial reply, defined as $0.5 \times (\# \text{ cordial words}) + 0.5 \times (\# \text{ characters of all cordial words})$, that takes more into account the “intensive” margin of cordiality. This variable produces a coefficient estimate of -0.110 (s.e. 0.030).

second uses the socioeconomic background signaled by the occupation of the sender. Moreover, we also explore the correlation of a discrimination index constructed using our data to measures of racial prejudice used in the literature.

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., African-American names.

As a first attempt at assessing this view, consider Figure ??, which plots a discrimination index from our data against the share of black population in employment, both at the state level. The discrimination index is obtained by aggregating responses at the state level and then calculating the racial gap in response rates (with larger values indicating higher discrimination). To increase the precision of our measure, we pooled data from both the first and second wave. The relationship between these two variables – weighted by the number of emails sent in each state – appears to be negative.¹⁶ Since we do not have exact information about the race of the recipient, we try to proxy for it by using two methodologies. In the first approach, for which we report the results in Table 7, we proxy for the probability of the person who receives the email being black using four measures: the share of blacks among employed individuals in the county (columns I and II); the share of blacks among employed individuals in the public sector in the county (columns III and IV); the share of blacks working as education administrators, sheriffs and librarians in the county (columns V and VI); the share of blacks among school principals in the state (columns VII and VIII).¹⁷

We observe that the share of blacks in the local area (columns I, III, V and VII) is associated with a significant reduction in the probability of receiving a response.¹⁸ When we interact the share of blacks with the black dummy variable (columns II, IV, VI and VIII), we obtain a positive and statistically significant coefficient, indicating that the higher the probability of the recipient being African American, the higher the probability that a black email receives a response. This result is consistent across the various measures considered. In additional tests, we estimated models interacting the quartiles of the share of blacks with the black dummy variable, finding similar

¹⁶ To improve readability, we excluded from the graph Washington D.C., where only 44 emails were sent (23 in wave I and 21 in wave II). In Washington D.C. the share of blacks in employment is 38% (vis-à-vis 8% at the national level) and the share of blacks in public employment is 34% (while only 11% for the whole U.S.). In any case, including Washington D.C. does not change the slope of the regression line in the Figure.

¹⁷ The correlation between the first two measures is 0.91. Note that, with the exception of the share of blacks in employment, data are available only for a subset of counties and states. The fact that the estimate in column III is lower than our baseline 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.

¹⁸ For example, moving from the 1st to the 3rd quartile of the share of blacks in employment (column I) implies a reduction in the response rate of nearly 2 percentage points (from 71.19% to 69.55%).

Table 7: Mechanisms: Race of Recipient

	% of blacks:							
	in employment		in public sector		among educators, librarians and sheriffs		among school principals	
Black	-0.038*** (0.006)	-0.051*** (0.008)	-0.019** (0.009)	-0.048*** (0.013)	-0.023** (0.010)	-0.051*** (0.013)	-0.032*** (0.009)	-0.073*** (0.015)
% of blacks	-0.243*** (0.072)	-0.344*** (0.078)	-0.181 (0.115)	-0.323*** (0.124)	-0.165 (0.132)	-0.316** (0.140)	-0.090 (0.158)	-0.265* (0.132)
Black × % of blacks		0.208*** (0.070)		0.272*** (0.090)		0.288*** (0.101)		0.366*** (0.122)
\bar{Y}	0.698	0.698	0.725	0.725	0.724	0.724	0.744	0.744
R^2	0.050	0.051	0.028	0.029	0.028	0.029	0.009	0.010
N	19,079	19,079	7,406	7,406	6,047	6,047	6,793	6,793

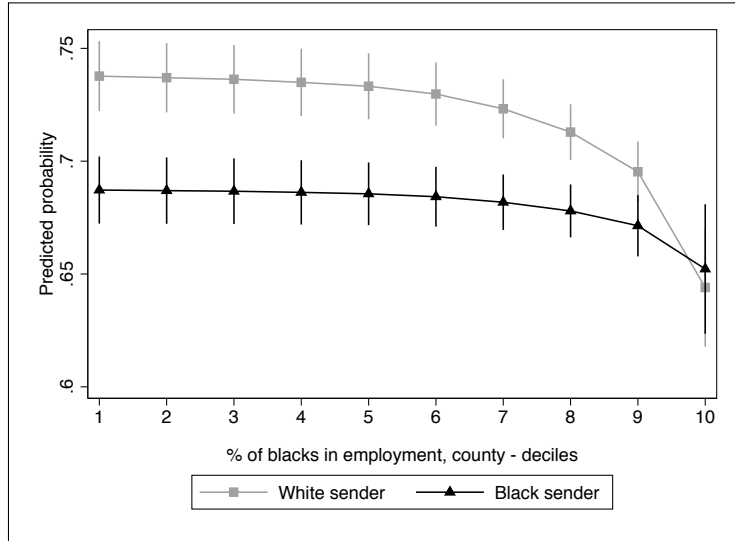
Robust standard errors in parentheses clustered at the state/public service type level. Dependent variable is a binary variable indicating whether a response to the email was provided (linear probability model). \bar{Y} refers to the average response rate. % of blacks in employment refers to the county share of blacks among the employed population. Data have been obtained from the 2006-2010 American Community Survey (ACS). % of blacks in employment refers to the county share of blacks among the population employed in the public sector. Data have been obtained from the 2006-2010 American Community Survey (ACS) and are available for 573 counties. % of blacks among educators, librarians and sheriffs refers to the county share of blacks working in the following occupations: Education Administrators, Librarians, Library Technicians, Library Assistants, Clerical and Sheriffs, Bailiffs, Correctional Officers, and Jailers. Data have obtained from Census 2000 and ACS 2001 to 2015 and are available for 469 counties. % of blacks among school principals refers to the state share of blacks working as school principals. Data have obtained from the 2011-2012 Schools and Staffing Survey and are available for 28 states. All regressions include controls of col IV, Table 3.

results.¹⁹

To facilitate an interpretation of these estimates, Figure 3 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% (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 African-American recipients are less likely to ignore black emails, supporting the interpretation that the estimated gap reflects taste-based 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 data (except for Job Centers) is provided with a contact person, we have name and surname information for both respondents and non-respondents. For each surname, we compute two indices for the

¹⁹ In particular, in the case of the model with the share of blacks in employment, the estimate for the baseline category (1st quartile) is -0.073 (s.e. 0.014) and the estimates for the interactions between the 2nd, 3rd and 4th quartiles and the black dummy are: 0.028 (s.e. 0.017), 0.045 (s.e. 0.020) and 0.063 (s.e. 0.020).

Figure 3: Race of Recipient



Predictions from col II of Table 7.

The x-axis represents deciles of the share of blacks in employment. 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.


“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.²⁰ The idea is to proxy for the 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 African Americans, 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 Sickler). Lower thresholds include surnames that are less characteristic, 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 4, which

²⁰ Since we could not match all surnames, we use both waves for this exercise. We were able to match surnames for 33,372 observations (87.43% of the sample). Out of the unmatched 4,796 observations, 1,452 relate to Job Centers (for which we did not have a contact name) and 3,344 pertain to surnames that were present in our data but not in the Census (which reports only surnames with a frequency higher than 100).

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 imprecise. Again, these results are consistent with the taste-based discrimination interpretation of the differential response rate that we find.

We also explore whether there exists a negative association between racial segregation and discrimination, as the contact hypothesis (Allport, 1954) would suggest. We estimate models similar to Table 7 using indices for residential and school segregation instead of share of blacks. In the case of residential segregation, we found a negative, statistically significant coefficient for the interaction term (-0.018, s.e. 0.006). In the case of school segregation, we found a positive but weaker and statistically insignificant coefficient for the interaction term (0.015, s.e. 0.014). We report these results in Table C.2 of the Appendix.

Figure 4: 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.

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).²¹ The literature generally finds no significant income differences by race.²² 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. The idea behind this test is that, similar to Oreopoulos (2011), if providing additional information regarding the sender changes the racial gap in terms of responsiveness by public officials, then this would be an indication of statistical discrimination playing a role. We may have an increase or a decrease in the gap, depending on how black and white real estate agents compare to generic citizens with black and white sounding names in terms of – for instance – the perceived effort that a public officer expects to exert when dealing with a query from them. Also, the overall response rate may increase or decrease. The critical aspect behind this test is that information about socioeconomic background may affect only statistical discrimination, as this depends on information other than race, while animus should depend only on information about race. However, it is possible that providing information about socioeconomic status shifts the expected cost of dealing with blacks and whites in a parallel way or that the signal of socioeconomic status is uninformative or goes completely unnoticed, thus leaving the gap unchanged even if statistical discrimination plays a role. Thus, while finding a change in the gap would be an indication of the importance of statistical discrimination, not detecting a change is not necessarily an indication of a lack of importance. In Appendix D, we provide a simple analytical framework to illustrate these points.

Looking at the results, the overall response rate to the second wave’s emails is slightly lower than the first wave (63.24%). 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 3.68 percentage points, almost identical to what we estimated in the first wave where there was no occupation signal. Under the assumption that recipients notice the profession indicated in the signature, this provides evidence that the differential response to white versus African-American 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 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.²³ In columns IV to VI, we estimate the same regressions using cordial reply as the

²¹ http://www.bls.gov/oes/current/oes_nat.htm.

²² For instance, in their study on almost 7,000 real estate licensees in the U.S. for 1999, Benjamin et al. (2007) find no significant differences in hourly wages by race. The same is true for Sirmans and Swicegood (2000), who focuses on Texas, and Sirmans and Swicegood (1997), who look at Florida.

²³ In addition, exploiting the fact that half of our recipients in the second wave have already seen an email signed by a white name and the other half have received an email signed with an African-American name, we checked whether the racial difference in response rates in the second wave depends on whether the email in the first wave was signed by a black or a white name. To this end, we estimate the baseline model using only wave II, including an indicator

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.

To summarize, through this exercise we fail to detect the presence of statistical discrimination.²⁴

Table 8: Fixed Socioeconomic Background

	Reply			Cordial reply		
	Wave II	Pooled waves		Wave II	Pooled waves	
Black	-0.037*** (0.007)	-0.038*** (0.005)	-0.039*** (0.006)	-0.061*** (0.008)	-0.061*** (0.006)	-0.062*** (0.007)
Black \times Wave II			0.002 (0.008)			0.002 (0.011)
\bar{Y}	0.632	0.665	0.665	0.702	0.705	0.705
R^2	0.051	0.052	0.052	0.105	0.107	0.107
N	19,089	38,168	38,168	12,072	25,393	25,393

Robust standard errors in parentheses 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 (linear probability model). \bar{Y} refers to the average response rate (columns I-III) and to the average of the cordiality outcome (columns IV-VI).

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

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

All regressions include controls of col IV, Table 3.

4.3 Relationship with Existing Racial Prejudice Measures

We next explore the correlation of our measure of discrimination with racial prejudice measures used in the literature. In particular, we considered three distinct measures: the average prejudice index used by Charles and Guryan (2008), the racially charged search rate used by Stephens-Davidowitz (2014) and an index capturing the race implicit association test created within the Harvard Project Implicit and described in Xu et al. (2014). In Table 9 we report the results of models where we regress our discrimination index (the same that appears in Figure ?? with opposite sign) on racial prejudice indices used in the above-mentioned studies, and which are available at the state level.

for whether the recipient has received an email signed by a white name in wave I and its interaction with our key indicator *Black*. The estimate of the interaction term is very small and insignificant (0.001, s.e. 0.013), suggesting that the racial gap we estimate in wave II does not depend upon the race of sender in wave I.

²⁴ To further probe the potential sources of statistical discrimination in our setting, we looked into the possibility that in areas with high electoral competition public officials may be more likely to worry about a “vocal voter” (as signaled by the sender’s occupation in Wave II) who is more prone to complain or communicate with politicians. Using county-level data from the 2012 presidential elections, we constructed the “effective number of parties” – an index measuring how tight election competitiveness is and that is widely employed in political sciences (see Laakso and Taagepera, 1979). Adding this index to our baseline regression and interacting it with a wave II dummy, we find that the interaction is statistically insignificant, suggesting that the above-mentioned mechanism may not be a relevant driver of statistical discrimination in our case.

Despite that existing measures capture dimensions of discrimination different than ours, and that the samples and time periods are different, we find a moderate positive correlation with our index. This correlation becomes stronger when controlling for geographic area by introducing dummies for the four U.S. regions. Together with the results presented in the previous Subsection, this evidence suggests that prejudice-based discrimination is an important explanation behind our finding.

Table 9: Relationship with racial animus indices

	Average prejudice index		Racially charged search rate		Race implicit association test	
Discrimination index	0.257* (0.143)	0.401*** (0.097)	0.187 (0.140)	0.260** (0.109)	0.282** (0.137)	0.292*** (0.085)
R^2	0.070	0.643	0.035	0.519	0.080	0.708
N	45	45	51	51	51	51
Region dummies	N	Y	N	Y	N	Y

Standard errors in parentheses.

The dependent variable is indicated in the column headers.

The independent variable is the index of discrimination in local public services at the state level constructed using Wave I and II. See text for details.

Average prejudice index refers to the measure used in Charles and Guryan (2008); Racially charged search rate refers to the measure used in Stephens-Davidowitz (2014); Race implicit association test refers to the index created within Project Implicit and described in Xu et al. (2014).

5 Discrimination and Racial Gaps in Attitudes towards the Government and Wages

To provide an initial exploration into whether the discrimination that we found has real consequences, in this Section, we explore if our measure of discrimination is associated with racial gaps in attitudes towards government and wages.

5.1 Attitudes towards government

To analyze attitudes towards government, we use data from the American National Election Studies (ANES), an opinion survey on voting attitudes and behavior. To increase sample size and since some of the questions of interest are asked only in some years, we pooled all available waves from 1948 to 2012. We constructed several indicators ranging from trust in government to interest in elections. The details on how the variables have been constructed and the years for which they are available are reported in Table A.4 in the Appendix. We then match information on the state of residence of respondents in the ANES with our discrimination index and estimate the following

regression:

$$Y_i = \alpha + \beta_1 \text{DiscriminationIndex}_s + \beta_2 \text{Black}_i + \beta_3 \text{Black}_i \times \text{DiscriminationIndex}_s + X_i' \gamma + \varepsilon_i, \quad (2)$$

where Y is the outcome of interest (e.g., faith in government) of individual i , $\text{DiscriminationIndex}$ is the discrimination index at the state level derived from our data (as described in Section 4.3) and Black is a variable equal to 1 if the respondent's reported race is black and 0 if it is white. X is a vector that includes the following control variables: sex, age, education, marital status, social class, occupation, work status, race-specific year dummies and four indicators for the region of residence. The coefficient of interest is β_3 , which represents how the white-black gap in the outcome of interest varies with our measure of discrimination in local public services. The results of this analysis are reported in Table 10.

Table 10: Black-white Gaps in Attitudes towards Government

	Faith in local government		Government performance rating			Govt poor attention to people	Not interested in elections
	Most	Least	Local	State	Federal		
Black	-0.101** (0.044)	0.154** (0.059)	-0.081* (0.046)	-0.031 (0.042)	-0.058 (0.042)	0.135** (0.055)	-0.070 (0.044)
Discrimination index	-0.299 (0.260)	0.161 (0.271)	-0.024 (0.251)	0.043 (0.263)	-0.204 (0.235)	0.048 (0.137)	-0.216** (0.087)
Black \times Discrimination index	-0.548* (0.286)	0.057 (0.294)	-0.470* (0.253)	-0.426 (0.430)	-0.958** (0.454)	0.441 (0.313)	0.433*** (0.147)
\bar{Y}	0.353	0.393	0.771	0.788	0.700	0.294	0.256
R^2	0.024	0.025	0.024	0.032	0.042	0.055	0.091
N	4,547	4,156	4,314	4,330	4,412	19,872	27,832

Robust standard errors clustered at the state level in parentheses. Data are from the 1948-2012 American National Election Studies. The sample includes individuals aged 16 or above.

The dependent variable is indicated in the column headers.

The key independent variable is the index of discrimination in local public services at the state level constructed using Wave I and II. See text for details.

Controls include a dummy for sex, age, race-specific year dummies, region dummies and dummies for: education, marital status social class, occupation and work status.

Observations are weighted by the weights provided in the ANES.

The outcome in column I is an indicator which is equal to one if the individual has most faith in local government and zero if the individual has most faith in state/federal government. The second column considers a similar outcome, but individuals are asked about the level of government in which they have least trust. The estimates of β_3 for both regressions indicate that the white-black gap in terms of faith in local government (compared to other levels of government) is larger in states where discrimination in public services is higher, albeit results are statistically significant at the 10% level only in column I.

The outcomes in columns III to V capture individuals' subjective rating of government's performance. We consider all levels of government (local/state/federal) and construct an indicator

which is equal to 1 if the rating is above the median and 0 otherwise. The point estimates of β_3 suggest that African Americans have an even lower rating of government performance if they reside in states where the value of the discrimination index is higher. However, coefficient estimates are statistically significant for local and federal governments only.

Next, we consider individuals' opinions about how much the government pays attention to what people think. The outcome variable is defined as 1 if the individual thinks that the government is not paying sufficient attention and 0 otherwise. Estimates suggest that while African Americans living in states where discrimination in public services is low are less likely to report a negative opinion, the opposite happens for those living in states with higher discrimination. However, the estimate of β_3 is not statistically significant.

Finally, we consider the link between discrimination in local public services and voting attitudes and behavior by using an indicator which equals 1 if individuals are not interested in elections and 0 otherwise. The estimate of β_3 is positive and statistically significant at conventional values, suggesting that in states with higher discrimination, African Americans are more likely to show lower political interest.²⁵

5.2 Racial wage gap

To look at racial wage gap, we use data from the Current Population Survey (CPS) for individuals aged 16 to 64 in full-time employment. In Table 11 we relate log weekly wages to race and the discrimination index. What emerges is that African Americans earn less, both overall and looking at subsamples characterized by gender or education (above/below high school). The interactions between the black dummy and our index are all negative, indicating that the gap is wider in states displaying higher discrimination. The coefficient of the interaction is however significant only for the low-education subsample.

Overall, in this Section we find that our measure of discrimination has some correlation with attitudes towards government and racial wage gaps, with states characterized by higher discrimination displaying more negative attitudes towards government by African Americans, as well as wider wage gaps.

²⁵ We also explored the correlation between our discrimination index and racial gaps in voting turnout. We accessed voting data by race related to the 2012 presidential elections – which are available for 34 states – and constructed a voting gap variable defined as the difference in the shares of white and black people who voted. We then regressed this variable on the discrimination index. We found a positive, albeit small and statistically insignificant correlation (0.088, s.e. 0.118). Results are nearly unchanged when adding four indicators for the U.S. regions in the regression (the estimate for the voting gap is 0.024, s.e. 0.085).

Table 11: Black-white Wage Gaps

	All	Males	Females	High school degree or below	Above high school degree
Black	-0.136*** (0.013)	-0.206*** (0.017)	-0.066*** (0.019)	-0.125*** (0.022)	-0.142*** (0.014)
Discrimination index	-0.242 (0.224)	-0.268 (0.228)	-0.200 (0.276)	-0.006 (0.241)	-0.338 (0.249)
Black \times Discrimination index	-0.379 (0.241)	-0.326 (0.321)	-0.459 (0.283)	-0.788* (0.394)	-0.199 (0.222)
R^2	0.334	0.324	0.325	0.205	0.283
N	27,799	15,102	12,697	9,717	18,082

Robust standard errors clustered at the state level in parentheses. Data are from the 2013-2015 May CPS Merged Outgoing Rotation Group. The sample includes individuals aged 16 to 64 who are in full-time employment.

The dependent variable is the log weekly wage. Top and bottom 1% of the wage distribution are trimmed. Controls include a dummy for sex, dummies for education categories, a quadratic on experience, year dummies and region dummies.

Observations are weighted by the CPS basic final weight.

6 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 to receive a reply than those from a person with a distinctively white name. In our context, statistical discrimination is unlikely to be of first order importance and a series of tests indeed points towards the role of animus in explaining our findings.

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 low-level bureaucrats are responsible for the implementation of policy enacted at both the federal and state level.

One 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. It is indeed 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 analysis in the previous Section showed that discrimination is associated with less engagement with government (as measured by interest in elections).

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). At the practical level, instituting modes of communication with citizens that do not disclose the name of the sender (e.g., through online forms) rather than through emails could also be useful. 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.

A Data Sources and Variable Definitions

Table A.1: 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 A.2: 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?

[†] Background checks are used, for instance, by prospective employers and landlords.

Table A.3: Data Sources of County Variables

Variable/data	Description	Observational unit	Source	Year	Weblink
% of blacks in employment	Number of blacks employed over total employed population	County	American Community Survey	2010	http://factfinder.census.gov/
% of blacks in public sector	Number of blacks employed in public sector over total population employed in public sector	County	American Community Survey	2006-2010	http://www.census.gov/people/eoetabulation/data/eoetables20062010.html
% of blacks among educators, librarians and sheriffs	Number of blacks employed as Education Administrators, Librarians, Library Technicians, Library Assistants, Clerical and Sheriffs, Bailiffs, Correctional Officers, and Jailers over total population employed in the same occupations	County	Census and American Community Survey	2000-2015	https://usa.ipums.org/usa/
% of blacks among school principals	Number of black school principals over total number of school principals	State	National Center for Education Statistics	2011/2012	https://nces.ed.gov/surveys/sass/tables/sass1112_2013313_p1s_001.asp
% of Hispanic	Number of hispanic over total population	County	Census: Profile of general population	2010	http://factfinder.census.gov/
Unemployment rate	Number of unemployed people over population in labor force	County	Local Area Unemployment Statistics	2014	http://www.bls.gov/lau/#cntyaa
Average labor income	Average weekly wage in US dollars based on the 12-monthly employment levels	County	Quarterly Census of Employment and Wages	2014	http://www.bls.gov/cew/datatoc.htm
Crime rate	Number of violent and propriety crime over total population	County	Uniform Crime Reporting Program Data	2012	http://enceladus.isr.umich.edu/race/racestart.asp
% of Dem votes	Number of votes given to Barack Obama over total votes during 2012 Presidential elections.	County	The Guardian and Alaska.gov	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 "Large central metro, Large fringe metro, Medium metro, Small metro and Metropolitan"	County	National Center for Health Statistics	2013	http://www.cdc.gov/nchs/data_access/urban_rural.htm

Table A.4: Data Sources of Ancillary Variables

Variable/data	Description	Observational unit	Source	Year	Weblink
Dissimilarity index (residential segregation)	Calculated as $\frac{1}{2} \sum_i \left \frac{b_i}{B} - \frac{w_i}{W} \right $ where b_i is the black population of block i , B is the total black population in the tract, w_i is the white population of block i , W is the total white population in the tract. The county-level index is obtained as a weighted average of the tract-level dissimilarity indices	County	Racial Residential Segregation Measurement Project	2000	http://enceladus.isr.umich.edu/race/racestart.asp
Dissimilarity index (school segregation)	Calculated as $\frac{1}{2} \sum_i \left \frac{b_i}{B} - \frac{w_i}{W} \right $ where b_i is the number of black students in the school district, w_i is the number of white students in the school district, B is the number of black students in the school district, W is the number of white students in the school district.	School District	National Center for Education Statistics	2014/2015	https://nces.ed.gov/ccd/pubschuniv.asp
Average prejudice index	Based on mean of responses to prejudice questions (see Charles and Guryan, 2008 for details)	State	General Social Survey	1972-2004	Obtained from the authors
Racially charged search rate	Google searches that include a racial epithet over total Google searches	State	Google Trends	2014/2015	http://sethsd.com/research/
Race implicit association test	Index that measures how long it takes to a person to sort associations between African-American and Euro-American faces with evaluations (good/bad). Higher level of the index mean stronger implicit preference for Euro-American faces. Observations aggregated by state. African Americans are excluded.	State	Harvard Project Implicit	2014	https://osf.io/qbqq3/
Most (least) faith in local government	Equals one if the individual has most (least) faith in local government and zero if the individual has most faith in state/federal government and zero otherwise	Individual	American National Election Studies	1968/1996	http://www.electionstudies.org/study pages/anes-timeseries_cdf/anes-timeseries_cdf.htm
Local/State/Federal Government performance rating	Index ranging from 0 (lowest performance) to one (highest performance)	Individual	American National Election Studies	1974/1980	http://www.electionstudies.org/study pages/anes-timeseries_cdf/anes-timeseries_cdf.htm
Government poor attention to people	Equals one if the individual feels that government is paying not much attention to what people think and 0 otherwise	Individual	American National Election Studies	1964/2004	http://www.electionstudies.org/study pages/anes-timeseries_cdf/anes-timeseries_cdf.htm
Not interested in elections	Equals one if the individual is not much interested in the elections and 0 otherwise	Individual	American National Election Studies	1952/2012	http://www.electionstudies.org/study pages/anes-timeseries_cdf/anes-timeseries_cdf.htm
Wages	Log weekly wages of working-age population	Individual	Current Population Survey	2011-2015	https://cps.ipums.org/cps/

B Additional Tables and Figures

Table B.1: Details of Emails

Recipient	N. recipients	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%

Notes: N recipients refer to the existing number of recipients (i.e., potential recipients, including those with no email address). Undelivered/testing refers to emails that were not part of the final sample.

Table B.2: County Characteristics of Recipients in Sample and not in Sample

	School D		Library		Sheriff	
	In sample	Not in sample	In sample	Not in sample	In sample	Not in sample
% of blacks in employment	7.63 (13.28)	9.84*** (14.06)	5.02 (9.72)	11.28*** (16.15)	7.77 (13.02)	7.89 (13.83)
Unemployment rate (%)	6.23 (2.29)	6.68*** (2.38)	5.95 (2.15)	6.66*** (2.42)	6.11 (2.19)	6.50*** (2.44)
% of Hispanic	8.26 (12.98)	8.53 (15.29)	8.64 (13.55)	7.85* (12.72)	8.70 (13.5)	7.69** (12.72)
Average labor income (USD)	716 (190)	715 (201)	729 (199)	701*** (179)	721 (186)	710 (197)
Crime rate (%)	2.42 (1.48)	2.25* (1.64)	2.27 (1.37)	2.56*** (1.62)	2.48 (1.52)	2.29*** (1.44)
% of Dem votes	38.75 (14.76)	36.90* (15.43)	38.85 (14.49)	38.26 (15.24)	38.77 (14.35)	38.33 (15.48)
Urbanization	58.43 (49.29)	47.97*** (50.05)	58.40 (49.3)	56.44 (49.6)	61.11 (48.76)	52.49*** (49.96)
N counties	2,872	271	1,738	1,405	1,836	1,307

	Treasurer		Job Center V.R.		County Clerk	
	In sample	Not in sample	In sample	Not in sample	In sample	Not in sample
% of blacks in employment	3.97 (8.52)	9.97*** (15)	10.04 (13.63)	7.24*** (13.23)	3.84 (6.41)	8.79*** (14.4)
Unemployment rate (%)	5.53 (2.1)	6.69*** (2.31)	6.37 (2.04)	6.24 (2.37)	6.12 (1.92)	6.31* (2.39)
% of Hispanic	10.52 (15.86)	7.03*** (11.24)	9.22 (12.71)	8.04** (13.31)	10.49 (14.84)	7.75*** (12.7)
Average labor income (USD)	712 (175)	719 (199)	803 (213)	694*** (178)	713 (188)	717 (191)
Crime rate (%)	2.18 (1.42)	2.52*** (1.52)	3.44 (1.49)	2.13*** (1.37)	2.41 (1.33)	2.40 (1.53)
% of Dem votes	36.55 (14.44)	39.73*** (14.92)	44.75 (14.36)	36.98*** (14.53)	34.96 (13.6)	39.47*** (14.98)
Urbanization	53.59 (49.89)	59.73*** (49.06)	86.88 (33.79)	49.90*** (50.01)	56.66 (49.6)	57.74 (49.41)
N counties	1,129	2,014	648	2,495	616	2,527

Notes: Standard deviations in parentheses.

* p < 0.10; ** p < 0.05; *** p < 0.01. P-values refer to the statistical significance of the two-sample t-statistic for the difference between the mean characteristics of "In sample" and "Not in sample" counties.

Table B.3: Response Rates - by Sending Name

	DeShawn	Tyrone	Total	Greg	Jake	Total
	Jackson	Washington	Black	Walsh	Mueller	White
Response rate	69.05 (46.23)	66.91 (47.06)	67.96 (46.67)	71.57 (45.11)	71.74 (45.03)	71.66 (45.07)
N	4,637	4,835	9,472	4,918	4,689	9,607
Difference within race (abs)	2.15		0.17			
z-stat (p-val)	0.025		0.855			
Difference B-W	-2.60	-4.75				
z-stat (p-val)	0.001	0.000				

Notes: Figures refer to response rates multiplied by 100. Standard deviations in parentheses.

Table B.4: Response Rates - by Type of Recipient

	School D.	Library	Sheriff	Treasurer	Job Center	County Clerk	Total
White	76.51 (42.4)	69.08 (46.23)	53.23 (49.92)	73.90 (43.96)	71.93 (44.99)	65.72 (47.54)	71.66 (45.07)
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	-3.41	-4.12	-6.98	-4.33	1.14	-1.63	-3.70
z-stat (p-val)	0.000	0.002	0.003	0.107	0.730	0.672	0.000

Notes: Figures refer to response rates multiplied by 100. Standard deviations in parentheses.

Table B.5: Other Outcomes - Summary Statistics

	Black	White	t-test (pval)
Number of replies	1.03 (0.18)	1.03 (0.19)	0.78
	[1-3]	[1-6]	
Length of reply (# words)	171.4 (99.2)	170.7 (101.2)	0.68
	[0-4213]	[2-5324]	
Cordial reply	0.679 (0.467)	0.735 (0.441)	0
	[0-1]	[0-1]	
Delay in reply (# hours)	24.2 (66.8)	25.2 (73.9)	0.41
	[0-960]	[0-960]	

Notes: N=13321. Cordial reply refers to whether the sender was addressed by name or with salutations. Standard deviations in parentheses. Values in square brackets represent the min and max, respectively.

C Additional Analysis

Table C.1: Cordiality in response - by Type of Public Service

	School D.	Library	Sheriff	Treasurer	Job Center	County Clerk
Black	-0.060*** (0.009)	-0.047*** (0.014)	-0.063** (0.029)	-0.135*** (0.031)	-0.037 (0.047)	-0.098* (0.051)
\bar{Y}	0.771	0.722	0.462	0.503	0.643	0.502
R^2	0.088	0.069	0.168	0.143	0.174	0.192
N	7,385	3,281	914	811	530	400

Robust standard errors in parentheses clustered at the state level.

Dependent variable is a binary variable indicating whether the sender was addressed by name or with salutations (linear probability model). \bar{Y} refers to the average of the cordiality outcome

All regressions include controls of col IV, Table 3.

Table C.2: Residential and school segregation

	Residential		School	
Black	-0.038*** (0.006)	-0.038*** (0.006)	-0.031*** (0.010)	-0.031*** (0.011)
Segregation index	-0.005 (0.005)	0.004 (0.006)	-0.001 (0.006)	-0.008 (0.007)
Black \times Segregation index		-0.018*** (0.006)		0.015 (0.014)
R^2	0.049	0.050	0.040	0.041
N	19,079	19,079	5,216	5,216


Robust standard errors in parentheses clustered at the state/recipient 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, Table 3.

Residential segregation is measured using the dissimilarity index for all counties. The index has been obtained from the Racial Residential Segregation Measurement Project website and is defined as $\frac{1}{2} \sum_i \left| \frac{b_i}{B} - \frac{w_i}{W} \right|$ where b_i is the black population of block i , B is the total black population in the tract, w_i is the white population of block i , W is the total white population in the tract. The county-level index is obtained as a weighted average of the tract-level dissimilarity indices (see <http://enceladus.isr.umich.edu/race/racestart.asp> for details). School segregation is measured using the dissimilarity index for school districts obtained from the 2014-15 Common Core of Data from the National Center for Education Statistics. Districts with less than 10 black students have been dropped. The index of dissimilarity is defined as above, where: b_i is the number of black students of school i , B is the number of black students in the school district, w_i is the number of white students of school i , W is the number of white students in the school district. Both indices of segregation have been standardized.

Figure C.1: Difference in Response Rates by Census Divisions



Black/white gap in response rate is obtained by estimating the model of col IV in Table 3 for each Census Division. Bars represent confidence intervals.

NE includes Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island and Vermont; MA includes New Jersey, New York and Pennsylvania; ENC includes Illinois, Indiana, Michigan, Ohio and Wisconsin; WNC includes Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota and South Dakota; SA includes Delaware, District of Columbia, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia and West Virginia; ESC includes Alabama, Kentucky, Mississippi and Tennessee; WSC includes Arkansas, Louisiana, Oklahoma and Texas; MO includes Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah and Wyoming; PA includes Alaska, California, Hawaii, Oregon and Washington.

D Framework for Public Official’s Decision Problem

Here, we sketch a simple analytical framework, along the lines of Neumark (2012) and Heckman (1998), to think about the decision-making of a public official who has received an email and needs to decide whether to answer or not.

Suppose that the public official applies the following decision rule when it comes to responding to email queries: a response will be provided if the perceived cost of the response is below a certain threshold c . The cost here can be associated with the effort that the official would need to exert if the citizen asking the initial question (e.g., opening hours of library), comes back with further questions or actually decides to visit the office and interact with the official. In this perspective, the official does not respond to citizens that are likely to be very demanding in terms of the official’s time and effort in future interactions.²⁶

Let us denote the fussiness of the citizen by F and suppose that fussiness depends on the citizen’s characteristics, X . The treatment that the citizen receives, T , will depend on the fussiness of the citizen and possibly also on the race, R (with $R = B$ for blacks and $R = W$ for whites). Let us assume for simplicity that $T(F(X), R) = X + \gamma I_R$, with $I_R = 1$ if $R = B$ and $I_R = 0$ if $R = W$, and $\gamma > 0$ capturing taste-based discrimination against African Americans. The treatment in this example is continuous. However, in our setting the treatment of senders is dichotomous (reply or not reply). As mentioned, this can be modelled through a cut-off rule, so that a response is offered ($T = 1$) if the perceived cost of the response is below a certain threshold:

$$T(F(X_W), W) = 1 \text{ if } E(X_W) < c, \text{ for white senders} \quad (3)$$

and

$$T(F(X_B), B) = 1 \text{ if } E(X_B) + \gamma < c, \text{ for black senders,} \quad (4)$$

where c is the cost of effort for the official, γ is the taste-based discrimination factor, and $E(X_W), E(X_B)$ are, respectively, the expected personal characteristics of citizens with white- and black-sounding names.

The difference in responses to whites and blacks in wave 1 (what we refer to as the racial gap) is $\Delta T_{wave1} = T(F(X_W), W) - T(F(X_B), B) = E(X_W) - E(X_B) - \gamma$, and is a combination of taste-based discrimination, γ , and statistical discrimination, $E(X_W) - E(X_B)$.

In wave 2, we attempt to affect the extent of statistical discrimination, $E(X_W) - E(X_B)$, present in wave 1. The logic of what we do in wave 2 is to try to change the term $E(X_W) - E(X_B)$ by adding information about the identity of the sender. Then, if we see the racial gap change across the two waves, we can deduce that this change can be attributed to a change in the extent of

²⁶Alternatively, the cost may be interpreted as the penalty inflicted to the official by the supervisor in case of no response and complain by the citizen. In that case, the decision rule is to respond to queries that have a perceived cost that is above a certain threshold, that is, respond to citizens that are more likely to complain.

statistical discrimination across the two waves and, thus, obtain a lower (upper) bound for statistical (taste-based) discrimination.²⁷ On the other hand, taste-based discrimination depends only on the perceived race of the sender and, thus, is insensitive to additional information. From now on, to make the argument sharper we assume a continuous treatment.²⁸ The intuition that information about socioeconomic background may affect only statistical discrimination holds, however, also to the case with dichotomous treatment. This is because statistical discrimination depends on information other than race, while taste-based discrimination should depend only on information about race. Indicating with $E(X_W|RealEstate)$ and with $E(X_B|RealEstate)$ the expected fussiness when the real estate signal is present of an agent with a white and black sounding name respectively, then we should consider three cases.

Case (i):

$$E(X_W|RealEstate) - E(X_B|RealEstate) < E(X_W) - E(X_B), \quad (5)$$

that is, the official expects real estate agents to be closer in terms of expected fussiness than wave 1 senders. In this case, we would expect to see a reduction in the racial gap comparing wave 1 to wave 2, and this reduction would be evidence of the presence of statistical discrimination.

Case (ii):

$$E(X_W|RealEstate) - E(X_B|RealEstate) > E(X_W) - E(X_B), \quad (6)$$

that is, the official expects real estate agents to be further apart in terms of expected fussiness than wave 1 senders. In this case, we expect to find an increase in the racial gap comparing wave 1 to wave 2, and this increase would, again, be evidence of the presence of statistical discrimination.

Case (iii):

$$E(X_W|RealEstate) - E(X_B|RealEstate) = E(X_W) - E(X_B), \quad (7)$$

that is, there is a parallel shift in terms of expected fussiness. In this knife-edge case, we would expect to find no change in the racial gap, regardless of the presence or not of statistical discrimination, and comparing wave 1 to wave 2 would not provide any evidence regarding the presence of statistical discrimination.

²⁷Oreopoulos (2011) applies this logic in the context of a résumé correspondence study by adding information about the candidates' unobserved productivity and observing whether it leads to a reduction of the gap in call back rates across native and immigrant job candidates. He finds little evidence of statistical discrimination, similarly to us.

²⁸The assumption of a continuous treatment is not innocuous. Heckman and Siegelman (1993), Heckman (1998) and Neumark (2012) show that, when the treatment is nonlinear, as in our context, then even if there are no mean differences in unobservables between the two racial groups, differences in the variance of unobservables may give rise to differences in response due to statistical discrimination. To see this, note that the probability that a white and a black sender receives a response, assuming that X is normally distributed with mean zero, is: $\Phi(c/\sigma_w)$ and $\Phi((c - \gamma)/\sigma_b)$, where Φ denotes the standard normal distribution and σ_w, σ_b are the standard deviations for whites and blacks respectively. One can readily see that, if the standard deviations are unequal, then even in the absence of differences in mean characteristics, one can observe different probabilities of response across the two racial groups also in the absence of taste-based discrimination, i.e. $\gamma = 0$.

Notice that what has been discussed above is true also if the signal is negative, meaning that a real estate agent of a given race is expected to be more fussy than a generic sender of the same race, that is, $E(X_W|RealEstate) > E(X_W)$ and $E(X_B|RealEstate) > E(X_B)$. This would reduce the probability of responses in the second wave, which is something that we indeed find, but would still allow us to detect the presence of statistical discrimination by comparing wave 1 to wave 2, as long as we are not in case (iii) above, that is, as long as $E(X_W|RealEstate) - E(X_B|RealEstate) \neq E(X_W) - E(X_B)$.

To summarize, as long as the change in the perceived fussiness of a black citizen when the real estate identity is attached is different from that of his white counterpart, we can still detect the presence and provide a lower bound for the importance of statistical discrimination by comparing the racial gap in response rates across waves.

Finding, as we do, that the racial gap is constant across waves is consistent with absence of statistical discrimination if $E(X_W|RealEstate) - E(X_B|RealEstate) < E(X_W) - E(X_B)$ or $E(X_W|RealEstate) - E(X_B|RealEstate) > E(X_W) - E(X_B)$. If, instead, $E(X_W|RealEstate) - E(X_B|RealEstate) = E(X_W) - E(X_B)$ or if the signal goes completely unnoticed, then comparing the two waves is not informative about the presence or not of statistical discrimination.

References

- Abrams, David, Marianne Bertrand, and Sendhil Mullainathan (2012). “Do Judges Vary in Their Treatment of Race?” *Journal of Legal Studies*, 41(2), 347–383.
- Alesina, Alberto and Eliana La Ferrara (2014). “A Test of Racial Bias in Capital Sentencing.” *American Economic Review*, 104(11), 3397–3433.
- Allport, Gordon W. (1954). *The Nature of Prejudice*. New York: Addison.
- Altonji, Joseph G. and Rebecca M. Blank (1999). “Race and Gender in the Labor Market.” In *Handbook of Labor Economics 3C*, edited by O. Ashenfelter and D. Card, chap. 48, pp. 3143–3259. New York: Elsevier.
- Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross (2017). “What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders.” *Review of Financial Studies*, (forthcoming).
- Becker, Gary S. (1957). *The Economics of Discrimination*. University of Chicago Press, Chicago.
- Benjamin, John D., Peter T. Chinloy, G. Donald Jud, and Daniel T. Winkler (2007). “Do Some People Work Harder than Others? Evidence from Real Estate Brokerage.” *The Journal of Real Estate Finance and Economics*, 35(1), 95–110.

- Bertrand, Marianne and Esther Duflo (2017). “Field Experiments on Discrimination.” *Handbook of Economic Field Experiments*, 1, 309–393.
- 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.
- Bhargava, Saurabh and Dayanand Manoli (2015). “Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment.” *American Economic Review*, 105(11), 3489–3529.
- Broockman, David 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, 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.
- CDC (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, Kerwin Kofi and Jonathan Guryan (2008). “Prejudice and Wages: An Empirical Assessment of Becker’s The Economics of Discrimination.” *Journal of Political Economy*, 116(5), 773–809.
- Charles, Kerwin Kofi and Jonathan Guryan (2011). “Studying Discrimination: Fundamental Challenges and Recent Progress.” *Annual Review of Economics*, 3(1), 479–511.
- Daponte, Beth Osborne, Seth Sanders, and Lowell Taylor (1999). “Why do Low-Income Households not use Food Stamps? Evidence from an Experiment.” *Journal of Human Resources*, 34(3), 612–628.
- Distelhorst, Greg and Yue Hou (2014). “Ingroup Bias in Official Behavior: A National Field Experiment in China.” *Quarterly Journal of Political Science*, 9(2), 203–230.
- Doleac, Jennifer L. and Luke C.D. Stein (2013). “The Visible Hand: Race and Online Market Outcomes.” *Economic Journal*, 123(572), F469–F492.
- Duflo, Esther and Emmanuel 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, Michael, Bryan Tomlin, and Liang Choon Wang (2014). “Statistical Discrimination or Prejudice? A Large Sample Field Experiment.” *Review of Economics and Statistics*, 96(1), 119–134.
- Fryer, Roland G. (2011). “Racial Inequality in the 21st Century: The Declining Significance of Discrimination.” In *Handbook of Labor Economics 4B*, edited by O. Ashenfelter and D. Card, chap. 10, pp. 855–971. New York: Elsevier.
- Fryer, Roland G. and Steven D. Levitt (2004). “The Causes and Consequences of Distinctively Black Names.” *Quarterly Journal of Economics*, 119(3), 767–805.
- Fryer, Roland G and Paul Torelli (2010). “An Empirical Analysis of Acting White.” *Journal of Public Economics*, 94(5), 380–396.
- Glaeser, Edward L and Bruce Sacerdote (2003). “Sentencing in Homicide Cases and the Role of Vengeance.” *Journal of Legal Studies*, 32(2), 363–382.
- Gneezy, Uri, John List, and Michael K. Price (2012). “Toward an Understanding of why People Discriminate: Evidence from a Series of Natural Field Experiments.” NBER Working Paper 17855.
- Guryan, Jonathan and Kerwin Kofi Charles (2013). “Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots.” *Economic Journal*, 123(572), F417–F432.
- Harrell, Jules P, Sadiki Hall, and James Taliaferro (2003). “Physiological Responses to Racism and Discrimination: An Assessment of the Evidence.” *American Journal of Public Health*, 93(2), 243–248.
- Hastings, Justine S. and Jeffrey M. Weinstein (2008). “Information, School Choice, and Academic Achievement: Evidence from two Experiments.” *Quarterly Journal of Economics*, 123(4), 1373–1414.
- Heckman, James J. (1998). “Detecting Discrimination.” *The Journal of Economic Perspectives*, 12(2), 101–116.
- Heckman, James J. and Peter Siegelman (1993). “The Urban Institute Audit Studies: Their Methods and Findings.” In *Clear and Convincing Evidence: Measurement of Discrimination in America*, edited by M. Fix and R. Struyk. Urban Institute Press.
- Hirschman, Albert O (1970). *Exit, Voice and Loyalty*. Harvard University Press, Cambridge/Massachusetts.
- Hoxby, Caroline and Sarah Turner (2013). “Expanding College Opportunities for High-Achieving, Low Income Students.” Stanford Institute for Economic Policy Research Discussion Paper 12-014.

- Laakso, Markku and Rein Taagepera (1979). ““Effective” Number of Parties: A Measure with Application to West Europe.” *Comparative Political Studies*, 12(1), 3–27.
- Lang, Kevin (2015). “Racial Realism: A Review Essay on John Skrentny’s *After Civil Rights*.” *Journal of Economic Literature*, 53(2), 351–359.
- Levine, Ross, Yona Rubinstein, and Alexey Levkov (2014). “Bank Deregulation and Racial Inequality in America.” *Critical Finance Review*, 3(1), 1–48.
- Lipsky, Michael (1980). *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services*. New York, Russel Sage Foundation.
- Neumark, David (2012). “Detecting Discrimination in Audit and Correspondence Studies.” *Journal of Human Resources*, 47(4), 1128–1157.
- Neumark, David (2016). “Experimental Research on Labor Market Discrimination.” NBER Working Paper 22022.
- Norton, Michael I and Samuel R. Sommers (2011). “Whites See Racism as a Zero-Sum Game That They are now Losing.” *Perspectives on Psychological Science*, 6(3), 215–218.
- Oreopoulos, Philip (2011). “Why do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Résumés.” *American Economic Journal: Economic Policy*, 3(4), 148–171.
- Pew Research Center (2013). *How Americans Value Public Libraries in Their Community*.
- Riach, Peter A. and Judith Rich (2002). “Field Experiments of Discrimination in the Market Place.” *Economic Journal*, 112(483), F480–F518.
- Rich, Judith (2014). “What do Field Experiments of Discrimination in Markets Tell us? A Meta Analysis of Studies Conducted since 2000.” IZA Discussion Paper 8584.
- Sabol, William J., Heather C. West, and Matthew Cooper (2009). “Prisoners in 2008.” *Bureau of Justice Statistics Bulletin*, (NCJ 228417), 1–45.
- Sirmans, G and Philip Swicegood (1997). “Determinants of Real Estate Licensee Income.” *Journal of Real Estate Research*, 14(2), 137–153.
- Sirmans, G. and Philip Swicegood (2000). “Determining Real Estate Licensee Income.” *Journal of Real Estate Research*, 20(1-2), 189–204.
- Smith, Aaron (2010). “The Internet Gives Citizens New Paths to Government Services and Information.” *Pew Internet and American Life Project*.

- Stephens-Davidowitz, Seth (2014). “The Cost of Racial Animus on a Black Candidate: Evidence Using Google Search Data.” *Journal of Public Economics*, 118, 26–40.
- U.S. 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, 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), 1–14.
- Wong, Gloria, Annie O. Derthick, E.J.R. David, Anne Saw, and Sumie 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.
- Xu, Kaiyuan, Brian Nosek, and Anthony Greenwald (2014). “Psychology data from the race implicit association test on the project implicit demo website.” *Journal of Open Psychology Data*, 2(1).