Massachusetts Institute of Technology
Department of Economics
Working Paper Series

Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination

Marianne Bertrand
Sendhil Mullainathan

Working Paper 03-22
May 27, 2003

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
http://papers.ssrn.com/abstract_id=422902

Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination

Marianne Bertrand and Sendhil Mullainathan *

May 27, 2003

Abstract

We perform a field experiment to measure racial discrimination in the labor market. We respond with fictitious resumes to help-wanted ads in Boston and Chicago newspapers. To manipulate perception of race, each resume is randomly assigned either a very African American sounding name or a very White sounding name. The results show significant discrimination against African-American names: White names receive 50 percent more callbacks for interviews. We also find that race affects the benefits of a better resume. For White names, a higher quality resume elicits 30 percent more callbacks whereas for African Americans, it elicits a far smaller increase. Applicants living in better neighborhoods receive more callbacks but, interestingly, this effect does not differ by race. The amount of discrimination is uniform across occupations and industries. Federal contractors and employers who list "Equal Opportunity Employer" in their ad discriminate as much as other employers. We find little evidence that our results are driven by employers inferring something other than race, such as social class, from the names. These results suggest that racial discrimination is still a prominent feature of the labor market.

*University of Chicago Graduate School of Business, NBER and CEPR; MIT and NBER. Addresses: 1101 E. 58th Street, R0 229D, Chicago, IL 60637; 50 Memorial Drive, E52-380a, Cambridge, MA 02142. E-mail: marianne.bertrand@gsb.uchicago.edu; mullain@mit.edu. David Abrams, Victoria Bede, Simone Berkowitz, Hong Chung, Almudena Fernandez, Mary Anne Guediguian, Christine Jaw, Richa Maheswari, Beverley Martis, Alison Tisza, Grant Whitehorn, and Christine Yee provided excellent research assistance. We are also grateful to numerous colleagues and seminar participants for very helpful comments.

1 Introduction

Every measure of economic success reveals significant racial inequality in the US labor market. Compared to Whites, African Americans are twice as likely to be unemployed and earn nearly 25 percent less when they are employed (Council of Economic Advisers, 1998). This inequality has sparked a debate on whether employers discriminate by race. When faced with observably similar African American and White applicants, do they favor the White one? Some argue yes, citing either employer prejudice or employer perception that race signals lower productivity. Others argue that discrimination is a relic of the past, eliminated by some combination of employer enlightenment, affirmative action programs and the profit-maximization motive. In fact, many in this later camp even feel that stringent enforcement of affirmative action programs has produced an environment of reverse discrimination. They would argue that faced with identical candidates, employers might favor the African American one.[1] Data limitations make it difficult to empirically test these views. Since researchers possess far less data on workers than employers do, White and African American workers that appear similar to researchers may look very different to employers. So any racial difference in labor market outcomes could just as easily be attributed to these differences unobserved by researchers as to discrimination.

We conduct a field experiment to circumvent this difficulty. We send resumes in response to help-wanted ads in Chicago and Boston newspapers and measure the number of callbacks each resume receives for interviews. We experimentally manipulate perception of race via the name on the resume. We randomly assign very White sounding names (such as Emily Walsh or Greg Baker) to half the resumes and very African American sounding names (such as Lakisha Washington or Jamal Jones) to the other half. Because we are also interested in how credentials affect discrimination, we experimentally vary the quality of the resumes used in response to a given ad. Higher quality applicants have on average a little more labor market experience and fewer holes in their employment history; they are also more likely to have an email address, have completed some certification degree, possess foreign language skills or have been awarded some honors.[2] In practice, we typically send four resumes in response to each ad: two higher quality and two lower quality ones. We randomly assign to one of the higher and one of the lower quality resumes an African American sounding name. In total, we respond to over 1300 employment ads in the sales, administrative support, clerical and customer services job categories and send nearly 5000 resumes. The ads we respond to cover a large spectrum of job quality, from cashier work at retail establishments and clerical work in a mailroom to office and sales management positions.

We find large racial differences in callback rates.[3] Applicants with White names need to send about

[1] This camp often explains the poor performance of African Americans in terms of supply factors. If African Americans lack many basic skills entering the labor market, then they will perform worse, even with parity or favoritism in hiring.

[2] In creating the higher quality resumes, we deliberately made small changes in credentials so as to minimize the chance of over-qualification.

[3] For ease of exposition, we refer to the effects uncovered in this experiment as racial differences. Technically, however, these

10 resumes to get one callback whereas applicants with African American names need to send around 15 resumes to get one callback. This 50 percent gap in callback rates is statistically very significant. Based on our estimates, a White name yields as many more callbacks as an additional eight years of experience. Since applicants' names are randomly assigned, this gap can only be attributed to the name manipulation.

Race also affects the reward to having a better resume. Whites with higher quality resumes receive 30 percent more callbacks than Whites with lower quality resumes, a statistically significant difference. On the other hand, having a higher quality resume has a much smaller effect for African Americans. In other words, the gap between White and African-Americans widens with resume quality. While one may have expected that improved credentials may alleviate employers' fear that African American applicants are deficient in some unobservable skills, this is not the case in our data.[4] Discrimination therefore appears to bite twice, making it harder not only for African Americans to find a job but also to improve their employability.

The experiment also reveals several other aspects of discrimination. First, since we randomly assign applicants' postal addresses to the resumes, we can study the effect of neighborhood of residence on the probability of callback. We find that living in a wealthier (or more educated or more White) neighborhood increases callback rates. But, interestingly, African Americans are not helped more than Whites by living in a "better" neighborhood. Second, the amount of discrimination we measure by industry does not appear correlated to Census-based measures of the racial gap by industry. The same is true for the amount of discrimination we measure in different occupations. In fact, we find that discrimination levels are statistically indistinguishable across all the occupation and industry categories covered in the experiment. We also find that federal contractors, who are thought to be more severely constrained by affirmative action laws, do not discriminate less; neither do larger employers or employers who explicitly state that they are an "Equal Opportunity Employer" in their ads. In Chicago, we find that employers located in more African American neighborhoods are slightly less likely to discriminate.

The rest of the paper is organized as follows. Section 2 compares this experiment to prior work on discrimination, and most notably to the labor market audit studies. We describe the experimental design in Section 3 and present the results in Section 4.1. In Section 5, we discuss possible interpretations of our results, focusing especially on two issues. First, we examine whether the race-specific names we have chosen might also proxy for social class above and beyond the race of the applicant. Using birth certificates data on mother's education for the different names used in our sample, we find little relationship between social background and the name specific callback rates.[5] Second, we discuss how our results map back

effects are about the racial soundingness of names. We briefly discuss the potential confounds between name and race below and more extensively in Section 5.1.

[4]These results contrast with the view, mostly based on non-experimental evidence, that African Americans receive higher returns to skills. For example, estimating earnings regressions on several decades of Census data, Heckman et al. (2001) show that African Americans experience higher returns to a high school degree than Whites.

[5]We also argue that a social class interpretation would find it hard to explain all of our results, such as why living a better neighborhood does not increase callback rates more for African American names than for White names.

to the different models of discrimination proposed in the economics literature. In doing so, we focus on two important results: the lower returns to credentials for African Americans and the relative homogeneity of discrimination across occupations and industries. We conclude that existing models do a poor job of explaining the full set of findings. Section 6 concludes.

2 Prior Research on Discrimination

With conventional labor force and household surveys, it is difficult to measure racial discrimination or analyze its mechanics.[6] With survey data, researchers usually measure discrimination by comparing the labor market performance of Whites and African Americans who report a similar set of skills. But such comparisons can be quite misleading. Standard labor force surveys do not contain all the characteristics that employers observe when hiring, promoting or setting wages. So one can never be sure that the African American and White workers being compared are truly similar from the employer's perspective. As a consequence, any measured differences in outcomes could be attributed to these unobserved (to the econometrician) factors and not to discrimination.

This difficulty with conventional data has led some authors to use pseudo-experiments.[7] Goldin and Rouse (2000), for example, examine the effect of blind auditioning on the hiring process of orchestras. By looking at the treatment of female candidates before and after the introduction of blind auditions, they try to measure the amount of sex discrimination. When such pseudo-experiments can be found, the resulting study can be very informative, but finding such experiments has been extremely difficult.

A different set of studies, known as audit studies, attempt to place comparable minority and White subjects into actual social and economic settings and measure how each group fares in these settings.[8] Labor market audit studies send comparable minority (African American or Hispanic) and White auditors in for interviews and measure whether one is more likely to get the job than the other.[9] While the results vary somewhat across studies, minority auditors tend to perform worse on average: they are less likely to get called back for a second interview and, conditional on getting called back, less likely to get hired.

These audit studies provide some of the cleanest non-laboratory evidence of labor market discrimination. But they also have weaknesses, most of which have been highlighted in Heckman and Siegelman (1992) and Heckman (1998). First, these studies require that both members of the auditor pair are identical in all dimensions that might affect productivity in employers' eyes, except for race. To accomplish this, researchers

[6] See Altonji and Blank 1999) for a detailed review of the existing literature on racial discrimination in the labor market.

[7] Darity and Mason (1998) describe an interesting non-experimental study. Prior to the Civil Rights Act of 1964, employment ads would explicitly state racial biases, providing a direct measure of discrimination. Of course, as Arrow (1998) mentions, discrimination was at that time "a fact too evident for detection."

[8] Fix and Turner (1998) provide a survey of many such audit studies.

[9] Earlier hiring audit studies include Newman (1978) and McIntyre et al (1980). Three more recent studies are Cross et al (1990), Turner et al (1991), James and DelCastillo (1991). Altonji and Blank (1999), Heckman and Siegelman (1992) and Heckman (1998) summarize these studies. See also Neumark (1996) for a labor market audit study on sex discrimination.

typically match auditors on several characteristics (height, weight, age, dialect, dressing style, hairdo) and train them for several days to coordinate interviewing styles. Yet, critics note that this is unlikely to erase the numerous differences that exist between the auditors in a pair.

Another weakness of the audit studies is that they are not double-blind. Auditors know the purpose of the study. As Turner et al (1990) note: "The first day of training also included an introduction to employment discrimination, equal employment opportunity, and a review of project design and methodology." This may generate conscious or subconscious motives among auditors to generate data consistent or inconsistent with racial discrimination. As psychologists know very well, these demand effects can be strong. It is very difficult to insure that auditors will not want to do "a good job." Since they know the goal of the experiment, they can alter their behavior in front of employers to express (indirectly) their own views. Even a small belief by auditors that employers treat minorities differently can result in apparent discrimination. This effect is further magnified by the fact that auditors are not in fact seeking jobs and are therefore more free to let their beliefs affect the interview process.

Finally, audit studies are extremely expensive, making it difficult to generate large enough samples to understand the nuances and possible mitigating factors of discrimination. Also, these budgetary constraints worsen the problem of mismatched auditor pairs. Cost considerations force the use of a few pairs of auditors, meaning that any one mismatched pair can easily drive the results. In fact, these studies generally tend to find significant differences in outcomes across pairs.

Our study circumvents these problems. First, because we only rely on resumes and not people, we can be sure to generate comparability across race. In fact, since race is randomly assigned to each resume, the same resume will sometimes be associated with an African American name and sometimes with a White name. This guarantees that any differences we find are due solely to the race manipulation. Second, the use of paper resumes insulates us from demand effects. While the research assistants know the purpose of the study, our protocol allows little room for conscious or subconscious deviations from the set procedures. Moreover, we can objectively measure whether the randomization occurred as expected. This kind of objective measurement is impossible in the case of the previous audit studies. Finally, because of relatively low marginal cost, we can send out a large number of resumes. Besides giving us more precise estimates, this larger sample size also allows us to examine the mechanics of discrimination from many more angles.[10]

[10] A similar "correspondence" technique has been used in a few U.K. studies. See Jowell and Precott-Clarke (1970), Brown and Gay (1985) and Hubbock and Carter (1980). These earlier studies have very limited sample size and focus exclusively on documenting gap in callbacks between the minority and non-minority groups. Some of these studies fail to fully match skills between minority and non-minority resumes, for example by imposing differential education background by racial origin.

3 Experimental Design

3.1 Creating a Bank of Resumes

The first step of the experimental design is to generate templates for the resumes to be sent. The challenge is to produce a set of realistic and representative resumes without using resumes that belong to actual job seekers. To achieve this goal, we start with resumes of actual job searchers but alter them sufficiently to create distinct resumes. The alterations maintain the structure and realism of the initial resumes without compromising their owners.

We begin with resumes posted on two job search websites as the basis for our artificial resumes.[11] While the resumes posted on these websites may not be completely representative of the average job seeker, they provide a practical approximation.[12] We restrict ourselves to people seeking employment in our experimental cities (Boston and Chicago). We also restrict ourselves to four occupational categories: sales, administrative support, clerical services and customer services. Finally, we further restrict ourselves to resumes posted more than six months prior to the start of the experiment. We purge the selected resumes of the person's name and contact information.

During this process, we classify the resumes within each occupational category into two groups: high and low quality. In judging resume quality, we use criteria such as labor market experience, career profile, existence of gaps in employment and skills listed. Such a classification is admittedly subjective but it is made independently of any race assignment on the resumes (which occurs later in the experimental design). To further reinforce the quality gap between the two sets of resumes, we add to each high quality resume a subset of the following features: summer or while-at-school employment experience, volunteering experience, extra computer skills, certification degrees, foreign language skills, honors or some military experience. This resume quality manipulation needs to be somewhat subtle to avoid making a higher quality job applicant over-qualified for a given job. We try to avoid this problem by making sure that the features listed above are not all added at once to a given resume. This leaves us with a high quality and a low quality pool of resumes.[13].

To minimize similarity to actual job seekers, we use resumes from Boston job seekers to form templates for the resumes to be sent out in Chicago and used the Chicago resumes to form templates for the resumes to be sent out in Boston. To implement this migration, we alter the names of the schools and previous employers on the resumes. More specifically, for each Boston resume, we use the Chicago resumes to replace a Boston school by a Chicago school.[14] We also use the Chicago resumes to replace a Boston employer by a

[11] The sites are *www.careerbuilder.com* and *www.americasjobbank.com*.
[12] In practice, we found large variation in skill levels among people posting their resumes on these sites.
[13] In Section 4.2 and Table 3, we provide a detailed summary of resume characteristics by quality type.
[14] We try as much as possible to match high schools and colleges on quality and demographic characteristics.

Chicago employer in the same industry. We use a similar procedure to migrate Chicago resumes to Boston.[15] This produces distinct but realistic looking resumes, similar in their education and career profiles to this sub-population of job searchers.[16]

3.2 Identities of Fictitious Applicants

The next step is to generate identities for the fictitious job applicants: names, telephone numbers, postal addresses and (possibly) e-mail addresses. The choice of names is crucial to our experiment.[17] To decide on which names are uniquely African American and which are uniquely White, we use name frequency data calculated from birth certificates of all babies born in Massachusetts between 1974 and 1979. We tabulate these data by race to determine which names are distinctively White and which are distinctively African American. Distinctive names are those that have the highest ratio of frequency in one racial group to frequency in the other racial group.

As a check of distinctiveness, we conducted a survey in various public areas in Chicago. Each respondent was given a resume with a name and asked to assess features of the person, one of which being race. In general, the names led respondents to readily attribute the expected race for the person but there were a few exceptions and these names were disregarded.[18]

The final list of first names used for this study are reported in Appendix Table 1. The table also reports the frequency of these names in the Massachusetts birth certificates data.[19] The African American first names used in the experiment are remarkably common in the population. This suggests that by using these names as an indicator of race, we are actually covering a large segment of the African American population.[20]

Applicants in each race/sex/city/resume quality cell are allocated the same phone number. This guarantees that we can precisely track employer callbacks in each of these cells. The phone lines we use are virtual ones with only a voice mail box attached to it. A similar outgoing message is recorded on each of the voice mail boxes but each message is recorded by someone of the appropriate race and gender. Since we allocate the same phone number for applicants with different names, we cannot use a person name in the outgoing message.

[15] Note that for applicants with schooling or work experience outside of the Boston or Chicago areas, we leave the school or employer name unchanged.

[16] We also generate a set of different fonts, layouts and cover letters to further differentiate the resumes. These are applied at the time the resumes are sent out.

[17] We chose name over other potential manipulations of race, such as affiliation with a minority group, because we felt such affiliations may especially convey more than race.

[18] For example, Maurice and Jerome are distinctively African American names in a frequency sense yet are not perceived as such by many people.

[19] We also tried to use more White sounding last names for White applicants and more African American sounding last names for African American applicants. The last names used for White applicants are: Baker, Kelly, McCarthy, Murphy, Murray, O'Brien, Ryan, Sullivan and Walsh. The last names used for African American applicants are: Jackson, Jones, Robinson, Washington and Williams.

[20] One might however wonder whether this is an atypical segment of the African American population. In Section 5.1, we discuss whether the race effect could be interpreted as a social class effect.

While we do not expect positive feedback from an employer to take place via postal mail, resumes still need postal addresses. We therefore construct fictitious addresses based on real streets in Boston and Chicago using the White Pages. We select up to 3 addresses in each 5-digit zip code in Boston and Chicago. Within cities, we randomly assign addresses across all resumes. We also create 8 email addresses, 4 for Chicago and 4 for Boston.[21] These email addresses are neutral with respect to both race and sex. Not all applicants are given an email address. As we explained above, the email addresses are used almost exclusively for the higher quality resumes. This procedure leaves us with a bank of names, phone numbers, addresses and e-mail addresses which we can assign to the template resumes when responding to the employment ads.

3.3 Responding to Ads

The experiment was carried on between July 2001 and January 2002 in Boston and between July 2001 and May 2002 in Chicago.[22] Over that period, we collected all employment ads in the Sunday editions of *The Boston Globe* and *The Chicago Tribune* in the sales, administrative support, and clerical and customer services sections. We eliminate any ad where applicants are asked to call or appear in person. In fact, most ads we surveyed in these job categories asked for applicants to fax in or (more rarely) mail in their resume. We log the name (when available) and contact information for each employer, along with any information on the position advertised and specific requirements (such as education, experience, or computer skills). We also record whether or not the ad explicitly states that the employer is an equal opportunity employer.

For each ad, we use the bank of resumes to sample four resumes (two high-quality and two low-quality) that fit the job description and requirements as closely as possible.[23] In some cases, we slightly alter the resumes to improve the quality of the match, such as by adding the knowledge of a specific software program.

One of the high and one of the low quality resumes selected are then drawn at random to receive African American names, the other high and low resumes receive White names.[24] We use male and female names for sales jobs, whereas we use nearly exclusively female names for administrative and clerical jobs to increase callback rates.[25] Based on sex, race, city and resume quality, we assign a resume the appropriate phone number. We also select at random a postal address. Finally, e-mail addresses are added to most of the high quality resumes.[26] The final resumes are formatted, with fonts, layout and cover letter style chosen at

[21] The e-mail addresses are registered on *Yahoo.com*, *Angelfire.com* or *Hotmail.com*.

[22] This period spans both tight and slack labor markets. In our data, this is apparent as call-back rates (and number of new ads) dropped precipitously after September 11th, 2001. Interestingly, however, the amount of discrimination we measure in these two periods is the same.

[23] In some instances, our resume bank does not have four resumes that are appropriate for a given ad. In such instances, we send only two resumes.

[24] Though the same names are repeatedly used in our experiment, we guarantee that no given ad receives multiple resumes with the same name.

[25] Male names were used for a few administrative jobs in the first month of the experiment.

[26] In the first month of the experiment, a few high quality resumes were sent without email address and a few low quality resumes were given email addresses. See Table 3 for details.

random. The resumes are then faxed (or in a few cases mailed) to the employer.[27] All in all, we respond to more than 1300 employment ads over the entire sample period and send close to 5000 resumes.

3.4 Measuring Responses

We measure whether a given resume elicits a callback or e-mail back for an interview. For each phone or email response, we use the content of the message left by the employer (name of the applicant, company name, telephone number for contact) to match the response to the corresponding resume-ad pair.[28] Any attempt by employers to contact applicants via postal mail cannot be measured in our experiment since the addresses are fictitious. Several human resource managers confirmed to us that employers rarely, if ever, contact applicants via postal mail to set up interviews.

3.5 Weaknesses of the Experiment

We have already highlighted the strengths of this experiment relative to previous audit studies. We now discuss its weaknesses. First, our outcome measure is crude, even relative to the previous audit studies. Ultimately, one cares about whether an applicants gets the job and about the wage offered conditional on getting the job. Our procedure, however, simply measures callbacks for interviews. To the extent that the search process has even moderate frictions, one would expect that reduced interview rates would translate into reduced job offers. However, we are not able to translate our results into gaps in hiring rates or earnings.

Another weakness is that the resumes do not directly report race but instead suggest race through personal names. This leads to various sources of concern. First, while the names are chosen to make race salient, some employers may simply not notice the names or not recognize their racial content. As a result, our findings may under-estimate the extent of discrimination. Relatedly, because we are not assigning race but only race-specific name, our results are not representative of the average African American (who may not have such a racially distinct name).[29]

Finally, and this is an issue pervasive in both our study and the pair-matching audit studies, newspaper ads represent only one channel for job search. As is well known from the existing literature, social networks are a common means through which people find jobs and one that clearly cannot be studied here. This omission would affect our results if African Americans use social networks more or if less discriminating employers rely more on networks.[30]

[27] As part of the faxing process, we strip all identifiers from the outgoing fax to guarantee that employers could not see that all four faxes originate from the same locale.

[28] Very few employers used email to contact an applicant back.

[29] As Appendix Table 1 indicates, the African American names we use are however quite common among African Americans, making this less of a concern. We return to this issue in Section 5.1.

[30] In fact, there is some evidence that African Americans may rely *less* on social networks for their job search (Holzer, 1987).

4 Results

4.1 Is There Discrimination?

Table 1 tabulates average callback rates by racial soundingness of names. Included in brackets under each rate is the number of resumes sent in that cell. Row 1 presents our results for the full data set. Resumes with White names have a 10.08 percent chance of receiving a callback. Equivalent resumes with African American names have a 6.70 percent chance of being called back. This represents a difference in callback rates of 3.35 percentage points, or 50 percent, that can solely be attributed to the name manipulation. Column 4 shows that this difference is statistically significant.[31] Put in other words, these results imply that a White applicant should expect on average one callback for every 10 ads she or he applies to; on the other hand, an African American applicant would need to apply to 15 different ads to achieve the same result.[32]

How large are these effects? While the cost of sending additional resumes might not be large per se, this 50 percent gap could be quite substantial when compared to the rate of arrival of new job openings. In our own study, the biggest constraining factor was the limited number of new job openings each week. Another way to benchmark the measured return to a White name is to compare it to the returns to other resume characteristics. For example, in Table 5, we will show that, at the average number of years of experience in our sample, an extra year of experience increases the likelihood of a callback by .4 percentage point. Based on this point estimate, the return to a White name is equivalent to about 8 additional years of experience.

Rows 2 and 3 break down the full sample of sent resumes into the Boston and Chicago markets. About 20 percent more resumes were sent in Chicago than in Boston. The average callback rate (across races) is lower in Chicago than in Boston. This might reflect differences in labor market conditions across cities over the experimental period or maybe differences in the ability of the MIT and Chicago teams of research assistants in selecting resumes that were good matches for a given help wanted ad. The percentage difference in callback rates is however strikingly similar across both cities. White applicants are 48 percent more likely than African American applicants to receive a callback in Chicago and 52 percent more likely in Boston. These racial differences are statistically significant in both cities.

Finally, rows 4 to 7 break down the full sample into female and male applicants. Row 4 displays the average results for all female names while rows 5 and 6 break the female sample into administrative (row 5) and sales jobs (row 6); row 7 displays the average results for all male names. As noted earlier, female names were used in both sales and administrative job openings whereas male names were used close to exclusively for sales openings.[33] Looking across occupations, we find a significant racial gap in callbacks for both males

[31] These statistical tests assume independence of callbacks. We have however verified that the results stay significant when we assume that the callbacks are correlated either at the employer or first name level.

[32] This obviously assumes that African American applicants cannot assess a priori which firms are more likely to discriminate against them.

[33] Only about 6 percent of all male resumes were sent in response to an administrative job opening.

(49 percent) and females (50 percent). Comparing males to females in sales occupations, we find a larger racial gap among males (49 percent versus 25 percent). Interestingly, females in sales jobs appear to receive more callbacks than males; however, this (reverse) gender gap is statistically insignificant and economically smaller than any of the racial gaps discussed above.

Rather than studying the distribution of callbacks at the applicant level, one can also tabulate the distribution of callbacks at the employment ad level. In Table 2, we compute the fraction of employers that treat White and African American applicants equally, the fraction of employers that favor White applicants and the fraction of employers that favor African American applicants. Because we send up to four resumes in response to each sampled ad, the 3 categories above can each take 3 different forms. Equal treatment occurs when either no applicant gets called back, one White and one African American get called back or two Whites and two African Americans get called back. Whites are favored when either only one White gets called back, two Whites and no African American get called back or two Whites and one African American get called back. African Americans are favored in the other cases.

As Table 2 indicates, equal treatment occurs for about 87.5 percent of the help-wanted ads. As expected, the major source of equal treatment comes from the high fraction of ads for which no callbacks are recorded (82.5 percent of the ads). Whites are favored by nearly 9 percent of the employers, with a majority of these employers contacting exactly one White applicant. African Americans, on the other hand, are favored by only about 3.7 percent of employers. We formally test whether there is symmetry in the favoring of Whites over African Americans and African Americans over Whites by employers. We find that the difference between the fraction of employers favoring Whites and the fraction of employers favoring African Americans (5.11 percent) is statistically very significant ($p = .0000$).

4.2 Do African Americans Receive Different Returns to Resume Quality?

Our results so far suggest a substantial amount of discrimination in job recruitment. Next, we would like to learn more about how employers discriminate. More specifically, we ask how employers respond to improvements in African American applicants' credentials. To answer this question, we examine how the racial gap in callback rates varies by resume quality.

As we mentioned in section 3, for most of the employment ads we respond to, we send four different resumes: two higher quality and two lower quality ones. Table 3 gives a better sense of which factors enter into this subjective classification. It displays means and standard deviations of the most relevant resume characteristics for the full sample (column 1) as well as broken down by race (columns 2 and 3) and resume quality (columns 4 and 5). Since applicants' names are randomized, there is no difference in resume characteristics by race. Columns 4 and 5 document the objective differences between resumes subjectively classified as high and low. Higher quality applicants have on average an extra year of labor market experience,

fewer employment holes (where an employment hole is defined as a period of at least 6 months without a reported job), are more likely to have worked while at school and to report some military experience. Also, higher quality applicants are more likely to have an email address, to have received some honors and to list some computer skills and other special skills (such as a certification degree or foreign language skills) on their resume. Note that the higher and lower quality resumes do not differ on average with regard to applicants' education level.[34] About 70 percent of the sent resumes report a college degree.[35]

The last 5 rows of Table 3 show summary characteristics of the applicants' zip code address. Using 1990 Census data, we compute the fraction high-school dropouts, fraction college educated or more, fraction Whites, fraction African Americans and log(median per capita income) for each zip code used in the experiment. Since addresses are randomized within cities, these neighborhood quality measures are uncorrelated with race or resume quality.

The differences in callback rates between high and low resumes are presented in Table 4. The first thing to note is that the resume quality manipulation works: higher quality resumes receive higher callback rates. As row 1 indicates, we record a callback rate of more than 11 percent for White applicants with a higher quality resume, compared to 8.8 percent for White applicants with lower quality resumes. This is a statistically significant difference of 2.51 percentage points, or 30 percent. Most strikingly, African Americans experience much less of an increase in callback rate for similar improvements in their credentials. African Americans with higher quality resumes receive a callback 6.99 percent of the time, compared to 6.41 percent for African Americans with lower quality resumes. This is only a .58 percentage point, or 9 percent, increase and this increase is not statistically significant.

Instead of relying on the subjective quality classification, Panel B directly uses resume characteristics to classify the resumes. More specifically, we use a random subsample of one-third of the resumes to estimate a probit regression of the callback dummy on the resume characteristics listed in Table 3. We further control for a sex dummy, a city dummy, 6 occupation dummies and a vector of job requirements as listed in the employment ads.[36] We then use the estimated coefficients on the resume characteristics to rank the remaining two-thirds of the resumes by predicted callback. We classify as "high" resumes those that have above median predicted callback; similarly, we classify as "low" resumes those that have below median predicted callback. As one can see from Panel B, qualitatively similar results emerge from this analysis. While African Americans do appear to significantly benefit from higher quality resumes under this alternative classification, they benefit much less than Whites. The ratio of callback rates for high versus low quality resumes is 1.66 for African Americans, compared to 2.81 for Whites.

[34] This reflects the fact that all sent resumes, whether high or low quality, are chosen to be good matches for a given job opening.

[35] This varies from about 50 percent for the clerical and administrative support positions to more than 80 percent for the executive, managerial and sales representatives positions.

[36] See section 4.4 for more details on these occupation categories and job requirements.

In Table 5, we directly report the results of race-specific probit regressions of the callback dummy on resume characteristics. We however start in column 1 with results for the full sample of sent resumes. As one can see, many of the resume characteristics have the expected effect on the likelihood of a callback. The addition of an email address, honors and special skills all have a positive and significant effect on the likelihood of a callback.[37] Also, more experienced applicants are more likely to get called back: at the average number of years of experience in our sample (8 years), each extra year of experience increases the likelihood of a callback by about .4 percentage point. The most counterintuitive effects come from computer skills, which appear to negatively predict callback, and employment holes, which appear to positively predict callback.

The same qualitative patterns hold in column 2 where we focus on White applicants. More importantly, the estimated returns to an email address, additional work experience, honors and special skills appear economically stronger for that racial group. For example, at the average number of years of experience in our sample, each extra year of experience increases the likelihood of a call back by about .7 percentage point. Also, working while at school, while not a statistically significant determinant of callback in the full sample, yields positive returns for White applicants.

As might have been expected from the two previous columns, we find that the estimated returns on these resume characteristics are all economically and statistically weaker for African American applicants (column 3). In fact, all the estimated effects for African Americans are statistically insignificant, except for the return to special skills. Resume characteristics thus appear less predictive of callback rates for African Americans than they are for Whites. To see this more clearly, we predict callback rates using either the regression in column 2 or the regression in column 3. The standard deviation of the predicted callback from column 2 is .064, whereas it is only 0.034 from column 3. In summary, employers simply seem to pay less attention or discount more the characteristics listed on the resumes with African American sounding names. Taken at face value, these results suggest that African Americans may face relatively lower individual incentives to invest in higher skills.[38]

4.3 Applicants' Address

An incidental feature of our experimental design is the random assignment of address to the resumes. This allows us to examine whether and how an applicant's residential address, all else equal, affects the likelihood of a callback. In addition, and most importantly for our purpose, we can also ask whether African American applicants are helped more by residing in more affluent neighborhoods.

[37] Note that the e-mail address dummy, because it is close to perfectly correlated with the subjective resume quality variable, may in part capture some other unmeasured resume characteristics that may have led us to categorize a given resume as higher quality.

[38] This of course assumes that the changes in wage offers associated with higher skills are the same across races, or at least not systematically larger for African Americans.

We perform this analysis in Table 6. We start (columns 1, 3 and 5) by discussing the effect of neighborhood of residence across all applicants. Each of these columns reports the results of a probit regression of the callback dummy on a specific zip code characteristic and a city dummy. Standard errors are corrected for clustering of the observations at the employment ad level. We find a positive and significant effect of neighborhood quality on the likelihood of a callback. Applicants living in Whiter (column 1), more educated (column 3) or higher income (column 5) neighborhoods have a higher probability of receiving a call back. For example, a 10 percentage point increase in the fraction of college-educated in zip code of residence increases the likelihood of a callback by .6 percentage point.

In columns 2, 4 and 6, we further interact each of the zip code characteristic above with a dummy variable for whether the applicant is African American or not. Each of the probit regressions in these columns also includes an African American dummy, a city dummy and an interaction of the city dummy with the African American dummy. There is no evidence that African Americans benefit any more than Whites from living in a more White, more educated zip code. The estimated interactions between fraction White and fraction college educated with the African American dummy are economically very small and statistically insignificant. We do find an economically more meaningful effect of zip code median income level on the racial gap in callback; that effect, however, is statistically insignificant.

In summary, while the quality of the neighborhood of residence is a significant factor in employers' decision to contact a job applicant, African Americans do not appear to benefit more than Whites from living in better neighborhoods. These findings are interesting. Indeed, if ghettos and bad neighborhoods are particularly stigmatizing for African Americans, one might have expected African Americans to be helped more by having a "good" address. Our results do not lend support to this hypothesis.

4.4 Job Characteristics

In this section, we turn to job characteristics and ask whether discrimination varies based on the specific job requirements or occupation listed in the employment ads.

Column 1 of Table 7 gives a description of the specific requirements stated in the sample of ads we respond to. About 80 percent of the ads state some form of requirement. About 43 percent of the ads require some minimum experience, of which roughly 50 percent simply ask for "some experience," 25 percent less than two years and 25 percent at least 3 years of experience. About 44 percent of ads mention some computer knowledge requirement, which can range from Excel or Word to more esoteric software programs. Good communication skills are explicitly required in about 12 percent of the ads. Organization skills are mentioned 7 percent of the time. Finally, only about 11 percent of the ads list an explicit education requirement. Of these, 8.8 percent require a high school degree, 49 percent some college (such as an associate degree) and

the rest at least a 4-year college degree.[39]

How do these different job requirements affect the racial gap in callback? To answer this question, we show in column 2 the results of various probit regressions. Each entry in this column is the marginal effect of the requirement on the racial gap in callback. Specifically, each entry comes from a separate probit regression where we regress a callback dummy on an African American dummy, the requirement dummy and the interaction of the requirement dummy with the African American dummy. The reported coefficient is that on the interaction term. The point estimates show no consistent economic pattern and are all statistically insignificant. Measures of job quality, such as experience, computer skills or education requirements do not predict the extent of discrimination. Equally surprising, communication or other inter-personal skill requirements have no effect on the level of discrimination. In short, discrimination appears to vary little by job requirements.[40]

Panel A of Table 8 investigates whether discrimination varies significantly across occupations. As we mention earlier, the specific subsections of the Sunday newspapers help-wanted sections that we sample for this study broadly relate to sales and administrative positions. This is reflected in the occupational distribution reported in column 1. To form this classification, we use the job description listed in the employment ad and map it into one of six following categories: executives and managerial occupations; administrative supervisors; sales representatives; sales workers; secretaries and legal assistants; clerical occupations.[41] In the 1990 5 percent Census, these occupations account for about 46.5 percent of total salaried private sector employment in Chicago and Boston. African Americans account for about 4 percent of employment in these occupations in Boston and 10.5 percent in Chicago. For comparison, African Americans account for about 4.8 percent of total salaried private sector employment in Boston and 11.8 percent in Chicago.

We report in the second to last column of Table 8 average earnings measures for each of these occupation categories. These measures are computed from the 1990 5 percent Census for Boston and Chicago. More specifically, we first estimate a micro wage regression of log(annual earnings) on 8 education dummies, a quadratic in age, a sex dummy and a city dummy. We then compute the mean residual earnings in each occupation category. As expected, the sampled occupation categories are very different. The first two (about 22% of the data) correspond to better jobs on average that involve, among other things, the supervision of other workers; these two categories are associated with the highest earnings. Sales representatives and secretaries (about 50% of the data) involve less supervision but are still fairly well paying jobs. The last two occupations (sales workers and clerical workers) correspond to the lowest paying jobs in our sample.

[39] Other requirements sometimes mentioned include typing skills for secretaries (with specific words per minute minimum thresholds) and, more rarely, foreign language skills.

[40] Other ways of estimating these effects produce a similar non-result. Among other things, we considered including a city dummy or estimating the effects separately by city; we also estimated one single probit regression including all requirements at once.

[41] These categories respectively correspond to the following Census 1990 occupation codes: 000-042; 303-307; 253-262; 263-302; 234 and 313; 308-402 (excluding 313).

The second and third columns of Panel A of Table 8 respectively report the callback rates for White and African American applicants in each of these occupation categories. Column 4 reports the ratio of callback rates while column 5 reports the difference. The average callback rate across races seems to decrease with the level of job quality. Whites receive more callbacks than African Americans in all job categories. While the racial gap in callback rates varies somewhat across occupations, we cannot reject the null hypothesis that the gap is the same across all categories.

We however briefly discuss the point estimates by occupation and how they relate to Census measures of average earnings. The smallest gap appears to be for executive job positions, where Whites only face a 33% higher chance than African American of being called back. The second lowest gap is for clerical jobs, with Whites being called back 38% more often. Interestingly, these two job categories are at opposite ends of the job quality spectrum, as captured by our earnings measure. This suggests no monotonous relationship between job quality and the extent of discrimination. The highest discrimination ratio happens for administrative supervisors, the second highest job quality category in our data. For such jobs, Whites are 64% more likely to get a callback.

In the last column of Table 8, we report the race differences in earnings by occupation. These are again computed from the 1990 5 percent Census for Boston and Chicago. The racial gaps in earning are defined as the White-African American differences in mean residual log(annual earnings) by occupation, where residual earnings are estimated as explained above. Interestingly, there does not seem to be a monotonous association between these estimated racial gaps in earnings and the racial gaps in callbacks. For example, executive and managerial positions have the lowest racial gap in callbacks but the second highest racial gap in earnings; similarly, administrative supervisors have the highest racial gap in callbacks but the second lowest racial gap in earnings. On the other hand, sales representatives are associated with the second-highest racial gap in callbacks and the highest racial gap in earnings. These results are interesting and potentially suggest that readily available measures of racial differences, such as the racial gap in earnings used above, may not give a reliable depiction of discrimination patterns in the economy. However, because the occupational differences in callback are not significant, these results should not be overstated.

To summarize, we find that the level of discrimination is remarkably uniform across a variety of occupations and job requirements. This occurs despite the fact that some of the jobs we study are higher earnings, such as managerial positions or high end sales jobs.[42]

4.5 Employer Characteristics

The final set of possible determinants of discrimination we consider are employer characteristics. Collecting such employer information is not obvious as most of this information is not readily available from the

[42] Obviously, we cannot rule out that different patterns might exist when one moves even higher in the job quality distribution.

employment ads we respond to. In fact, the only piece of employer information we can directly collect from the employment ad is whether or not the employer explicitly states being an "Equal Opportunity Employer." In several cases, the name of the employer is not even mentioned in the ad and the only piece of information we can rely on is the fax number which applications must be submitted to.

We proceeded as follows. For employment ads that do not list a specific employer, we use the fax number to identify the company name via web reverse-lookup services. Based on the company names, we use three different data sources (*Onesource Business Browser*, *Thomas Register* and *Dun and Bradstreet Million Dollar Directory, 2001*) to track company information such as total employment, industry and ownership status. Using this same set of data sources, we also try to identify the specific zip code of the company (or company branch) that resumes are to be sent to. Finally, we use the Federal Procurement and Data Center website to find a list of companies that have federal contracts.[43] The racial difference in callback rates for the subsample where employer characteristics could be determined is very similar in magnitude to that in the full sample.

Column 1 of Table 9 reports summary employer characteristics for the available data. Sample sizes for each variable are reported in parentheses. Twenty-nine percent of all employers explicitly state that they are "Equal Opportunity Employers". Eleven percent are federal contractors and, therefore, might face greater scrutiny under affirmative action laws. The average company size is around 2000 employees but there is a lot of variation across firms. Finally, 73 percent of the firms are privately held, 15 percent are publicly traded and 12 percent are non-profits.

The second column of Table 9 presents the marginal effect of each of these characteristics on discrimination. As before, each entry corresponds to a separate probit regression where we regress a callback dummy on an African American dummy, the employer characteristic in that row and the interaction of the employer characteristic with the African American dummy. The reported coefficient is that on the interaction term. First, neither the "Equal Opportunity Employers" nor the federal contractors appear to discriminate less. In fact, each of these employer categories is associated with more discrimination, even though these effects are noisily estimated. Second, we find no effect of employer size on the degree of discrimination.[44] Point estimates indicate that publicly traded firms seem to discriminate more, while not-for-profit organizations seem to discriminate less; however, these effects are again noisily estimated.[45]

Panel B of Table 8 documents how the racial gap in callbacks varies across broad industry categories. Our experiment covers a variety of industries. Around 8 percent of the jobs are in manufacturing, 3 percent in transportation and communication, 22 percent are in wholesale and retail trade, 8 percent are in finance insurance and real estate, 27 percent are in business and personal services and 15 percent are in health,

[43] This website (www.fpdc.gov) is accurate up to and including March 21, 2000.

[44] Similar results hold when we measure employer size using a total sales measure rather than an employment measure.

[45] Of course, our measurement of discrimination by firm type may not be a good indicator of the fraction of African Americans actually employed in these firms. For example, "Equal Opportunity Employers" may receive a higher fraction of African American resumes. Their actual hiring may therefore look different from that of non "Equal Opportunity Employers" when one considers the full set of resumes they receive.

educational and social services. For 16 percent of the ads, we are not able to determine the recruiter's industry.

Columns 2 through 4 show the callback rates for Whites and African Americans, and the differences and ratios of these two. In every industry except for transportation and communication (which corresponds to a very small subsample of the jobs in our experiment), African Americans fare worse than Whites. The biggest gap appears to be in finance, insurance and real estate (where the ratio of callback is 2.44), while the smallest (outside of transportation) appears to be in health, educational and social services (where the ratio is 1.35). While the industrial differences in callbacks appear more pronounced than the occupational differences we discussed above, we cannot reject the null hypothesis that discrimination is the same across industries.

We however report how the point estimates by industry relate to Census measures of earnings and racial gap in earnings by industry. As we did for occupations, we compute such measures using the 1990 5 percent Census for Boston and Chicago. We follow a similar methodology and compute residual log(annual earnings) and mean residual log(annual earnings) by industry and industry-race cell. These Census measures are presented in the last two columns of the table.

Documenting the relationship between inter-industry wage differentials and discrimination is interesting in light of theories that predict that higher industry rents may allow for more discrimination. If higher rents translate into higher wages through rent-sharing, these theories predict that high wage industries should discriminate more (Katz And Summers 1989). Such a prediction could evolve from a straightforward taste-based model of discrimination where employers are homogeneous in their dislike of minority workers but are differentially constrained by industry-specific product market pressures in their ability to engage in costly discrimination (Becker 1961). If one abstracts from the small transportation and communication sector, it does appear that the two highest wage industries (manufacturing and finance, insurance and real estate) are also the two highest discrimination industries. However, the relationship is far from monotonous. For example, health, educational and social services have about average wage levels but also record the lowest level of discrimination.

Our final table, Table 10, focuses on the marginal effect of employer location on discrimination.[46] As we mentioned above, we use as a measure of employer location the zip code of the company (or company branch) resumes were to be sent to. More specifically, we ask whether discrimination varies with the fraction of African Americans in the employer's zip code. Each column in Table 10 corresponds to a different probit regression of a callback dummy on an African American dummy, the fraction African Americans in the zip code and the interaction of the African American dummy with fraction African Americans in the zip code. Reported in the last row of this table is the mean fraction African Americans in the relevant sample. In

[46] For previous work on the effect of employer location on labor market discrimination, see for example Raphael, Stoll and Holzer (2000).

columns 1 and 2, we use the full sample for which data are available. As we can see in column 1, the interaction term is positive, suggesting that employers located in more African American zip codes are more likely to call back African American applicants. However, the coefficient on this interaction term is not statistically significant. Column 2 estimates the same probit regression but allows for the degree of discrimination to vary by occupation categories, industry categories and city. This increases a bit the magnitude of the coefficient on the interaction term but that coefficient is still not significant.

Pooling Chicago and Boston in the regressions in columns 1 and 2 however hides an important difference across the two cities. The last four columns of Table 10 replicate the exercise above separately for Chicago (columns 3 and 4) and Boston (columns 5 and 6). In Chicago, there is a larger and statistically significant effect of employer location on discrimination. A 10 percentage point increases in the fraction African Americans in the employer's zip code reduces the racial gap in callbacks by close to 1 percentage point. Such a 10 percentage point increases corresponds to a little less than a move from the 25th percentile to the 75th percentile of the fraction African Americans in the Chicago sample. According to our estimates, equality of callbacks across races would occur in a zip code that is about 60 percent African Americans. This corresponds to the 97th percentile in the Chicago sample.

In contrast, we find no effect in Boston. The estimated coefficient on the interaction term is both economically and statistically insignificant. One possible reason for this differential effect might be differences in both the level of and variation in the fraction African Americans across these two cities. The Chicago sample contains many more employers located in fairly African American neighborhoods. The average Chicago employer is located in a zip code with 10.5 percent African Americans; the average Boston employer in a zip code with 4.5 percent African Americans. The 25th and 75th percentiles of fraction African Americans in Chicago are respectively 1 percent and 12.5 percent; these are 0 percent and 5 percent in Boston.

Finally, in regressions not reported here, we also investigated whether other employers' zip code characteristics affect the level of discrimination in Chicago. We found qualitatively similar results to the ones reported above when we use fraction Whites, fraction college educated, fraction high school dropouts or median per capita income in the employer's zip code. These results all suggest that employers located in less affluent neighborhoods are somewhat more amenable to hiring African American applicants.

5 Interpretation

Two main sets of issues arise when interpreting the results above. First, does our design isolate the effect of race or is the name manipulation conveying some other factors than race? Second, what do our results imply for different models of discrimination?

5.1 Potential Confounds

Though we have interpreted our results in terms of racial differences, we actually manipulate only the name on the resume. While these names clearly signal race, perhaps they also signal some other personal characteristics. More specifically, one might be concerned that employers are inferring social background from the personal name. When employers read a name like "Tyrone" or "Latoya," they may assume that the person comes from a disadvantaged background. In the extreme form of this social background interpretation, employers do not care at all about race but are discriminating only against the social background conveyed by the names we have chosen.[47]

While plausible, we feel that some of our earlier results are hard to reconcile with this interpretation. For example, in Table 6, we found that while employers value "better" addresses, African Americans are not helped more than Whites by living in Whiter or more educated neighborhoods. If the African American names mainly signal negative social background, one might have expected the estimated name-gap to be lower for the better addresses. Also, if the names mainly signal social background, one might have expected the name gap to be higher for jobs that rely more on soft skills or require more inter-personal interactions. We found no such evidence in Tables 6 or 7.

We however directly address this alternative interpretation by examining the average social background of babies born with the names used in the experiment. We were able to obtain birth certificate data on mother's education (less than high school, high school or more) for babies born in Massachusetts between 1970 and 1986.[48] For each first name in our experiment, we compute the fraction of babies with that name and in that gender-race cell whose mothers have at least completed a high-school degree.

In Table 11, we display the average callback rate for each first name along with this proxy for social background. Within each race-gender group, the names are ranked by increasing callback rate. Interestingly, there is significant variation in callback rates by name. Of course, chance alone could produce such variation because of the rather small number of observations in each cell (about 200 for the female names and 70 for the male names).[49]

[47] African Americans as a whole come from more disadvantaged backgrounds than Whites. For this social class effect to be something of independent interest, one must assert that African Americans with the African American names we have selected are from a lower social background than the average African American and/or that Whites with the White names we have selected are from a higher social background than the average White. We come back to this point below.

[48] This longer time span (compared to that used to assess name frequencies) was imposed on us for confidentiality reasons. When fewer than 10 births with education data available are recorded in a particular education-name cell, the exact number of births in that cell is not reported and we impute 5 births. Our results are not sensitive to this imputation. One African-American female name (Latonya) and two male names (Rasheed and Hakim) were imputed in this way. For one African American male name (Tremayne) we had too few births with education data at all and it is dropped from this analysis. Our results are qualitatively similar when we use a larger data set of California births for the years 1989 to 2000 (kindly provided to us by Steven Levitt).

[49] We formally tested whether this variation was significant by estimating a probit regression of the callback dummy on all the personal first names, allowing for clustering of the observations at the employment ad level. For all but African American females, we cannot reject the null hypothesis that all the first name effects in the same race-gender group are the same. Of course, a lack of a rejection does not mean there is no underlying pattern in the between-name variation in callbacks that might have been detectable with larger sample sizes.

The row labeled "Average" contains the average fraction of mothers that have at least completed high school for the set of names listed in that gender-race group. The row labeled "Overall" contains the average fraction of mothers that have at least completed high school for the full sample of births in that gender-race group. For example, 83.9 percent of White female babies born between 1970 and 1986 have mothers with at least a high school degree; 91.7 percent of the White female babies with one of the names used in the experiment have mothers with at least a high school degree.

Consistent with a social background interpretation, the African American names we have chosen fall below the African American average. For African American male names, however, the gap between the experimental names and the population average is negligible. For White names, both the male and female names are above the population average.

But, more interestingly for us, there is substantial between-name heterogeneity in social background. African American babies named Kenya or Jamal are affiliated with much higher mothers' education than African American babies named Latonya or Leroy. Conversely, White babies named Carrie or Neil have lower social background than those named Emily or Geoffrey. This allows for a direct test of the social background hypothesis within our sample: are names associated with a worse social background discriminated against more? In the last row in each gender-race group, we report the rank-order correlation between callback rates and mother's education. The social background hypothesis predicts a positive correlation. Yet, for all four categories, we find the exact opposite. The p-values indicate that we cannot reject independence at standard significance levels except in the case of African American males where we can almost reject it at the 10 percent level. In summary, this test suggests little evidence that social background drives the extent of discrimination.

Names might also influence our results through familiarity. It might be that these African American names simply appear odd to human resource managers and that any odd name faces discrimination. But as noted earlier, the names we have selected are not particularly uncommon among African Americans (see Appendix Table 1). We have also performed a similar exercise to that of Table 11 and measured the rank-order correlation between name-specific callback rates and name frequency for each gender-race group. We found no systematic positive correlation.

There is one final potential confound to our results. Perhaps what appears as discrimination is actually the result of *reverse discrimination.* If qualified African Americans are thought to be in high demand, then employers with average quality jobs might feel that an equally talented African American would never accept an offer from them and thereby never call her or him in for an interview. Such an argument might also explain why African Americans do not receive as strong of a return as Whites to better resumes, since higher qualification only strengthens this argument. But this interpretation would suggest that among the better jobs, we ought to see evidence of reverse, or at least less, discrimination. However, as we discussed in

Section 4.4, we do not find any such evidence. Discrimination does not vary across jobs with different skill requirements, nor does it vary across occupation categories. Even among the better jobs in our sample, we find quite a bit of discrimination against African American names.[50]

5.2 Relation to Existing Theories

What do the results in this paper imply for existing models of discrimination? Existing economic theories can be classified into two main categories: taste-based and statistical discrimination models. Both sets of models can obviously "explain" our average racial gap in callbacks by virtue of being discrimination models. But can these models explain our other findings? More specifically, we discuss the relevance of these models with a focus on two of these findings: (i) the lower returns to credentials for African Americans, and (ii) the relative uniformity of discrimination across occupations and job requirements and, to a lesser extent, industries.

Taste-based models differ in whose prejudiced "tastes" they emphasize: customers, co-workers or employers. Customer and co-worker discrimination models seem at odds with the lack of significant variation of the racial gap by occupation and industry categories, as the amount of customer contact and the fraction of White employees must vary across these categories. More precisely, we do not find more discrimination among jobs that explicitly require "communication skills" and jobs for which we expect either customer or co-worker contacts to be higher (see Table 7).

Because we do not know what drives employer tastes, employer discrimination models could be consistent with the lack of occupation and industry variation. Employer discrimination also matches the finding that Chicago employers located in more African American neighborhoods discriminate less. However, employer discrimination models would struggle to explain why African Americans get relatively lower returns to their credentials. Indeed, the cost of indulging the discrimination taste should increase as the minority applicants' credentials increase.[51]

Statistical discrimination models are the prominent alternative to the taste-based models in the economics literature. In one class of statistical discrimination models, employers use the observable race to proxy for *unobservable* skills (e.g. Phelps 1972, Arrow 1973). This class of models struggle to explain the credentials effect. Indeed, the added credentials should lead to a larger update for African Americans and hence greater returns to skills for that group.

A second class of statistical discrimination models "emphasize the precision of the information that

[50] One might argue that employers who reverse discriminate hire through less formal channels than help-wanted ads. But this would imply that African Americans are less likely to find jobs through formal channels, which does not appear consistent with the existing evidence (Holzer, 1987).

[51] One could however assume that employer tastes differ not just by race but also by race and skill, so that employers have greater prejudice against minority workers with better credentials. But the opposite preferences, employers having a particular distaste for low-skilled African Americans, also seem reasonable.

employers have about individual productivity" (Altonji and Blank, 1999). Specifically, in these models, employers believe that the same observable signal is more precise for Whites than for African Americans (Aigner and Cain 1977, Lundberg and Startz 1983, Cornell and Welch 1996). Under these models, African Americans should receive lower returns to observable skills because employers place less weight on these skills. However, how reasonable is this interpretation for our experiment? First, it is important to note that we are using the same set of resume characteristics for both racial groups. So the lower precision of information for African Americans cannot be that, for example, an employer does not know what a high school degree from a very African American neighborhood means (as in Aigner and Cain (1977)). Second, many of the credentials on the resumes are in fact externally and easily verifiable, such as a certification for a specific software.

An alternative version of these models would rely on bias in the observable signal rather than differential variance or noise of these signals by race. Perhaps the skills of African Americans are discounted because affirmative action makes it easier for African Americans to get these skills. While this is plausible for credentials such as an employee of the month honor, it is less clear why this would apply to more verifiable and harder skills. It is equally unclear why work experience would be less rewarded since our study suggests that getting a job is prone to discrimination rather than reverse discrimination.

The uniformity of discrimination across occupation is also troubling for a statistical discrimination interpretation. Numerous factors that should affect the importance of statistical discrimination, such as the importance of unobservable skills, the observability of qualifications, the precision of observable skills and the ease of performance measurement, may vary quite a lot across occupations.

These facts suggest that perhaps other models may do a better job at explaining our findings. One simple alternative model is lexicographic search by employers. Employers receive so many resumes that they may use quick heuristics in reading these resumes. One such heuristic could be to simply read no further when they see an African American name. Thus they may never see the skills of African American candidates and this could explain why these skills are not rewarded. This might also to some extent explain the uniformity of discrimination since the screening process (i.e. looking through a large set of resumes) may be quite similar across the variety of jobs we examine.[52]

6 Conclusion

This paper suggests that discrimination is an important factor in why African Americans do poorly in the labor market. Job applicants with African American names get far fewer callbacks for each resume they

[52] Another explanation could be based on employer stereotyping or categorizing. If employers have coarser stereotypes for African Americans, many of our results would follow. See Jones (2002) for the relevant psychology and Mullainathan (2002) for a formalization of categorization.

send out. Equally importantly, applicants with African American names find it hard to fight discrimination in callbacks by improving their observable skills or credentials.

Taken at face value, our results on differential returns to skill have possibly important policy implications. They suggest that training programs alone may not be enough to alleviate the barriers raised by discrimination. For training to work, some general equilibrium force outside the context of our experiment would have to be at play. So, while a massive training program at the national level may change the structure of discrimination, small training programs may not work. In fact, if African Americans recognize how employers reward their skills, they may be rationally more reluctant than Whites to even participate in these programs.

References

Aigner, Dennis J. and Glenn G. Cain, "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review*, 1977, 30 (1): 175-187.

Altonji, Joseph G. and Rebecca M. Blank, "Race and Gender in the Labor Market," in Orley Ashenfelter and David Card eds, *Handbook of Labor Economics*, vol. 3, Elsevier Science B.V., 1999.

Arrow, Kenneth J., "What Has Economics to Say About Racial Discrimination?" *The Journal of Economic Perspectives*, 1998, 12 (2): 91-100.

Arrow, Kenneth J., "The Theory of Discrimination," in Orley Ashenfelter and Albert Rees, eds., *Discrimination in Labor Markets*, Princeton, NJ: Princeton University Press, 1973.

Becker, Gary S., *The Economics of Discrimination*, 2nd Edition, University of Chicago Press: Chicago, IL (1961).

Brown, C. and P. Gay, *Racial Discrimination 17 Years After the Act*, London, UK: Policy Studies Institute, 1985.

Council of Economic Advisers, *Changing America: Indicators of Social and Economic Well-Being by Race and Hispanic Origin*, September 1998. *http://w3.access.gpo.gov/eop/ca/pdfs/ca.pdf.*

Cornell Bradford and Ivo Welch, "Culture, Information and Screening Discrimination," *The Journal of Political Economy*, 1996, 104 (3): 542-571.

Cross, Harry, Genevieve Kenney, Jane Mell and Wendy Zimmermann, *Employer Hiring Practices: Differential Treatment of Hispanic and Anglo Job Applicants*, Washington, DC: Urban Institute Press, 1990.

Darity, William A. Jr. and Patrick L. Mason, "Evidence on Discrimination in Employment: Codes of Color, Codes of Gender," *The Journal of Economic Perspectives*, 1998, 12 (2): 63-90.

Fix, Michael and Margery Austin Turner (eds.), *A National Report Card on Discrimination in America: The Role of Testing*, 1998.

Goldin, Claudia and Cecilia Rouse, "Orchestrating Impartiality: The Impact of Blind Auditions on Female Musicians," *The American Economic Review*, 2000, 90 (4): 715-741.

Heckman, James J., "Detecting Discrimination," *The Journal of Economic Perspectives,* 1998, 12 (2): 101-116.

Heckman, James J. and Peter Siegelman, "The Urban Institute Audit Studies: Their Methods and Findings," in Michael Fix and Raymond J. Struyk, *Clear and Convincing Evidence: Measurement of Discrimination in America*, Lanham, MD: Urban Institute Press, 1992.

Heckman, James J., Lance J. Lochner and Petra E. Todd, "Fifty Years of Mincer Earnings Regressions," Mimeo, University of Chicago, Chicago, 2001.

Holzer, Harry J., "Informal Job Search and Black Youth Unemployment," *American Economic Review,* 1987, 77 (3): 446-452.

Hubbock, J. and S. Carter, *Half a Chance? A Report on Job Discrimination against Young Blacks in Nottingham*, London UK: Commission for Racial Equality, 1980.

James F. and S. W. DelCastillo, "Measuring Job Discrimination by Private Employers Against Young Black and Hispanic Seeking Entry Level Work in the Denver Metropolitan Area," Mimeo, University of Colorado at Denver, Denver, 1991.

Jones, Melinda, *Social Psychology of Prejudice*, Saddle River NJ: Pearson Education, 2002.

Jowell, R. and Prescott-Clark P, "Racial Discrimination and White-Collar Workers in Britain," *Race*, 1970, 11: 397-417.

Katz, Lawrence F. and Lawrence H. Summers, "Industry Rents: Evidence and Implications," *Brookings Papers on Economic Activity*, 1989: 209-275.

Lundberg, Shelly J. and Richard Startz, "Private Discrimination and Social Intervention in Competitive Labor Market," *The American Economic Review*, 1983, 73 (3): 340-347.

McIntyre, Shelby J., Dennis J. Moberg and Barry Z. Posner, "Discrimination in Recruitment: An Empirical Analysis: Comment,"*Industrial and Labor Relations Review*, 1980, 33 (4): 543-547.

Mullainathan, Sendhil, "Thinking Through Categories," Mimeo, Massachusetts Institute of Technology, Cambridge, 2003.

Neumark, David, "Sex Discrimination in Restaurant Hiring: An Audit Study," *The Quarterly Journal of*

Economics, 1996, 111 (3): 915-942.

Newman, Jerry M., "Discrimination in Recruitment: An Empirical Analysis," *Industrial and Labor Relations Review*, 1980, 32 (1): 15-23.

Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *The American Economic Review*, 1972, 62: 659-661.

Raphael, Steven, Michael A. Stoll and Harry J. Holzer, "Are Suburban Firms More Likely to Discriminate against African Americans?," *Journal of Urban Economics*, 2000, 48 (3): 485-508.

Turner, Margery A., Michael Fix and Raymond J. Struyk, *Opportunities Denied, Opportunities Diminished: Racial Discrimination in Hiring*, Washington, DC: Urban Institute Press, 1991.

Turner, Margery A., Raymond J. Struyk and J. Yinger, *Housing Discrimination Study Synthesis*, Washington DC: Urban Institute Press, 1991.

Table 1
Mean Call-Back Rates By Racial Soundingness of Names [a]

	Call-Back Rate for White Names	*Call-Back Rate for African American Names*	*Ratio*	*Difference (p-value)*
Sample:				
All sent resumes	**10.06%** [2445]	**6.70%** [2445]	**1.50**	**3.35%** (.0000)
Chicago	**8.61%** [1359]	**5.81%** [1359]	**1.48**	**2.80%** (.0024)
Boston	**11.88%** [1086]	**7.83%** [1086]	**1.52**	**4.05%** (.0008)
Females	**10.33%** [1868]	**6.87%** [1893]	**1.50**	**3.46%** (.0001)
Females in administrative jobs	**10.93%** [1363]	**6.81%** [1364]	**1.60**	**4.12%** (.0001)
Females in sales jobs	**8.71%** [505]	**6.99%** [529]	**1.25**	**1.72%** (.1520)
Males	**9.19%** [577]	**6.16%** [552]	**1.49**	**3.03%** (.0283)

[a] Notes:

1. The table reports, for the entire sample and different subsamples of sent resumes, the call-back rates for applicants with a White sounding name (column 1) and an African American sounding name (column 2), as well as the ratio (column 3) and difference (column 4) of these call-back rates. In brackets in each cell is the number of resumes sent in that cell.
2. Column 4 also reports the p-value for a test of proportion testing the null hypothesis that the call-back rates are equal across racial groups.

Table 2
Distribution of Call-Backs By Employment Ad [a]

Equal Treatment:	*No Call-back*	*1W+1B*	*2W+2B*
87.37%	82.56%	3.46%	1.35%
[1162]	[1098]	[46]	[18]
Whites Favored (WF):	*1W+0B*	*2W+0B*	*2W+1B*
8.87%	5.93%	1.50%	1.43%
[118]	[79]	[20]	[19]
African Americans Favored (BF):	*1B+0W*	*2B+0W*	*2B+1W*
3.76%	2.78%	0.45%	0.53%
[50]	[37]	[6]	[7]
Ho: WF=BF			
p=.0000			

[a] Notes:

1. This table documents the distribution of call-backs at the employment ad level. "No Call-Back" is the fraction of ads for which none of the fictitious applicants received a call-back. "1W+1B" is the fraction of ads for which exactly one White and one African American applicant received a call-back. "2W+2B" is the fraction of ads for which exactly two White applicants and two African American applicants received a call-back. "Equal Treatment" is defined as the sum of "No Call-Back," "1W+1B," "2W+2B." "1W+0B" is the fraction of ads for which exactly one White applicant and no African American applicant received a call back. "2W+0B" is the fraction of ads for which exactly two White applicants and no African American applicant received a call-back. "2W+1B" is the fraction of ads for which exactly two White applicants and one African American applicant received a call-back. "Whites Favored" is defined as the sum of "1W+0B," "2W+0B," and "2W+1B." "1B+0W" is the fraction of ads for which exactly one African American applicant and no White applicant received a call-back. "2B+0W" is the fraction of ads for which exactly two African American applicants and no White applicant received a call-back. "2B+1W" is the fraction of ads for which exactly two African American applicants and one White applicant received a call-back. "African Americans Favored" is defined as the sum of "1B+0W," "2B+0W," and "2B+1W."
2. In brackets in each cell is the number of employment ads in that cell.

Table 3
Resume Characteristics: Summary Statistics [a]

Sample:	**All Resumes**	**White Names**	**African American**	**Higher Quality**	**Lower Quality**
Characteristic:					
College degree	.72	.72	.72	.72	.72
(Y=1)	(.45)	(.45)	(.45)	(.45)	(.45)
Years of experience	7.82	7.84	7.81	8.27	7.38
	(5.04)	(5.07)	(5.00)	(5.28)	(4.75)
Volunteering experience?	.42	.41	.41	.79	.03
(Y=1)	(.49)	(.49)	(.49)	(.41)	(.16)
Military experience?	.10	.09	.10	.18	.00
(Y=1)	(.30)	(.29)	(.30)	(.39)	(.06)
Email address?	.48	.48	.48	.92	.03
(Y=1)	(.50)	(.50)	(.50)	(.27)	(.17)
Employment holes?	.45	.45	.45	.34	.56
(Y=1)	(.50)	(.50)	(.50)	(.47)	(.50)
Work in school?	.56	.56	.56	.72	.40
(Y=1)	(.50)	(.50)	(.50)	(.45)	(.49)
Honors?	.05	.05	.05	.07	.03
(Y=1)	(.22)	(.23)	(.22)	(.25)	(.18)
Computer skills?	.82	.81	.83	.91	.73
(Y=1)	(.38)	(.39)	(.37)	(.29)	(.44)
Special skills?	.33	.33	.33	.36	.30
(Y=1)	(.47)	(.47)	(.47)	(.48)	(.46)
Fraction high school dropouts	.19	.19	.19	.19	.18
in applicant's zip code	(.08)	(.08)	(.08)	(.08)	(.08)
Fraction college or more	.21	.21	.21	.21	.21
in applicant's zip code	(.17)	(.17)	(.17)	(.17)	(.17)
Fraction Whites	.54	.54	.54	.54	.55
in applicant's zip code	(.33)	(.33)	(.33)	(.33)	(.33)
Fraction African Americans	.31	.31	.31	.32	.31
in applicant's zip code	(.33)	(.33)	(.33)	(.33)	(.33)
Log(median per capita income)	9.55	9.55	9.55	9.54	9.56
in applicant's zip code	(.56)	(.56)	(.55)	(.54)	(.57)
Sample Size	4890	2445	2445	2458	2432

[a] Notes:

1. The table reports means and standard deviations for the resume characteristics. Column (1) refers to all resumes sent; column (2) refers to resumes with White sounding names; column (3) refers to resumes with African American sounding names; column (4) refers to higher quality resumes; column (5) refers to lower quality resumes. See text for details.

Table 4
Average Call-Back Rates
By Racial Soundingness of Names and Resume Quality [a]

Panel A: Subjective Measure of Quality

	Low	High	Ratio	Difference (p-value)
White Names	**8.80%**	**11.31%**	**1.29**	**2.51%**
	[1216]	[1229]		(.0391)
African American Names	**6.41%**	**6.99%**	**1.09**	**0.58%**
	[1216]	[1229]		(.5644)

Panel B: Predicted Measure of Quality

	Low	High	Ratio	Difference (p-value)
White Names	**5.04%**	**14.18%**	**2.81**	**9.14%**
	[834]	[804]		(.0000)
African American Names	**5.14%**	**8.58%**	**1.66**	**3.44%**
	[817]	[816]		(.0060)

[a] Notes:

1. Panel A reports the mean call-back rates for applicants with a White sounding name (raw 1) and African American sounding name (raw 2) depending on whether the resume was subjectively qualified as a lower quality (column 1) or higher quality (column 2). In brackets is the number of resumes sent for each race/quality group. Column 4 reports the p-value of a test of proportion testing the null hypothesis that the call-back rates are equal across quality groups within each racial group.
2. For Panel B, we use a third of the sample to estimate a probit regression of the call-back dummy on the set of resume characteristics as displayed in Table 3. We further control for a sex dummy, a city dummy, 6 occupation dummies and a vector of dummy variables for job requirements as listed in the employment ad (see Section 4.4 for details). We then use the estimated coefficients on the set of resume characteristics to estimate a predicted call-back for the remaining resumes (2/3 of the sample). We call "high quality" resumes the resumes that rank above the median predicted call-back and "low quality" resumes the resumes that rank below the median predicted call-back. In brackets is the number of resumes sent for each race/quality group. Column 4 reports the p-value of a test of proportion testing the null hypothesis that the call-back rates are equal across quality groups within each racial group.

Table 5
Effect of Resume Characteristics on Likelihood of Call-Back [a]

Dependent Variable: Call-Back Dummy

Sample:	**All Resumes**	**White Names**	**African American Names**
Years of experience (*10)	.07	.13	.02
	(.03)	(.04)	(.03)
Years of experience^2 (*100)	-.02	-.04	-.00
	(.01)	(.02)	(.01)
Volunteering? (Y=1)	-.01	-.01	.00
	(.01)	(.02)	(.01)
Military experience? (Y=1)	-.00	.01	-.01
	(.02)	(.02)	(.02)
Email? (Y=1)	.02	.03	.00
	(.01)	(.01)	(.01)
Employment holes? (Y=1)	.02	.03	.01
	(.01)	(.02)	(.01)
Work in school? (Y=1)	.01	.02	-.00
	(.01)	(.01)	(.01)
Honors? (Y=1)	.05	.07	.02
	(.02)	(.03)	(.02)
Computer skills? (Y=1)	-.02	-.03	-.00
	(.01)	(.02)	(.01)
Special skills? (Y=1)	.05	.07	.04
	(.01)	(.02)	(.01)
Ho: Resume characteristics effects are all zero (p-value)	55.73 (.0000)	59.83 (.0000)	20.78 (.0227)
Standard deviation of predicted call-back	.047	.064	.034
Sample size	4890	2445	2445

[a] Notes:

1. Each column gives the results of a probit regression where the dependent variable is the call-back dummy. Reported in the table are the estimated marginal change in probability for the continuous variables and the estimated discrete change for the dummy variables. Also included in each regression are a city dummy, a sex dummy, 6 occupation dummies and a vector of dummy variables for job requirements as listed in the employment ad (see Section 4.4 for details).
2. Sample in column (1) is the entire set of sent resumes; sample in column (2) is the set of resumes with White sounding names; sample in column (3) is the set of resumes with African American sounding names.
3. Standard errors are corrected for clustering of the observations at the employment ad level.
4. Reported in the second to last row are the p-value for a χ^2 testing that the effects on the resume characteristics are all zero.
5. Reported in the last row is the standard deviation of the predicted call-back rate.

Table 6
Effect of Applicant's Address on Likelihood of Call-Back [a]

Dependent Variable: Call-Back Dummy

Zip code characteristic:	**Fraction Whites**		**Fraction college or more**		**Log(per capita income)**	
Zip code characteristic	.021	.023	.057	.055	.019	.014
	(.012)	(.016)	(.023)	(.031)	(.007)	(.010)
Zip code characteristic* African American name	—	-.005	—	.002	—	.010
		(.025)		(.050)		(.015)
African American name	—	-.030	—	-.033	—	-.134
		(.014)		(.013)		(.157)

[a] Notes:

1. Each column gives the results of a probit regression where the dependent variable is the call-back dummy. Reported in the table are the estimated marginal change in probability. Also included in columns (1), (3) and (5) is a city dummy; also included in columns (2), (4) and (6) is a city dummy and a city dummy interacted with a race dummy.
2. Sample in all regressions is the entire set of sent resumes ($N = 4890$).
3. Standard errors are corrected for clustering of the observations at the employment ad level.

Table 7
Effect of Job Requirements
on Racial Differences in Call-Backs [a]

Requirement:	**Sample Mean (st. dev.)**	**Marginal Effect on Call-Backs for African-American Names**
Any requirement? (Y=1)	.80	.024
	(.41)	(.015)
Experience? (Y=1)	.43	-.019
	(.49)	(.012)
of which:		
some	50.2%	
two years or less	24.8%	
three years or more	25.0%	
Computer skills? (Y=1)	.44	-.004
	(.50)	(.013)
Communication skills? (Y=1)	.12	.000
	(.33)	(.018)
Organization skills? (Y=1)	.07	.019
	(.26)	(.026)
Education? (Y=1)	.11	-.025
	(.31)	(.019)
of which:		
high school degree	8.8%	
some college	48.5%	
4-year college degree	42.7%	
Total number of requirements	1.18	.001
	(.93)	(.006)

[a] Notes:

1. Column (2) reports means and standard deviations (in parentheses) for the job requirements.
2. Column (3) reports the marginal effect of the job requirement listed in that row on discrimination. Specifically, each cell in column (3) corresponds to a different probit regression of the call-back dummy on on African American name dummy, a dummy for the requirement listed in that row and the interaction of the requirement dummy with the African American name dummy. Reported in each cell is the estimated discrete change for the interaction term. Standard errors are corrected for clustering of the observations at the employment ad level.
3. Sample is all sent resumes ($N = 4890$).

Table 8
Racial Gap in Call-Back by Occupation and Industry [a]

Panel A: Occupation Break-Down

	% of Ads	*Call-back Rates for*		*Ratio*	*Diff.*	*1990 Census*	
		White Names	*Afr. Am. Names*			*Log(W)*	*Racial Gap in Log(W)*
Executives and managers	14.5%	7.91%	5.95%	1.33	1.96%	.29	.18
Administrative supervisors	7.7%	9.57%	5.85%	1.64	3.72%	.23	-.02
Sales representatives	15.2%	8.04%	5.09%	1.58	2.95%	.17	.31
Sales workers	16.8%	10.46%	7.05%	1.48	3.41%	-.45	.08
Secretaries	33.9%	10.49%	6.63%	1.58	3.86%	.11	-.07
Clerical workers	11.9%	13.75%	9.96%	1.38	3.79%	-.08	-.01

H_0: Racial gap is the same across occupations
p-value=.975

Panel B: Industry Break-Down

	% of Ads	*Call-back Rates for*		*Ratio*	*Diff.*	*1990 Census*	
		White Names	*Afr. Am. Names*			*Log(W)*	*Racial Gap in Log(W)*
Manufacturing	8.3%	6.93%	3.96%	1.75	2.97%	.14	.15
Transportation and communication	3.0%	12.16%	14.86%	.82	-2.70%	.21	.11
Wholesale and retail trade	21.5%	8.76%	5.71%	1.53	3.05%	-.17	.19
Finance, insurance and real estate	8.5%	10.63%	4.35%	2.44	6.28%	.17	.11
Business and personal services	26.8%	11.30%	6.71%	1.68	4.59%	-.15	.21
Health, educational and social services	15.5%	12.14%	9.50%	1.28	2.64%	-.04	.13
Other/unknown	16.4%	8.71%	6.47%	1.35	2.24%	—	—

H_0: Racial gap is the same across industries
p-value=.1923

[a] Notes:

1. This table reports call-back rates by race and occupation (Panel A) and by race and industry (Panel B). Sample is all sent resumes ($N = 4890$).
2. The two tests reported in the table are log-likelihood tests obtained from two separate probit regressions. In Panel A, we regress the call-back dummy on 6 occupation dummies, a black dummy and the interaction of the black dummy with the six occupation dummies. In Panel B, we regress the call-back dummy on 7 industry dummies, a black dummy and the interactions of the black dummy with the 7 industry dummies. In each case, the null hypothesis tested is that the interaction term effects are all the same.
3. The last two columns of the table report earnings statistics from the 1990 5 percent Census for Boston and Chicago. The first of these two columns reports mean log(residual annual earnings) in the occupation or industry category. The second column reports mean white-black gap in log(residual annual earnings) in the occupation or industry category. Log(residual annual earnings) are obtained from a micro wage regression of log(annual earnings) on 8 education dummies, a quadratic in age, a sex dummy and a city dummy. See text for details.

Table 9
Effect of Employer Characteristics on Racial Differences in Call-Backs[a]

Characteristic:	**Sample Mean (st. dev.)**	**Marginal Effect on Call-Backs for African-American Names**
Equal opportunity employer? (Y=1)	.29	-.010
(N=4890)	(.45)	(.012)
Federal contractor? (Y=1)	.11	-.027
(N=3118)	(.31)	(.018)
Log(employment)	5.74	-.000
(N=1702)	(1.74)	(.032)
Ownership status:		
(N=2894)		
Privately held	73.0%	.003
		(.019)
Publicly traded	15.5%	-.023
		(.016)
Not-for-profit	11.5%	.040
		(.045)

[a]Notes:

1. Column (2) reports means and standard deviations (in parentheses) for the employer characteristics. Sample sizes for each characteristic are reported in column (1).
2. Column (3) reports the marginal effect of the employer characteristic listed in that row on discrimination. Specifically, each cell in column (3) corresponds to a different probit regression of the call-back dummy on an African American name dummy, a dummy for the employer characteristic listed in that row and the interaction of the employer characteristic with the African American name dummy. Reported in each cell is the estimated coefficient on the interaction term. Standard errors are corrected for clustering of the observations at the employment ad level.

Table 10
Effect of Employer's Address
on Racial Differences in Call-Backs [a]

Dependent Variable: Call-Back Dummy

Sample:	**Both Cities**		**Chicago**		**Boston**	
African American name	-.039	—	-.055	—	-.019	—
	(.010)		(.013)		(.019)	
% blacks in employer's zip code	-.008	-.012	-.044	-.058	.205	.171
	(.054)	(.048)	(.048)	(.047)	(.172)	(.110)
African American name*	.059	.075	.087	.086	.016	-.032
%blacks in employer's zip code	(.071)	(.060)	(.044)	(.046)	(.307)	(.201)
Industry dummies	No	Yes	No	Yes	No	Yes
African American name* Industry dummies	No	Yes	No	Yes	No	Yes
Occupation dummies	No	Yes	No	Yes	No	Yes
African American name* Occupation dummies	No	Yes	No	Yes	No	Yes
City dummy	No	Yes	No	Yes	No	Yes
African American name* City dummy	No	Yes	No	Yes	No	Yes
Mean % African Americans in employer's zip code	.082		.106		.047	
	(.154)		(.185)		(.085)	
Sample size	1930		1142		788	

[a] Notes:

1. Each column gives the results of a probit regression where the dependent variable is a call-back dummy. Reported in the table are the estimated marginal change in probability for the continous variable and the estimated discrete change for the dummy variables.
2. Sample in all regressions is the set of sent resumes for which we could determine the employer's zip code.
3. Standard errors are corrected for clustering of the observations at the employment ad level.

Table 11
Call-Back Rates and Mother's Education by First Name[a]

White Female			African American Female		
Name	Call-back	Mother Education	Name	Call-back	Mother Education
Emily	8.3%	96.6%	Aisha	2.2%	77.2%
Anne	9.0%	93.1%	Keisha	3.8%	68.8%
Jill	9.3%	92.3%	Tamika	5.4%	61.5%
Allison	9.4%	95.7%	Lakisha	5.5%	55.6%
Sarah	9.8%	93.4%	Tanisha	6.3%	64.0%
Meredith	10.6%	97.9%	Latoya	8.8%	55.5%
Laurie	10.8%	81.8%	Kenya	9.1%	70.2%
Carrie	13.1%	80.7%	Latonya	9.1%	31.3%
Kristen	13.6%	93.4%	Ebony	10.5%	65.6%
Average		91.7%	Average		61.0%
Overall		83.9%	Overall		70.2%
Correlation	-.350	(p=.3558)	Correlation	-.326	(p=.391)

White Male			African American Male		
Name	Call-back	Mother Education	Name	Call-back	Mother Education
Neil	6.6%	85.7%	Rasheed	3.0%	77.3%
Geoffrey	6.8%	96.0%	Tremayne	4.3%	—
Brett	6.8%	93.9%	Kareem	4.7%	67.4%
Brendan	7.7%	96.7%	Darnell	4.8%	66.1 %
Greg	7.8%	88.3%	Tyrone	5.3%	64.0%
Todd	8.7%	87.7%	Jamal	6.6%	73.9%
Matthew	9.0%	93.1%	Hakim	7.3%	73.7%
Jay	13.2%	85.4%	Leroy	9.4%	53.3%
Brad	15.9%	90.5%	Jermaine	11.3%	57.5%
Average		91.7%	Average		66.7%
Overall		83.5%	Overall		68.9%
Correlation	-.276	(p=.472)	Correlation	-.619	(p=.102)

[a]Notes:

1. This table reports, for each first name used in the experiment, call-back rate and average mother education. Average mother education for a given first name is defined as the fraction of babies born with name in Massachusetts between 1970 and 1986 whose mother had at least completed a high school degree (see text for details). Within each sex/race group, first names are ranked by increasing call-back rate. In brackets in each cell is the number of resumes sent in that cell.
2. "Average" reports, within each race-gender group, the average mother education for all the babies born with one of the names used in the experiment. "Overall" reports, within each race-gender group, average mother education for all babies born in Massachusetts between 1970 and 1986 in that race-gender group. "Correlation" reports the Spearman rank order correlation betwen call-back rates and mother education *within* each race-gender group as well as the p-value for the test of independence.

Appendix Table 1
First Names Used in Experiment[a]

White Female		**African American Female**	
Name	Frequency	Name	Frequency
Allison	4.7%	Aisha	3.6%
Anne	5.0%	Ebony	4.3%
Carrie	3.5%	Keisha	3.7%
Emily	4.7%	Kenya	4.0%
Jill	4.2%	Latonya	4.7%
Laurie	4.0%	Lakisha	4.1%
Kristen	4.4%	Latoya	4.6%
Meredith	3.9%	Tamika	5.3%
Sarah	3.9%	Tanisha	4.2%
Fraction of all births:		Fraction of all births	
3.8%		7.1%	

White Male		**African American Male**	
Name	Frequency	Name	Frequency
Brad	1.3%	Darnell	0.9%
Brendan	1.3%	Hakim	1.1%
Geoffrey	1.2%	Jermaine	1.1%
Greg	1.0%	Kareem	1.3%
Brett	1.2%	Jamal	1.2%
Jay	1.4%	Leroy	1.3%
Matthew	1.4%	Rasheed	1.4%
Neil	1.6%	Tremayne	1.4%
Todd	1.4%	Tyrone	1.6%
Fraction of all births:		Fraction of all births	
1.7%		3.1%	

[a] Notes:

1. This table tabulates the different first names used in the experiment and the frequencies with which each of these names was used. Also reported for each race-sex category is the fraction of all births in that race-sex category with these first names (from the Massachusetts birth certificates, 1974 to 1979).