

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

By MARIANNE BERTRAND AND SENDHIL MULLAINATHAN*

We study race in the labor market by sending fictitious resumes to help-wanted ads in Boston and Chicago newspapers. To manipulate perceived race, resumes are randomly assigned African-American- or White-sounding names. White names receive 50 percent more callbacks for interviews. Callbacks are also more responsive to resume quality for White names than for African-American ones. The racial gap is uniform across occupation, industry, and employer size. We also find little evidence that employers are inferring social class from the names. Differential treatment by race still appears to still be prominent in the U.S. labor market. (JEL J71, J64).

Every measure of economic success reveals significant racial inequality in the U.S. 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 as to whether employers treat members of different races differentially. 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 differential treatment by race 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 latter camp even feel that stringent enforcement of affirmative action programs has produced an environment of reverse discrimination. They would argue that faced with identical candi-

dates, employers might favor the African-American one.¹ Data limitations make it difficult to empirically test these views. Since researchers possess far less data 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 differences that are observable to employers but unobservable to researchers.

To circumvent this difficulty, we conduct a field experiment that builds on the correspondence testing methodology that has been primarily used in the past to study minority outcomes in the United Kingdom.² We send resumes in response to help-wanted ads in Chicago and Boston newspapers and measure callback for interview for each sent resume. We

* Bertrand: Graduate School of Business, University of Chicago, 1101 E. 58th Street, R0 229D, Chicago, IL 60637, NBER, and CEPR (e-mail: marianne.bertrand@gsb.uchicago.edu); Mullainathan: Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, E52-380a, Cambridge, MA 02142, and NBER (e-mail: 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 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 See Roger Jowell and Patricia Prescott-Clarke (1970), Jim Hubback and Simon Carter (1980), Colin Brown and Pat Gay (1985), and Peter A. Riach and Judith Rich (1991). One caveat is that some of these studies fail to fully match skills between minority and nonminority resumes. For example some impose differential education background by racial origin. Doris Weichselbaumer (2003, 2004) studies the impact of sex-stereotypes and sexual orientation. Richard E. Nisbett and Dov Cohen (1996) perform a related field experiment to study how employers' response to a criminal past varies between the North and the South in the United States.

experimentally manipulate perception of race via the name of the fictitious job applicant. 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 the racial gap in callback, 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 e-mail address, have completed some certification degree, possess foreign language skills, or have been awarded some honors.³ 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 1,300 employment ads in the sales, administrative support, clerical, and customer services job categories and send nearly 5,000 resumes. The ads we respond to cover a large spectrum of job quality, from cashier work at retail establishments and clerical work in a mail room, to office and sales management positions.

We find large racial differences in callback rates.⁴ Applicants with White names need to send about 10 resumes to get one callback whereas applicants with African-American names need to send about 15 resumes. This 50-percent gap in callback is statistically significant. A White name yields as many more callbacks as an additional eight years of experience on a resume. 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 nearly 30-percent more callbacks than

Whites with lower-quality resumes. On the other hand, having a higher-quality resume has a smaller effect for African-Americans. In other words, the gap between Whites and African-Americans widens with resume quality. While one may have expected improved credentials to alleviate employers' fear that African-American applicants are deficient in some unobservable skills, this is not the case in our data.⁵

The experiment also reveals several other aspects of the differential treatment by race. First, since we randomly assign applicants' postal addresses to the resumes, we can study the effect of neighborhood of residence on the likelihood of callback. We find that living in a wealthier (or more educated or Whiter) neighborhood increases callback rates. But, interestingly, African-Americans are not helped more than Whites by living in a "better" neighborhood. Second, the racial gap we measure in different industries does not appear correlated to Census-based measures of the racial gap in wages. The same is true for the racial gap we measure in different occupations. In fact, we find that the racial gaps in callback are statistically indistinguishable across all the occupation and industry categories covered in the experiment. Federal contractors, who are thought to be more severely constrained by affirmative action laws, do not treat the African-American resumes more preferentially; neither do larger employers or employers who explicitly state that they are "Equal Opportunity Employers." In Chicago, we find a slightly smaller racial gap when employers are located in more African-American neighborhoods.

The rest of the paper is organized as follows. Section I compares this experiment to earlier work on racial discrimination, and most notably to the labor market audit studies. We describe the experimental design in Section II and present the results in Section III, subsection A. In Section IV, we discuss possible interpretations of our results, focusing especially on two issues. First, we examine whether the

³ In creating the higher-quality resumes, we deliberately make small changes in credentials so as to minimize the risk of overqualification.

⁴ For ease of exposition, we refer to the effects uncovered in this experiment as racial differences. Technically, however, these effects are about the racial soundingness of names. We briefly discuss below the potential confounds between name and race. A more extensive discussion is offered in Section IV, subsection B.

⁵ These results contrast with the view, mostly based on nonexperimental evidence, that African-Americans receive higher returns to skills. For example, estimating earnings regressions on several decades of Census data, James J. Heckman et al. (2001) show that African-Americans experience higher returns to a high school degree than Whites do.

race-specific names we have chosen might also proxy for social class above and beyond the race of the applicant. Using birth certificate data on mother's education for the different first names used in our sample, we find little relationship between social background and the name-specific callback rates.⁶ Second, we discuss how our results map back 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 the racial gap across occupations and industries. We conclude that existing models do a poor job of explaining the full set of findings. Section V concludes.

I. Previous Research

With conventional labor force and household surveys, it is difficult to study whether differential treatment occurs in the labor market.⁷ Armed only with survey data, researchers usually measure differential treatment by comparing the labor market performance of Whites and African-Americans (or men and women) for which they observe similar sets 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 minority and nonminority workers being compared are truly similar from the employers' perspective. As a consequence, any measured differences in outcomes could be attributed to these unobserved (to the researcher) factors.

This difficulty with conventional data has led some authors to instead rely on pseudo-experiments.⁸ Claudia Goldin and Cecilia

Rouse (2000), for example, examine the effect of blind auditioning on the hiring process of orchestras. By observing 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 proven to be extremely challenging.

A different set of studies, known as audit studies, attempts to place comparable minority and White actors into actual social and economic settings and measure how each group fares in these settings.⁹ 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.¹⁰ 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 nonlaboratory evidence of differential treatment by race. 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 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. (1991)

⁶ We also argue that a social class interpretation would find it hard to explain some of our findings, such as why living in a better neighborhood does not increase callback rates more for African-American names than for White names.

⁷ See Joseph G. Altonji and Rebecca M. Blank (1999) for a detailed review of the existing literature on racial discrimination in the labor market.

⁸ William A. Darity, Jr. and Patrick L. Mason (1998) describe an interesting nonexperimental study. Prior to the Civil Rights Act of 1964, employment ads would explicitly state racial biases, providing a direct measure of differential treatment. Of course, as Arrow (1998) mentions, discrimination was at that time "a fact too evident for detection."

⁹ Michael Fix and Marjery A. Turner (1998) provide a survey of many such audit studies.

¹⁰ Earlier hiring audit studies include Jerry M. Newman (1978) and Shelby J. McIntyre et al. (1980). Three more recent studies are Harry Cross et al. (1990), Franklin James and Steve W. DelCastillo (1991), and Turner et al. (1991). Heckman and Peter Siegelman (1992), Heckman (1998), and Altonji and Blank (1999) summarize these studies. See also David Neumark (1996) for a labor market audit study on gender discrimination.

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 their beliefs about race issues in America. As psychologists know very well, these demand effects can be quite 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 measured differences in treatment. 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 nuances and possible mitigating factors. Also, these budgetary constraints worsen the problem of mismatched auditor pairs. Cost considerations force the use of a limited number of 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 caused solely by 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 nature of the differential treatment from many more angles.

II. Experimental Design

A. *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 Web sites as the basis for our artificial resumes.¹¹ While the resumes posted on these Web sites 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 detailed 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

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.

experience. This resume quality manipulation needs to be somewhat subtle to avoid making a higher-quality job applicant overqualified 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 use resumes from Chicago job seekers 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 with a Chicago school.¹⁴ We also use the Chicago resumes to replace a Boston employer with a 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 subpopulation of job searchers.¹⁶

B. 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 asked to assess features of a person with a particular name, one of which is race. For each name, 30 respondents were asked to identify the name as either “White,” “African-American,” “Other,” or “Cannot Tell.” 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 is shown in Appendix Table A1. The table reports the relative likelihood of the names for the Whites and African-Americans in the Massachusetts birth certificates data as well as the recognition rate in the field survey.¹⁹ As Appendix Table A1 indicates, the African-American first names used in the experiment are quite common in the population. This suggests that by using these names as an indicator of race, we are actually covering a rather 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 mailbox attached to them. A similar outgoing message is recorded on each of the voice mailboxes but each message is recorded by someone of the appropriate race and gender.

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 So many of names show a likelihood ratio of ∞ because there is censoring of the data at five births. If there are fewer than five babies in any race/name cell, it is censored (and we do not know whether a cell has zero or was censored). This is primarily a problem for the computation of how many African-American babies have “White” names.

²⁰ 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.

¹³ In Section III, subsection B, and Table 3, we provide a detailed summary of resume characteristics by quality level.

¹⁴ We try as much as possible to match high schools and colleges on quality and demographic characteristics.

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.

Since we allocate the same phone number for applicants with different names, we cannot use a person name in the outgoing message.

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 three addresses in each 5-digit zip code in Boston and Chicago. Within cities, we randomly assign addresses across all resumes. We also create eight e-mail addresses, four for Chicago and four for Boston.²¹ These e-mail addresses are neutral with respect to both race and sex. Not all applicants are given an e-mail address. The e-mail 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 that we can assign to the template resumes when responding to the employment ads.

C. Responding to Ads

The experiment was carried out between July 2001 and January 2002 in Boston and between July 2001 and May 2002 in Chicago.²² Over that period, we surveyed 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 were asked to call or appear in person. In fact, most of the ads we surveyed in these job categories ask 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-quality 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 random. The resumes are then faxed (or in a few cases mailed) to the employer. All in all, we respond to more than 1,300 employment ads over the entire sample period and send close to 5,000 resumes.

D. Measuring Responses

We measure whether a given resume elicits a callback or e-mail back for an interview. For each phone or e-mail 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.²⁷ 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

²³ In some instances, our resume bank does not have four resumes that are appropriate matches 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 e-mail addresses and a few low-quality resumes were given e-mail addresses. See Table 3 for details.

²⁷ Very few employers used e-mail to contact an applicant back.

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

²² This period spans tighter and slacker labor markets. In our data, this is apparent as callback rates (and number of new ads) dropped after September 11, 2001. Interestingly, however, the racial gap we measure is the same across these two periods.

TABLE 1—MEAN CALLBACK RATES BY RACIAL SOUNDINGNESS OF NAMES

	Percent callback for White names	Percent callback for African-American names	Ratio	Percent difference (<i>p</i> -value)
Sample:				
All sent resumes	9.65 [2,435]	6.45 [2,435]	1.50	3.20 (0.0000)
Chicago	8.06 [1,352]	5.40 [1,352]	1.49	2.66 (0.0057)
Boston	11.63 [1,083]	7.76 [1,083]	1.50	4.05 (0.0023)
Females	9.89 [1,860]	6.63 [1,886]	1.49	3.26 (0.0003)
Females in administrative jobs	10.46 [1,358]	6.55 [1,359]	1.60	3.91 (0.0003)
Females in sales jobs	8.37 [502]	6.83 [527]	1.22	1.54 (0.3523)
Males	8.87 [575]	5.83 [549]	1.52	3.04 (0.0513)

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

employers rarely, if ever, contact applicants via postal mail to set up interviews.

E. 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 applicant 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 gaps in 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. On a related note, because we are not assigning race but only race-specific names, our results are not representative of the average African-American (who may not have such a racially distinct

name).²⁸ We return to this issue in Section IV, subsection B.

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 previous work, social networks are another common means through which people find jobs and one that clearly cannot be studied here. This omission could qualitatively affect our results if African-Americans use social networks more or if employers who rely more on networks differentiate less by race.²⁹

III. Results

A. Is There a Racial Gap in Callback?

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

²⁸ As Appendix Table A1 indicates, the African-American names we use are, however, quite common among African-Americans, making this less of a concern.

²⁹ In fact, there is some evidence that African-Americans may rely less on social networks for their job search (Harry J. Holzer, 1987).

names have a 9.65 percent chance of receiving a callback. Equivalent resumes with African-American names have a 6.45 percent chance of being called back. This represents a difference in callback rates of 3.20 percentage points, or 50 percent, that can solely be attributed to the name manipulation. Column 4 shows that this difference is statistically significant.³⁰ 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 about 15 different ads to achieve the same result.³¹

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 in sending more resumes 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 a 0.4 percentage point. Based on this point estimate, the return to a White name is equivalent to about eight 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 the two 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 49 percent more likely

than African-American applicants to receive a callback in Chicago and 50 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.³² Looking across occupations, we find a significant racial gap in callbacks for both males (52 percent) and females (49 percent). Comparing males to females in sales occupations, we find a larger racial gap among males (52 percent versus 22 percent). Interestingly, females in sales jobs appear to receive more callbacks than males; however, this (reverse) gender gap is statistically insignificant and economically much 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 three categories above can each take three 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 all other cases.

As Table 2 indicates, equal treatment occurs for about 88 percent of the help-wanted ads. As expected, the major source of equal treatment comes from the high fraction of ads for which

³⁰ 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.

³¹ This obviously assumes that African-American applicants cannot assess a priori which firms are more likely to treat them more or less favorably.

³² Only about 6 percent of all male resumes were sent in response to an administrative job opening.

TABLE 2—DISTRIBUTION OF CALLBACKS BY EMPLOYMENT AD

	No Callback	1W + 1B	2W + 2B
Equal Treatment:			
88.13 percent	83.37	3.48	1.28
[1,166]	[1,103]	[46]	[17]
Whites Favored (WF):	1W + 0B	2W + 0B	2W + 1B
8.39 percent	5.59	1.44	1.36
[111]	[74]	[19]	[18]
African-Americans Favored (BF):	1B + 0W	2B + 0W	2B + 1W
3.48 percent	2.49	0.45	0.53
[46]	[33]	[6]	[7]
<i>Ho: WF = BF</i>			
<i>p</i> = 0.0000			

Notes: This table documents the distribution of callbacks at the employment-ad level. “No Callback” is the percent of ads for which none of the fictitious applicants received a callback. “1W + 1B” is the percent of ads for which exactly one White and one African-American applicant received a callback. “2W + 2B” is the percent of ads for which exactly two White applicants and two African-American applicants received a callback. “Equal Treatment” is defined as the sum of “No Callback,” “1W + 1B,” and “2W + 2B.” “1W + 0B” is the percent of ads for which exactly one White applicant and no African-American applicant received a call back. “2W + 0B” is the percent of ads for which exactly two White applicants and no African-American applicant received a callback. “2W + 1B” is the percent of ads for which exactly two White applicants and one African-American applicant received a callback. “Whites Favored” is defined as the sum of “1W + 0B,” “2W + 0B,” and “2W + 1B.” “1B + 0W” is the percent of ads for which exactly one African-American applicant and no White applicant received a callback. “2B + 0W” is the percent of ads for which exactly two African-American applicants and no White applicant received a callback. “2B + 1W” is the percent of ads for which exactly two African-American applicants and one White applicant received a callback. “African-Americans Favored” is defined as the sum of “1B + 0W,” “2B + 0W,” and “2B + 1W.” In brackets in each cell is the number of employment ads in that cell. “*Ho: WF = WB*” reports the *p*-value for a test of symmetry between the proportion of employers that favor White names and the proportion of employers that favor African-American names.

no callbacks are recorded (83 percent of the ads). Whites are favored by nearly 8.4 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 3.5 percent of employers. We formally test whether there is symmetry in the favoring of Whites over African-Americans and African-Americans over Whites. We find that the difference between the fraction of employers favoring Whites and the fraction of employers favoring African-Americans is statistically very significant ($p = 0.0000$).

B. Do African-Americans Receive Different Returns to Resume Quality?

Our results so far demonstrate a substantial gap in callback based on applicants’ names. Next, we would like to learn more about the factors that may influence this gap. 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 varies by resume quality.

As we explained in Section II, 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. Table 3 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 quality. Higher-quality applicants have on average close to an extra year of labor market experience, fewer employment holes (where an employment hole is defined as a period of at least six 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 e-mail 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

1000

THE AMERICAN ECONOMIC REVIEW

SEPTEMBER 2004

TABLE 3—RESUME CHARACTERISTICS: SUMMARY STATISTICS

Sample:	All resumes	White names	African-American	Higher quality	Lower quality
Characteristic:					
College degree	0.72	0.72	0.72	0.72	0.71
(Y = 1)	(0.45)	(0.45)	(0.45)	(0.45)	(0.45)
Years of experience	7.84	7.86	7.83	8.29	7.39
	(5.04)	(5.07)	(5.01)	(5.29)	(4.75)
Volunteering experience?	0.41	0.41	0.41	0.79	0.03
(Y = 1)	(0.49)	(0.49)	(0.49)	(0.41)	(0.16)
Military experience?	0.10	0.09	0.10	0.19	0.00
(Y = 1)	(0.30)	(0.29)	(0.30)	(0.39)	(0.06)
E-mail address?	0.48	0.48	0.48	0.92	0.03
(Y = 1)	(0.50)	(0.50)	(0.50)	(0.27)	(0.17)
Employment holes?	0.45	0.45	0.45	0.34	0.56
(Y = 1)	(0.50)	(0.50)	(0.50)	(0.47)	(0.50)
Work in school?	0.56	0.56	0.56	0.72	0.40
(Y = 1)	(0.50)	(0.50)	(0.50)	(0.45)	(0.49)
Honors?	0.05	0.05	0.05	0.07	0.03
(Y = 1)	(0.22)	(0.23)	(0.22)	(0.25)	(0.18)
Computer skills?	0.82	0.81	0.83	0.91	0.73
(Y = 1)	(0.38)	(0.39)	(0.37)	(0.29)	(0.44)
Special skills?	0.33	0.33	0.33	0.36	0.30
(Y = 1)	(0.47)	(0.47)	(0.47)	(0.48)	(0.46)
Fraction high school dropouts in applicant's zip code	0.19	0.19	0.19	0.19	0.18
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
Fraction college or more in applicant's zip code	0.21	0.21	0.21	0.21	0.22
	(0.17)	(0.17)	(0.17)	(0.17)	(0.17)
Fraction Whites in applicant's zip code	0.54	0.55	0.54	0.53	0.55
	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)
Fraction African-Americans in applicant's zip code	0.31	0.31	0.31	0.32	0.31
	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)
Log(median per capital income) in applicant's zip code	9.55	9.55	9.55	9.54	9.56
	(0.56)	(0.56)	(0.55)	(0.54)	(0.57)
Sample size	4,870	2,435	2,435	2,446	2,424

Notes: The table reports means and standard deviations for the resume characteristics as listed on the left. Column 1 refers to all resumes sent; column 2 refers to resumes with White names; column 3 refers to resumes with African-American names; column 4 refers to higher-quality resumes; column 5 refers to lower-quality resumes. See text for details.

applicants' education level. This reflects the fact that all sent resumes, whether high or low quality, are chosen to be good matches for a given job opening. About 70 percent of the sent resumes report a college degree.³³

The last five rows of Table 3 show summary characteristics of the applicants' zip code address. Using 1990 Census data, we compute the fraction of high school dropouts, fraction of college educated or more, fraction of Whites, fraction of African-Americans and log(median per capital 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-quality resumes are presented in Panel A of Table 4. The first thing to note is that the resume quality manipulation works: higher-quality resumes receive more callbacks. As row 1 indicates, we record a callback rate of close to 11 percent for White applicants with a higher-quality resume, compared to 8.5 percent for White applicants with lower-quality resumes. This is a statistically significant difference of 2.29 percentage points, or 27 percent ($p = 0.0557$). Most strikingly, African-Americans experience much less of an increase in callback

³³ 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.

TABLE 4—AVERAGE CALLBACK RATES BY RACIAL SOUNDINGNESS OF NAMES AND RESUME QUALITY

Panel A: Subjective Measure of Quality (Percent Callback)				
	Low	High	Ratio	Difference (<i>p</i> -value)
White names	8.50 [1,212]	10.79 [1,223]	1.27	2.29 (0.0557)
African-American names	6.19 [1,212]	6.70 [1,223]	1.08	0.51 (0.6084)
Panel B: Predicted Measure of Quality (Percent Callback)				
	Low	High	Ratio	Difference (<i>p</i> -value)
White names	7.18 [822]	13.60 [816]	1.89	6.42 (0.0000)
African-American names	5.37 [819]	8.60 [814]	1.60	3.23 (0.0104)

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

rate for similar improvements in their credentials. African-Americans with higher-quality resumes receive a callback 6.7 percent of the time, compared to 6.2 percent for African-Americans with lower quality resumes. This is only a 0.51-percentage-point, or 8-percent, difference and this difference is not statistically significant ($p = 0.6084$).

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, six occupation dummies, and a vector of job requirements as listed in the employment ads.³⁴ We then use the estimated coefficients on the resume characteristics to rank the remaining two-thirds of the resumes by predicted callback. In Panel B, we classify as “high” those resumes that have above-median-predicted callback; similarly, we classify as “low” those resumes

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 less than Whites. The ratio of callback rates for high- versus low-quality resumes is 1.60 for African Americans, compared to 1.89 for Whites.

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 e-mail address, honors, and special skills all have a positive and significant effect on the likelihood of a callback.³⁵ Also, more experienced applicants are more likely to get called back: at the average number of years of experience in our sample (eight years), each

³⁴ See Section III, subsection D, for more details on these occupation categories and job requirements.

³⁵ 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.

TABLE 5—EFFECT OF RESUME CHARACTERISTICS ON LIKELIHOOD OF CALLBACK

Dependent Variable: Callback Dummy			
Sample:	All resumes	White names	African-American names
Years of experience (*10)	0.07 (0.03)	0.13 (0.04)	0.02 (0.03)
Years of experience ² (*100)	-0.02 (0.01)	-0.04 (0.01)	-0.00 (0.01)
Volunteering? (Y = 1)	-0.01 (0.01)	-0.01 (0.01)	0.01 (0.01)
Military experience? (Y = 1)	-0.00 (0.01)	0.02 (0.03)	-0.01 (0.02)
E-mail? (Y = 1)	0.02 (0.01)	0.03 (0.01)	-0.00 (0.01)
Employment holes? (Y = 1)	0.02 (0.01)	0.03 (0.02)	0.01 (0.01)
Work in school? (Y = 1)	0.01 (0.01)	0.02 (0.01)	-0.00 (0.01)
Honors? (Y = 1)	0.05 (0.02)	0.06 (0.03)	0.03 (0.02)
Computer skills? (Y = 1)	-0.02 (0.01)	-0.04 (0.02)	-0.00 (0.01)
Special skills? (Y = 1)	0.05 (0.01)	0.06 (0.02)	0.04 (0.01)
<i>H₀</i> : Resume characteristics effects are all zero (<i>p</i> -value)	54.50 (0.0000)	57.59 (0.0000)	23.85 (0.0080)
Standard deviation of predicted callback	0.047	0.062	0.037
Sample size	4,870	2,435	2,435

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

extra year of experience increases the likelihood of a callback by about a 0.4 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 e-mail 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 callback by about a 0.7 percentage point.

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 illustrate this more saliently, we predict callback rates using either regression estimates in column 2 or regression estimates in column 3. The standard deviation of the predicted callback from column 2 is 0.062, whereas it is only 0.037 from column 3. In summary, employers simply seem to pay less attention or discount more the characteristics listed on the

TABLE 6—EFFECT OF APPLICANT'S ADDRESS ON LIKELIHOOD OF CALLBACK

Dependent Variable: Callback Dummy						
Zip code characteristic:	Fraction Whites		Fraction college or more		Log(per capital income)	
Zip code characteristic	0.020 (0.012)	0.020 (0.016)	0.054 (0.022)	0.053 (0.031)	0.018 (0.007)	0.014 (0.010)
Zip code characteristic*	—	−0.000 (0.024)	—	−0.002 (0.048)	—	0.008 (0.015)
African-American name	—	−0.031 (0.015)	—	−0.031 (0.013)	—	−0.112 (0.152)

Notes: Each column gives the results of a probit regression where the dependent variable is the callback dummy. Reported in the table is 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. Standard errors are corrected for clustering of the observations at the employment-ad level.

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.³⁶

C. Applicants' Address

An incidental feature of our experimental design is the random assignment of addresses 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 relatively more by residing in more affluent neighborhoods.

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 callback. For example, a 10-percentage-point increase in the fraction of college-educated in zip code of residence in-

creases the likelihood of a callback by a 0.54 percentage point (column 3).

In columns 2, 4, and 6, we further interact the zip code characteristic 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 Whiter, 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; this effect, however, is statistically insignificant.

In summary, while neighborhood quality affects callbacks, African-Americans do not benefit more than Whites from living in better neighborhoods. If ghettos and bad neighborhoods are particularly stigmatizing for African-Americans, one might have expected African-Americans to be helped more by having a "better" address. Our results do not support this hypothesis.

D. Job and Employer Characteristics

Table 7 studies how various job requirements (as listed in the employment ads) and employer characteristics correlate with the racial gap in callback. Each row of Table 7 focuses on a specific job or employer characteristic, with

³⁶ This of course assumes that the changes in job and wage offers associated with higher skills are the same across races, or at least not systematically larger for African-Americans.

TABLE 7—EFFECT OF JOB REQUIREMENT AND EMPLOYER CHARACTERISTICS ON RACIAL DIFFERENCES IN CALLBACKS

Job requirement:	Sample mean (standard deviation)	Marginal effect on callbacks for African-American names
Any requirement? (Y = 1)	0.79 (0.41)	0.023 (0.015)
Experience? (Y = 1)	0.44 (0.49)	0.011 (0.013)
Computer skills? (Y = 1)	0.44 (0.50)	0.000 (0.013)
Communication skills? (Y = 1)	0.12 (0.33)	−0.000 (0.015)
Organization skills? (Y = 1)	0.07 (0.26)	0.028 (0.029)
Education? (Y = 1)	0.11 (0.31)	−0.031 (0.017)
Total number of requirements	1.18 (0.93)	0.002 (0.006)
Employer characteristic:	Sample mean (standard deviation)	Marginal effect on callbacks for African-American names
Equal opportunity employer? (Y = 1)	0.29 (0.45)	−0.013 (0.012)
Federal contractor? (Y = 1) (N = 3,102)	0.11 (0.32)	−0.035 (0.016)
Log(employment) (N = 1,690)	5.74 (1.74)	−0.001 (0.005)
Ownership status: (N = 2,878)		
Privately held	0.74	0.011 (0.019)
Publicly traded	0.15	−0.025 (0.015)
Not-for-profit	0.11	0.025 (0.042)
Fraction African-Americans in employer's zip code (N = 1,918)	0.08 (0.15)	0.117 (0.062)

Notes: Sample is all sent resumes (N = 4,870) unless otherwise specified in column 1. Column 2 reports means and standard deviations (in parentheses) for the job requirement or employer characteristic. For ads listing an experience requirement, 50.1 percent listed “some,” 24.0 percent listed “two years or less,” and 25.9 percent listed “three years or more.” For ads listing an education requirement, 8.8 percent listed a high school degree, 48.5 percent listed some college, and 42.7 percent listed at least a four-year college degree. Column 3 reports the marginal effect of the job requirement or employer characteristic listed in that row on differential treatment. Specifically, each cell in column 3 corresponds to a different probit regression of the callback dummy on an African-American name dummy, a dummy for the requirement or characteristic listed in that row and the interaction of the requirement or characteristic dummy with the African-American name dummy. Reported in each cell is the estimated change in probability for the interaction term. Standard errors are corrected for clustering of the observations at the employment-ad level.

summary statistics in column 2. Column 3 shows the results of various probit regressions. Each entry in this column is the marginal effect of the specific characteristic listed in that row on the racial gap in callback. More specifically, each entry is from a separate probit regression of a callback dummy on an African-American dummy, the characteristic listed in that row and the interaction of that characteristic with the

African-American dummy. The reported coefficient is that on the interaction term.

We start with job requirements. About 80 percent of the ads state some form of requirement. About 44 percent of the ads require some minimum experience, of which roughly 50 percent simply ask for “some experience,” 24 percent less than two years, and 26 percent at least three 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, 48.5 percent some college (such as an associate degree), and the rest at least a four-year college degree.³⁷

Despite this variability, we find little systematic relationship between any of the requirements and the racial gap in callback. The point estimates in column 3 show no consistent economic pattern and are all statistically weak. Measures of job quality, such as experience or computer skills requirements, do not predict the extent of the racial gap. Communication or other interpersonal skill requirements have no effect on the racial gap either.³⁸

We also study employer characteristics. Collecting such information is a more difficult task since it is not readily available from the employment ads we respond to. 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 therefore have to turn to supplemental data sources. For employment ads that do not list a specific employer, we first use the fax number to try to identify the company name via Web reverse-lookup services. Based on 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 Web site to find a list of companies that have federal contracts.³⁹ The racial difference in callback rates for the subsamples where employer characteristics could be determined is very similar in magnitude to that in the full sample.

Employer characteristics differ significantly across ads. 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 2,000 employees but there is a lot of variation across firms. Finally, 74 percent of the firms are privately held, 15 percent are publicly traded, and 11 percent are not-for-profit organizations.

Neither “Equal Opportunity Employers” nor federal contractors appear to treat African-Americans more favorably. In fact, each of these employer characteristics is associated with a larger racial gap in callback (and this effect is marginally significant for federal contractors). Differential treatment does not vary with employer size.⁴⁰ Point estimates indicate less differential treatment in the not-for-profit sector; however, this effect is very noisily estimated.⁴¹

In an unpublished Appendix (available from the authors upon request), we also study how the racial gap in callback varies by occupation and industry. Based on the employment ad listings, we classify the job openings into six occupation categories: executives and managers; administrative supervisors; sales representatives; sales workers; secretaries and legal assistants; clerical workers. We also, when possible,

³⁷ Other requirements sometimes mentioned include typing skills for secretaries (with specific words-per-minute minimum thresholds), and, more rarely, foreign language skills.

³⁸ Other ways of estimating these effects produce a similar nonresult. 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.

³⁹ This Web site (www.fpdc.gov) is accurate up to and including March 21, 2000.

⁴⁰ Similar results hold when we measure employer size using a total sales measure rather than an employment measure.

⁴¹ Our measurement of the racial gap by firm or employer 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.

classify employers into six industry categories: manufacturing; transportation and communication; wholesale and retail trade; finance, insurance, and real estate; business and personal services; health, educational, and social services. We then compute occupation and industry-specific racial gaps in callback and relate these gaps to 1990 Census-based measures of occupation and industry earnings, as well as Census-based measures of the White/African-American wage gap in these occupations and industries.

We find a positive White/African-American gap in callbacks in all occupation and industry categories (except for transportation and communication). While average earnings vary a lot across the occupations covered in the experiment, we find no systematic relationship between occupation earnings and the racial gap in callback. Similarly, the industry-specific gaps in callback do not relate well to a measure of inter-industry wage differentials. In fact, while the racial gap in callback rates varies somewhat across occupations and industries, we cannot reject the null hypothesis that the gap is the same across all these categories.

The last row of Table 7 focuses on the marginal effect of employer location on the racial gap in callback.⁴² 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 differential treatment by race varies with the fraction of African-Americans in the employer's zip code. We find a marginally significant positive effect of employer location on African-American callbacks but this effect is extremely small. In regressions not reported here (but available from the authors upon request), we reestimate this effect separately by city. While the point estimates are positive for both cities, the effect is only statistically significant for Chicago.

IV. Interpretation

Three main sets of questions arise when interpreting the results above. First, does a higher callback rate for White applicants imply that employers are discriminating against African-

Americans? Second, does our design only isolate the effect of race or is the name manipulation conveying some other factors than race? Third, how do our results relate to different models of racial discrimination?

A. Interpreting Callback Rates

Our results indicate that for two identical individuals engaging in an identical job search, the one with an African-American name would receive fewer interviews. Does differential treatment within our experiment imply that employers are discriminating against African-Americans (whether it is rational, prejudice-based, or other form of discrimination)? In other words, could the lower callback rate we record for African-American resumes *within our experiment* be consistent with a racially neutral review of the *entire pool* of resumes the surveyed employers receive?

In a racially neutral review process, employers would rank order resumes based on their quality and call back all applicants that are above a certain threshold. Because names are randomized, the White and African-American resumes we send should rank similarly on average. So, irrespective of the skill and racial composition of the applicant pool, a race-blind selection rule would generate equal treatment of Whites and African-Americans. So our results must imply that employers use race as a factor when reviewing resumes, which matches the legal definition of discrimination.

But even rules where employers are not trying to interview as few African-American applicants as possible may generate observed differential treatment in our experiment. One such hiring rule would be employers trying to interview a target level of African-American candidates. For example, perhaps the average firm in our experiment aims to produce an interview pool that matches the population base rate. This rule could produce the observed differential treatment if the average firm receives a higher proportion of African-American resumes than the population base rate because African-Americans disproportionately apply to the jobs and industries in our sample.⁴³

⁴² For previous work on the effect of employer location on labor market discrimination, see, for example, Steven Raphael et al. (2000).

⁴³ Another variant of this argument is that the (up to) two African-American resumes we sent are enough to signifi-

Some of our other findings may be consistent with such a rule. For example, the fact that “Equal Opportunity Employers” or federal contractors do not appear to discriminate any less may reflect the fact that such employers receive more applications from African-Americans. On the other hand, other key findings run counter to this rule. As we discuss above, we find no systematic difference in the racial gap in callback across occupational or industry categories, despite the large variation in the fraction of African-Americans looking for work in those categories. African-Americans are underrepresented in managerial occupations, for example. If employers matched base rates in the population, the few African-Americans who apply to these jobs should receive a higher callback rate than Whites. Yet, we find that the racial gap in managerial occupations is the same as in all the other job categories. This rule also runs counter to our findings on returns to skill. Suppose firms are struggling to find White applicants but overwhelmed with African-American ones. Then they should be less sensitive to the quality of White applicants (as they are trying to fill in their hiring quota for Whites) and much more sensitive to the quality of Black applicants (when they have so many to pick from). Thus, it

cantly distort the racial composition of the entire applicant pool. This is unlikely for two reasons. First, anecdotal evidence and the empirically low callback rates we record suggest that firms typically receive many hundreds of resumes in response to each ad they post. Hence, the (up to) four resumes we send out are unlikely to influence the racial composition of the pool. Second, the similar racial gap in callback we observe across the two cities goes counter to this interpretation since the racial composition base rates differ quite a lot across these two cities. Another variant of this argument is that, for some reason, the average firm in our sample receives a lot of high-quality resumes from African-American applicants and much fewer high-quality resumes from White applicants. Hypothetically, this might occur if high-quality African-Americans are much more likely to use help-wanted ads rather than other job search channels. If employers perform within-race comparisons and again want to target a certain racial mix in their interviewing and hiring, our African-American resumes may naturally receive lower callbacks as they are competing with many more high-quality applicants. This specific argument would be especially relevant in a case where the average sampled employer is “known” to be good to African-Americans. But our selection procedure for the employment ads did not allow for such screening: we simply responded to as many ads as possible in the targeted occupational categories.

is unlikely that the differential treatment we observe is generated by hiring rules such as these.

B. *Potential Confounds*

While the names we have used in this experiment strongly signal racial origin, they may also signal some other personal trait. 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.⁴⁴ 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.⁴⁵

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 we have chosen mainly signal negative social background, one might have expected the estimated name gap to be lower for 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 interpersonal interactions. We found no such evidence in Table 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

⁴⁴ Roland Fryer and Steven Levitt (2003) provide a recent analysis of social background and naming conventions amongst African-Americans.

⁴⁵ 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.

TABLE 8—CALLBACK RATE AND MOTHER'S EDUCATION BY FIRST NAME

White female			African-American female		
Name	Percent callback	Mother education	Name	Percent callback	Mother education
Emily	7.9	96.6	Aisha	2.2	77.2
Anne	8.3	93.1	Keisha	3.8	68.8
Jill	8.4	92.3	Tamika	5.5	61.5
Allison	9.5	95.7	Lakisha	5.5	55.6
Laurie	9.7	93.4	Tanisha	5.8	64.0
Sarah	9.8	97.9	Latoya	8.4	55.5
Meredith	10.2	81.8	Kenya	8.7	70.2
Carrie	13.1	80.7	Latonya	9.1	31.3
Kristen	13.1	93.4	Ebony	9.6	65.6
Average		91.7	Average		61.0
Overall		83.9	Overall		70.2
Correlation	−0.318	($p = 0.404$)	Correlation	−0.383	($p = 0.309$)

White male			African-American male		
Name	Percent callback	Mother education	Name	Percent callback	Mother education
Todd	5.9	87.7	Rasheed	3.0	77.3
Neil	6.6	85.7	Tremayne	4.3	—
Geoffrey	6.8	96.0	Kareem	4.7	67.4
Brett	6.8	93.9	Darnell	4.8	66.1
Brendan	7.7	96.7	Tyrone	5.3	64.0
Greg	7.8	88.3	Hakim	5.5	73.7
Matthew	9.0	93.1	Jamal	6.6	73.9
Jay	13.4	85.4	Leroy	9.4	53.3
Brad	15.9	90.5	Jermaine	9.6	57.5
Average		91.7	Average		66.7
Overall		83.5	Overall		68.9
Correlation	−0.0251	($p = 0.949$)	Correlation	−0.595	($p = 0.120$)

Notes: This table reports, for each first name used in the experiment, callback rate and average mother education. Mother education for a given first name is defined as the percent of babies born with that 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 callback rate. “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 between callback rate and mother education *within* each race-gender group as well as the p -value for the test of independence.

1986.⁴⁶ 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 8, 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

⁴⁶ 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 five 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. One African-American male name (Tremayne) had too few births with available education data and was therefore dropped from this analysis. Our results are quali-

tatively similar when we use a larger data set of California births for the years 1989 to 2000 (kindly provided to us by Steven Levitt).

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).⁴⁷

The row labeled “Average” reports 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” reports 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 to 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 ($p = 0.120$). In summary, this test suggests little evidence that social background drives the measured race gap.

Names might also influence our results through familiarity. One could argue that the African-American names used in the experiment simply appear odd to human resource managers and that any odd name is discriminated against. But as noted earlier, the names we have selected are not particularly uncommon among African-Americans (see Appendix Table A1). We have also performed a similar exercise to that of Table 8 and measured the rank-order correlation between name-specific callback rates and name frequency within each gender-race group. We found no systematic positive correlation.

There is one final potential confound to our results. Perhaps what appears as a bias against African-Americans 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 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 discrimination, or at least a smaller racial gap. However, as we discussed in Section III, subsection D, we do not find any such evidence. The racial gap 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 that employers significantly favor applicants with White names.⁴⁸

⁴⁷ 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.

⁴⁸ 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. The evidence on exit out of unemployment does not paint a clear picture in this direction (Holzer, 1987).

C. Relation to Existing Theories

What do these results imply for existing models of discrimination? Economic theories of discrimination can be classified into two main categories: taste-based and statistical discrimination models.⁴⁹ Both sets of models can obviously “explain” our average racial gap in callbacks. But can these models explain our other findings? More specifically, we discuss the relevance of these models with a focus on two of the facts that have been uncovered in this paper: (i) the lower returns to credentials for African-Americans; (ii) the relative uniformity of the race gap across occupations, job requirements and, to a lesser extent, employer characteristics and industries.

Taste-based models (Gary S. Becker, 1961) differ in whose prejudiced “tastes” they emphasize: customers, coworkers, 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 vary quite a lot across these categories. We do not find a larger racial gap among jobs that explicitly require “communication skills” and jobs for which we expect either customer or coworker contacts to be higher (retail sales for example).

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 employers located in more African-American neighborhoods appear to discriminate somewhat 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.⁵⁰

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 (observable) race to proxy for *unobservable* skills (e.g., Edmund S. Phelps, 1972; Kenneth J. Arrow, 1973). This class of models struggle to explain the credentials effect as well. 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 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 (Dennis J. Aigner and Glenn G. Cain, 1977; Shelly J. Lundberg and Richard Startz, 1983; Bradford Cornell and Ivo Welch, 1996). Under such models, African-Americans 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 unclear 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 more, not less, difficult for African-Americans.

The uniformity of the racial gap across occupations is also troubling for a statistical discrimination interpretation. Numerous factors that should affect the level of statistical discrimination, such as the importance of unobservable skills, the observability of qualifications, the precision of observable skills and the ease of

⁴⁹ Darity and Mason (1998) provide a more thorough review of a variety of economic theories of discrimination.

⁵⁰ 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.

performance measurement, may vary quite a lot across occupations.

This discussion suggests 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 the race gap since the screening process (i.e., looking through a large set of resumes) may be quite similar across the variety of jobs we study.⁵¹

⁵¹ 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 Melinda Jones (2002) for the relevant psychology and Mullainathan (2003) for a formalization of the categorization concept.

V. Conclusion

This paper suggests that African-Americans face differential treatment when searching for jobs and this may still be a factor in why they do poorly in the labor market. Job applicants with African-American names get far fewer callbacks for each resume they send out. Equally importantly, applicants with African-American names find it hard to overcome this hurdle 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 racial gap in labor market outcomes. For training to work, some general-equilibrium force outside the context of our experiment would have to be at play. In fact, if African-Americans recognize how employers reward their skills, they may rationally be less willing than Whites to even participate in these programs.

TABLE A1—FIRST NAMES USED IN EXPERIMENT

White female			African-American female		
Name	L(W)/L(B)	Perception White	Name	L(B)/L(W)	Perception Black
Allison	∞	0.926	Aisha	209	0.97
Anne	∞	0.962	Ebony	∞	0.9
Carrie	∞	0.923	Keisha	116	0.93
Emily	∞	0.925	Kenya	∞	0.967
Jill	∞	0.889	Lakisha	∞	0.967
Laurie	∞	0.963	Latonya	∞	1
Kristen	∞	0.963	Latoya	∞	1
Meredith	∞	0.926	Tamika	284	1
Sarah	∞	0.852	Tanisha	∞	1
Fraction of all births:			Fraction of all births:		
3.8 percent			7.1 percent		

White male			African-American male		
Name	L(W)/L(B)	Perception White	Name	L(B)/L(W)	Perception Black
Brad	∞	1	Darnell	∞	0.967
Brendan	∞	0.667	Hakim		0.933
Geoffrey	∞	0.731	Jamal	257	0.967
Greg	∞	1	Jermaine	90.5	1
Brett	∞	0.923	Kareem	∞	0.967
Jay	∞	0.926	Leroy	44.5	0.933
Matthew	∞	0.888	Rasheed	∞	0.931
Neil	∞	0.654	Tremayne	∞	0.897
Todd	∞	0.926	Tyrone	62.5	0.900
Fraction of all births:			Fraction of all births:		
1.7 percent			3.1 percent		

Notes: This table tabulates the different first names used in the experiment and their identifiability. The first column reports the likelihood that a baby born with that name (in Massachusetts between 1974 and 1979) is White (or African-American) relative to the likelihood that it is African-American (White). The second column reports the probability that the name was picked as White (or African-American) in an independent field survey of people. The last row for each group of names shows the proportion of all births in that race group that these names account for.

REFERENCES

- Aigner, Dennis J. and Cain, Glenn G.** “Statistical Theories of Discrimination in Labor Markets.” *Industrial and Labor Relations Review*, January 1977, 30(1), pp. 175–87.
- Altonji, Joseph G. and Blank, Rebecca M.** “Race and Gender in the Labor Market,” in Orley Ashenfelter and David Card, eds., *Handbook of labor economics*, Vol. 30. Amsterdam: North-Holland, 1999, pp. 3143–259.
- Arrow, Kenneth, J.** “The Theory of Discrimination,” in Orley Ashenfelter and Albert Rees, eds., *Discrimination in labor markets*. Princeton, NJ: Princeton University Press, 1973, pp. 3–33.
- . “What Has Economics to Say about Racial Discrimination?” *Journal of Economic Perspectives*, Spring 1998, 12(2), pp. 91–100.
- Becker, Gary S.** *The economics of discrimination*, 2nd Ed. Chicago: University of Chicago Press, 1961.
- Brown, Colin and Gay, Pat.** *Racial discrimination 17 years after the act*. London: Policy Studies Institute, 1985.
- Cornell, Bradford and Welch, Ivo.** “Culture, Information, and Screening Discrimination.” *Journal of Political Economy*, June 1996, 104(3), pp. 542–71.
- 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>.
- Cross, Harry; Kenney, Genevieve; Mell, Jane and Zimmerman, Wendy.** *Employer hiring practices: Differential treatment of Hispanic and Anglo job applicants*. Washington, DC: Urban Institute Press, 1990.