Neighborhood Affluence or Long Commutes: Testing Why Employers Discriminate Against Applicants from Poor Neighborhoods Using an Audit Experiment

David C. Phillips[1]
Hope College

November 2014

Abstract

I use an audit study of the low-wage labor market in Washington, DC to test whether employers discriminate against applicants who live further from the job location. I find that fictional résumés randomly assigned to have addresses far from the job location receive 14% fewer callbacks than nearby addresses that are on average 2.6 miles closer. This effect is economically large; the measured distance penalty in callback rates equals 40% of the penalty experienced by applicants with stereotypically black names. On the other hand, evidence that employers respond to neighborhood affluence is mixed. The results have two major implications. First, previously documented discrimination against applicants from less affluent neighborhoods can mostly be accounted for by the fact that poor neighborhoods tend to be far from jobs. Second, the results provide evidence for one mechanism by which urban labor markets may exhibit spatial mismatch effects. Employer discrimination by commute distance will reduce economic prospects of everyone in a neighborhood, leading to concentrated poverty in locations far from jobs.

Keywords: employer discrimination, spatial mismatch, urban poverty, audit study

---

[1] E-mail: phillipsd@hope.edu. This project would not be possible without excellent research assistance from Alex Belica, Josh Coulter, Hayden Davis, and Brennan Mange. Funding for this project was made possible in part by grant number 1H79AE000100-1 to the UC Davis Center for Poverty Research from the U.S. Department of Health and Human Services, Office of the Assistant Secretary for Planning and Analysis (ASPE), which was awarded by the Substance Abuse and Mental Health Services Administration (SAMHSA). The views expressed are those of the authors and do not necessarily reflect the official policies of the Department of Health and Human Services. Funds from Hope College also supported this project. All remaining errors and omissions remain the responsibility of the author.

1. Introduction

The urban poor tend to be concentrated in a small number of neighborhoods. American Community Survey data compiled by Kneebone (2014) indicate that 23% of poor people in US cities live in census tracts with poverty rates over 40%, and 71% of the urban poor live in census tracts with poverty rates over 20%. Furthermore, recent research indicates that intergenerational economic mobility correlates negatively and strongly with residential segregation across U.S. cities (Chetty, et. al., 2014). Unfortunately, discrimination by employers could reinforce such concentrated urban poverty. If employers discriminate against job applicants from poor neighborhoods (Bertrand and Mullainathan, 2004), then living in a poor neighborhood could directly hinder job prospects, making it difficult to escape poverty.

Economic theory provides several explanations for why employers might discriminate in this manner. Employers may statistically discriminate against applicants from poor neighborhoods because this residential location signals low average productivity. On the other hand, poor neighborhoods tend to be geographically distant from job locations, and employers may wish to avoid workers with long commutes that sap productivity. The mechanism driving employers' discrimination matters for public policy. Statistical discrimination against fixed neighborhood attributes may be countered by promoting mixed-income neighborhoods. Concerns about long commutes require moving workers closer to jobs or improving public transit. While the existing literature provides some evidence that employers respond to a job applicant's residential address, it provides little guidance on why.

In the present study, I confirm that employers discriminate according to the residential location of job applicants. I find that concerns regarding commute distance, rather than neighborhood poverty, drive most of such discrimination. I use the standard audit study methodology[2] to examine the low-wage labor market in Washington, DC, sending fictional résumés to actual job vacancies. Résumés listing residential addresses far from the job location receive 14 percent fewer callbacks than résumés listing addresses in nearby neighborhoods with similar levels of affluence (i.e. income, education, and fraction white). For addresses on average only 2.6 miles further from the job location, this difference represents an economically large effect. I can also measure the standard racial discount in callback rates and find that living 2.6 miles further from the job reduces callback rates by 40% of the discount faced by an applicant with a stereotypically black name. I also correct for imperfect matching on affluence using applicant address fixed effects, and this correction indicates that the above results are, if anything, conservative estimates. I also measure lower callback rates for addresses in less affluent neighborhoods with the same commute distance but find smaller and statistically insignificant effects. The evidence indicates that employers discriminate according to an applicant's listed residential address, and commuting distance drives most of the observed effect.

Most narrowly, these results provide evidence that employers discriminate against applicants from poor neighborhoods, not necessarily because of the neighborhood's poverty, but because poor neighborhoods tend to be far from the job. Such evidence resolves a tension between existing audit studies of the labor market. Bertrand and Mullainathan (2004) find strong evidence that employer response rates correlate with the neighborhood affluence of addresses randomly assigned to fictional résumés in Boston and Chicago. However, Tunstall, et. al. (2013) find that addresses from poor neighborhoods in the UK receive similar callback rates as

[2] This well-established method has been used to study labor market discrimination by many different factors including but not limited to race (Bertrand and Mullainathan, 2004; Arceo-Gomex and Campos-Vazquez, 2014), immigrant status (Oreopoulos, 2011), unemployment duration (Eriksson and Rooth, 2014; Kroft, et. al. 2013), and age (Lahey, 2006).

addresses in “bland” neighborhoods.[3] The key difference between these two studies regards how they treat the distance of the applicant’s address to the job. While Tunstall, et. al. (2013) send applications that are matched to have similar commute distances, Bertrand and Mullainathan (2004) do not control for distance. By varying commuting distance and neighborhood affluence separately, the present study provides an explanation: employers may only respond to commuting distance and not neighborhood affluence. When not controlling for distance, employers will appear to respond to affluence because commuting distance and affluence are correlated in many cities. Thus, concerns regarding commuting distance drive at least a portion of employer discrimination against applicants from poor neighborhoods.

More broadly, the present results provide evidence for one mechanism behind spatial mismatch in the low-wage labor market. The spatial mismatch hypothesis (Kain, 1968) contends that living in a neighborhood far from employment opportunities will harm job prospects, leading to concentrated urban poverty. However, the large scale Moving to Opportunity project surprisingly found that moving public housing residents from high-poverty neighborhoods to low poverty neighborhoods provided no improvement in their employment prospects (Kling, et. al., 2007; Ludwig, et. al. 2012). A large debate has ensued regarding whether spatial mismatch does not matter for explaining urban poverty or whether the housing the vouchers used in MTO were ineffective in addressing spatial mismatch effects (e.g. Quigley, et. al., 2008; Aliprantis and Richter, 2012). One useful response to this debate is to drill deeper, testing whether any of the many potential mechanisms (Gobillon, et. al. 2007) that could drive a spatial mismatch effect are operable. The present results indicate that employer discrimination represents one such mechanism. If the housing market tends to sort poor and/or minority individuals into

[3] Duguet, et. al. (2010) also find no effect of a the listed town of residence for fictional applicants to French accounting positions; however, their focus on a skilled labor market makes direct comparison difficult.

neighborhoods far from available jobs, then employer discrimination against distant applicants will lead to persistent poverty and persistent ethnic differences in labor market outcomes. Thus, documented aspects of the housing market such as racial discrimination in rental housing (Ewens, et. al., 2014) can translate into gaps in labor market prospects because employers discriminate against applicants from distant neighborhoods.

In the remainder of the paper, section 2 provides some background on the geography of employment in Washington, DC and the relevant literature on spatial mismatch and employer discrimination. Section 3 describes the design of the experiment, and section 4 presents the results of the experiment. Finally, section 5 provides concluding remarks.

2. Background

2.1. Context

Neighborhood poverty correlates strongly with geographic access to jobs in Washington, DC, as in many other cities. Figure 1 displays the poverty rate for different zip codes across the city. The red outline displays the outline of the city itself with the Virginia suburbs beyond the Potomac River to the South and Maryland suburbs to the North and East. Poverty rates display a strong tendency to increase as one travels south and east, as evidenced by the darker shades on those zip codes. The three zip codes just inside the southeast boundary of the city coincide with the part of the city beyond the Anacostia River. As displayed, this area exhibits the city's highest poverty rates.

The same neighborhoods tend to be distant from job locations as well. Figure 2 displays a heat map of job locations for the jobs applied to in this experiment (see below for definition of sample). Jobs cluster downtown as evidence by the large dark circles in this area. Notably, few firms locate jobs east of the Anacostia River. Thus, these areas remain both high poverty and distant from job vacancies. In Figure 3, I summarize this relationship for all census tracts in the city. The data exhibit a strong negative relationship between the fraction of residents with at least a bachelor's degree and the average distance from the tract centroid to the jobs in my experimental sample. On average, being one mile further from the average job is associated with 13 percentage points fewer people with college degrees. Similar results obtain for tract median incomes or fraction white.

2.2. Spatial Mismatch: The Job Applicant's Decision

The observed negative correlation between distance to employment and neighborhood affluence motivates the spatial mismatch hypothesis. Kain (1968) and Wilson (1996) argue that concentrated poverty directly results from living in neighborhoods that are geographically far from job vacancies. A large empirical literature debates whether and in what contexts spatial mismatch effects actually exist. The large scale Moving To Opportunity experiment found no effect of housing moves on employment (Ludwig, et. al. 2012) while many studies with observational data find negative effects of spatial mismatch (Aslund, et. al. 2010; Sanchis-Guarner, 2014; Andersson, et. al., 2014). Others argue that spatial mismatch effects only matter in interaction with race (Hellerstein, et. al. 2008).

Given disagreement regarding the overall effects of spatial mismatch, one useful path forward is to more directly test whether potential mechanisms behind spatial mismatch are operational. Many different mechanisms could generate such an effect (Gobillon, et. al. 2007). From the point of view of the worker, job search could be less effective in distant locations due to transportation costs, lack of information, or more limited job search networks. The ineffectiveness of search could then lead those living in distant neighborhoods to reduce their search intensity. Phillips (2014) and Franklin (2014) both find evidence in field experiments that subsidizing transportation costs can induce greater search intensity for those living in neighborhoods far from jobs. Standard search models also predict that workers will require higher wages to work further from home (Zenou, 2009), and empirical evidence from firm relocations supports this idea (Mulalic, Van Ommeren, and Pilegaard, 2014). If the wage is inflexible job applicants may also reject job offers far from home, never apply in the first place, or quit jobs when the location changes because commuting costs erode net take-home pay. For instance, Zax and Kain (1996) find that firms moving to the suburbs tend to lose black employees. All of these mechanisms match the general empirical finding that workers tend to search for jobs close to home (Manning and Petrongolo, 2013; Marinescu and Rathelot, 2013), and this is especially true for poor minority workers (Holzer and Reaser, 2000). Thus, living in a neighborhood far from job vacancies could limit job prospects by affecting the job applicant's behavior.

2.2. Spatial Mismatch: Employer Discrimination

On the other hand, employer behavior could also generate spatial mismatch effects if employers discriminate based on the residential location of the job applicant (Zenou and Boccard, 2000). I will focus on testing this mechanism. Bertrand and Mullainathan (2004) send matched fictional résumés to real jobs and find that employers are less likely to call back applicants who list addresses in neighborhoods that are less wealthy/educated/white. Employers may not care about neighborhood attributes per se but still engage in statistical discrimination by poverty or any other fixed neighborhood attribute that can be extracted from a residential address. They may use neighborhood poverty to proxy average productivity differences in workers across neighborhoods (Phelps, 1972). However, as documented in the previous section, commuting distance and neighborhood affluence are strongly correlated. Observed employer discrimination against applicants from poor neighborhoods could result from employers discriminating against applicants who live far away from the particular job in question rather than discrimination based on any fixed neighborhood attribute. Supporting this theory, Tunstall et. al. (2013) find no difference in callback rates for job applicants listing addresses in neighborhoods w203ith differing levels of poverty but the same distance from the job.

Employers may wish to account for commuting distance when making hiring decisions for various reasons. First, employers may be concerned that long commutes directly decrease productivity due to fatigue or unreliability of public transit systems (Zenou, 2002). On the other hand, employers may be aware of the effect of long commutes on the employee’s behavior, leading to concerns that applicants with high commuting costs will not attend an interview, not accept the job, or quit the job in the future. Both of these mechanisms could generate discrimination by employers based on commute distance. However, these two mechanisms should have differing effects on observably more productive versus less productive applicants. In an environment with a binding minimum wage, employer concerns about direct productivity effects should fall most severely on low quality applicants for whom transit-related productivity losses causes their productivity to fall below the minimum threshold required to be hired. On the other hand, attractive applicants with better outside options should face greater distance-related discrimination if employers are concerned about distant applicants quickly quitting in favor of a new job.

In sum, employers may wish to discriminate according to an applicant’s residential location for many different reasons. They may wish to statistically discriminate against applicants from poor neighborhoods, or they may wish to discriminate against applicants who live far from the job. Economic theories of discrimination provide ample justification for either possibility. Because distance to employment tends to negatively correlate with neighborhood affluence, either of these mechanisms will lead to employers calling back applicants from poor neighborhoods at lower rates. In the present study, I undertake an experiment to disentangle these to effects and determine why employers discriminate by residential location.

3. Experimental Design

I use a correspondence audit experiment in the pattern of Bertrand and Mullainathan (2004) to study employer discrimination by residential location. From May 2014 through August 2014, I send fictional résumés to actual jobs. I independently and randomly assign different characteristics listed on the fictional résumés. Since the experiment can control and randomly assign all information observed by the employer, any correlation of employers' responses with résumé characteristics can be attributed to employer discrimination based on that attribute. I measure employer responses using e-mail and voicemail accounts according to the information listed on the job applications. I record whether employers positively respond to the application; the vast majority of positive responses are requests to setup interview times, requests for specific information about the applicant, or general requests to call back. Henceforth, these will all be generally referred to as "callbacks" and "responses." I do not include negative responses (e.g. rejection e-mails) or automated messages in this measure. I can then interpret differences in callback rates as employer discrimination.

3.1. Treatment

I focus on the address listed at the top of the résumé. The natural occurrence of such addresses on résumés provides a straightforward way to manipulate employer perceptions of the applicant's residential location. Importantly, the residential address provides information to the employer regarding both the affluence of the applicant's neighborhood and the applicant's commute distance. As noted above, these two characteristics tend to be correlated with each other such that employer discrimination based on one cannot be, in general, disentangled from discrimination based on the other. Thus, I adopt a 2x2 research design to separately vary neighborhood affluence and commute distance. I randomly assign each job application to have an address in one of four categories: near and poor (NP), near and affluent (NA), far and poor (FP), or far and affluent (FA). Details for how these addresses are chosen can be found in the Appendix. Figure 4 displays this strategy graphically. I set addresses so that NA and NP are the same distance from the job. Comparing callback rates for such addresses allows me to measure the effect of neighborhood affluence separately from commuting distance. The same holds for types FA and FP, and greater statistical precision can be had by pooling NA and FA types and comparing to NP and FP types. Likewise, distance effects can be measured by comparing callback rates for types NP and FP, which are both addresses in poor neighborhoods but differ in their distance to the job site; likewise for types NA and FA.

The first panel of Table 1 quantifies how the four types of addresses differ. The columns show average characteristics for all four address types. For instance, fictional applicants from NA addresses live on average 3.0 miles from the jobs to which they apply. NP addresses are also 3.0 miles away while FA and FP addresses are in fact further away at 5.3 miles and 5.8 miles. The final two columns measure the pooled differences between treatment types. The remaining rows display similar results for variables related to neighborhood affluence. The results indicate that the chosen addresses do generate significant variation in both distance and affluence that matches their assigned treatments. Far addresses are 2.6 miles further away from jobs than near addresses, and poor addresses are in neighborhoods with $74,000 lower median income, 50 percentage points fewer college graduates, and 40 percentage points fewer whites.

The research team conducted a small-scale public survey in Washington, DC to confirm that such variation in actual attributes of addresses does lead to perceived differences. A sample of 52 individuals were each presented with 2 addresses in Washington, DC and prompted to respond to a series of questions regarding their characteristics. Respondents demonstrated knowledge of both location and affluence subject to reasonable noise. Travel time (p-value = 0.02), neighborhood median income (0.23), fraction college educated (0.03), and fraction white (0.01) all correlate positively with actual values.[4] Combined with the documented variation in actual commuting distance and neighborhood affluence, these survey results allow us to reasonably conclude that the experiment shifts perceptions of hiring managers observing résumés.

Table 1 also demonstrates that the addresses are not perfectly matched. For instance, addresses classified as poor versus near should be the same distance from jobs. In fact, poor addresses are 0.2 miles further away. Similarly, far and near addresses should have similar affluence. While this is true for median income, far addresses tend to be less educated and less white. These remaining differences occur because the available variation in these variables does not always allow for a perfect match to be made (see Appendix for details). The matching process does, though, significantly reduce the correlation between commuting distance and neighborhood affluence. Median income and distance to the job are no longer correlated. Even when correlation between distance and measures of affluence remain, it has been reduced. Recall from Figure 3 that in a representative sample of addresses, an address one mile further from the average job tends to be in a neighborhood with 13 percentage points fewer people with college degrees. In the experimental sample this falls to 3 percentage points per mile $\left(\frac{9}{2.6}\right)$.

[4]More detailed results available upon request.

Thus, the confounding relationship between distance and educational attainment has been deflated by at least one quarter. In any case, these small differences in address characteristics that do not vary from application to application can be absorbed by including applicant address fixed effects.

3.2. Designing Fictional Job Applications

The research team composes fictional job applications in a manner similar to previous studies (e.g. Bertrand and Mullainathan 2004; Lahey 2008; Oreopoulos, 2011). A detailed experimental protocol defines the process by which research assistants apply to jobs. The overarching goal of the process is twofold. First, when possible I keep the process similar to previous audit studies of the labor market. Second, I tailor the process to studying the labor market for low-wage work by applying to different job categories and only jobs with lower skill requirements than previous studies. I generate fictional applicants with only high school education and do not apply to jobs requiring more than high school.

In particular, eight different job low-wage job categories (administrative assistant, cook, fast food, janitor, building maintenance, retail, server, and valet driver) are randomly distributed to the different research assistants and randomly ordered. Each research assistant identifies the most recent advertisement in their first assigned category on a popular website for posting job vacancies. Jobs must be located within the District of Columbia (not the suburbs), must request an e-mailed résumé or online application (not in-person application), must have an identifiable location, must not require more than high school education, and must not have been the subject of an application within the previous two weeks. If no new appropriate jobs have been posted in the job category, the research assistant moves onto their next category. Each research assistant continues through their list until meeting a daily quota of 2-4 new jobs. Using different job categories and a lower level of education leads to a pool of jobs substantially different from previous studies. Even in situations when the job categories of the present study overlap with previous studies (e.g. retail and administrative), the education requirement leads my team to apply to a different subset of such jobs. Fitting an urban poverty research question, I focus on jobs requiring limited formal education.

Once a job vacancy has been identified, the research assistant sends four separate applications to the job with at least one hour between each application. The four fictional applications include one of each address type (NP, FP, FA, NA) with specific addresses chosen according to the computerized algorithm described in the Appendix and the sending order of the applications sorted randomly. Research assistants insert the four addresses into four different résumé templates drawn from online databases of job applicants and a local employment agency in DC. Occasionally, errors in entering the inputs of the address selection algorithm result in incorrect address assignment. However, since address selection was completed correctly for 98% of applications, I will measure intent-to-treat effects using the intended address type.

The templates also require applicant names, phone numbers, e-mail addresses, prior employment information, and education information. Listed applicant names fall in three categories: white, black, or ambiguous. In each category there are male and female names. Each job vacancy receives applications evenly split between male and female. Each vacancy receives one name from all three racial categories with the fourth randomly selected from white or black. The specific names I use are the same first names Bertrand and Mullainathan (2004) use to indicate stereotypically white or black first names. Ambiguous first names were drawn using data on baby names in New York City (NYC Open Data, 2009) and chosen to be common (at least 1,000 babies per year) and have as close to equal distribution as possible between black and white. White last names come from Bertrand and Mullainathan (2004) as do most black last names. Since they use fewer black last names, I supplement their last names list to include a few more last names that have the highest proportion of black to white with at least 160,000 people having the last name. Similarly, ambiguous last names are chosen to have at least 160,000 people having the name and a black to white ratio of close to 1 to 2.[5] Finally, I randomly assign first names to last names within the same ethnic group. E-mail addresses then correspond to the name on the résumé, and phone numbers are matched to a voicemail box with a generic message recorded by a person of the appropriate sex. I use eight voicemail boxes in total so that an application can be matched both by the sex of the applicant and by each of the four address types. This ensures that all callbacks will be matched to the appropriate address type.

I design prior employment information to fit the low-wage jobs that are the subject of this study but also to indicate highly qualified applicants who should receive non-negligible callback rates. For each job category, the research team designed four separate work history profiles which are randomly assigned to the four different applications. Following the previous literature, we drew actual work histories from an online job applicant database (Indeed.com) from cities other than Washington, DC. Work histories were selected to include positive features such as experience relevant to the job category, promotion within the same organization, and increasing level of responsibility. We modified these work histories if necessary to reflect actual employers in Washington, DC and sometimes shortened job responsibility descriptions to fit our

[5] Because 1 to 1 essentially gives the list of black last names. Altogether the names are: Black male: Tremayne Jones, Leroy Thomas, Rasheed Jackson, Jamal Coleman, Kareem Robinson, Darnell Washington, Hakim Harris, Jermaine James, and Tyrone Williams. Black female: Aisha Washington, Ebony Jackson, Keisha Robinson, Kenya James, Lakisha Harris, Latonya Thomas, Latoya Williams, Tamika Jones, and Tanisha Coleman. White male: Geoffrey Kelly, Jay Sullivan, Neil Baker, Todd O’Brien, Brett McCarthy, Brendan Murphy, Matthew Ryan, Brad Walsh, and Greg Murray. White female: Allison Sullivan, Anne Walsh, Carrie Ryan, Emily Murray, Jill Murphy, Laurie McCarthy, Kristen Kelly, Meredith O’Brien, and Sarah Baker. Ambiguous male: Tyler Richardson, Jason Brooks, Eric Scott, Antonio Sanders, Raymond Bell, Brian Mitchell, Richard Ford, Joel Butler, Kyle Davis. Ambiguous female: Alyssa Richardson, Ashley Brooks, Danielle Scott, Amanda Sanders, Morgan Bell, Brianna Mitchell, Erin Ford, Christina Butler, and Paige Davis.

four templates. Work dates were chosen at random. First, the end date of the most recent job was determined by randomly drawing a current ongoing unemployment duration of zero to six months from a uniform distribution.[6] Then, the applicant shows continuous employment over three separate jobs. The length of each job was set to be at least 6 months and then randomly drawn from the empirical distribution of job lengths of the group of low-wage job seekers in the sample of Phillips (2014).

I set education information to fit the low wage labor market. In particular, all résumés list only high school graduation. This differs significantly from the previous literature, which studies college graduates or applicants with some college. I do, though, list high schools that signal high quality by selecting four schools from local parochial schools and public magnet schools and randomly assigning these to each application. I also list a GPA selected from a random uniform distribution from 3 to 4 when the template requires it. The date of graduation communicates information about age, and I select it at random to match the distribution of ages in Phillips (2014). If the graduation date and work history conflict such that the person would be working as a child, I truncate the work history at age 16.

This process for setting names, contact information, work history, and education history encapsulates all information displayed on the fictional résumés. Beyond this information, jobs requesting an e-mailed résumé also require a cover letter. We compose four standard cover letters based on publicly available templates and randomly assign these to job applications. Some job vacancies require more extensive online applications asking further information. To meet this need, each work history profile also includes wage information (based on estimates from glassdoor.com) and reasons for leaving each job. Each fictional applicant is also assigned

[6] Previous literature indicates this is the relevant range in which variation in work gaps matter (Eriksson and Rooth, 2014; Kroft, et. al., 2013)

three references from the experiment names not used for other applicants to the job. These applications often require either an IQ/skills test or personality questionnaire which is completed by the research assistants in a manner communicating a high quality applicant (i.e. to the best of their ability). For any other question idiosyncratic to the specific job application, the research assistant composes four different answers and randomly assigns them to the different applications.

Altogether, the research team sent 2,260 fictional applications to 565 job vacancies.[7] The final two panels of Table 1 present summary statistics of the various résumé characteristics as well as their balance across address treatment types. Panel B shows characteristics of the job location's census tract. Jobs tend to be in high-income, well-educated, and white neighborhoods near downtown. These variables are perfectly balanced by construction because I stratify the address treatments by job vacancy. The typical fictional applicant has graduated from high school, is 41 years old, has been unemployed for 3 months, and has 8 years of listed work experience. The sample is evenly split between male and female; 25% have ambiguous names with the remainder split evenly between black and white names.

As expected, most applicant characteristics show differences that are small both economically and statistically. All employer characteristics, having a white name, age, work experience, and sex are all statistically balanced. By chance, résumés with the "far" treatment are 6 percentage points more likely to have black names and have work gaps that are 5 days longer. These differences are statistically significant though economically small and ultimately not of major concern. One might be concerned that this imbalance could lead to lower callback rates for far addresses, leading to an overestimate of the effect of discrimination by commuting

[7] This value was chosen based on ex-ante power calculations taking into account the 2x2 research design and previous effect sizes in the literature.

distance. However, controlling for these characteristics does not change the main results (see results below). Additionally, it appears that randomly high values of these "negative" characteristics are counterbalanced by other factors. I measure overall quality of all the applicant characteristics on the résumé by regressing a callback dummy on the listed applicant characteristics and a set of 32 dummies for the interaction of the 4 different job experience profiles with the 8 different job types. The fitted values of this regression measure the overall quality of non-address related characteristics on the résumé. The final row of Table 1 displays balance on this measure of overall résumé quality. All four types of addresses have predicted callback rates between 18.3% and 19.2% based on observable characteristics, and the difference between near and far addresses is statistically insignificant and small. Randomization of résumé characteristics ensures that résumés in different treatment categories are on average of similar quality, except for the listed address.

4. Results

4.1. Main Results

Table 2 provides the simplest presentation of the experimental results. As expected near affluent addresses have the highest callback rate at .207. However, near poor addresses receive only a slightly lower callback rate of .195, indicating a 1.2 percentage point decrease in callback rates for an applicant living in a neighborhood with a similar commute time but lower affluence level. The difference between far affluent and far poor addresses is smaller. Discrimination against applicants listing distant addresses appears roughly three times larger. Applicants from far affluent addresses are called back at a rate 3.0 percentage points lower than those with near affluent addresses, and applicants from far poor neighborhoods receive 2.5 percentage points fewer callbacks than applicants from near poor addresses. Thus, employment discrimination by commute distance appears to be larger than discrimination by neighborhood affluence. Finally, any interaction between the two effects seems relatively small.

Table 3 unpacks these differences more carefully using linear regression. Column (1) reports how much lower callback rates are for the different types of addresses relative to near affluent addresses. I measure these differences by regressing a callback dummy on treatment dummies with near affluent as the omitted category. Applicants from distant poor neighborhoods receive 3.9 percentage points fewer callbacks than those from near affluent neighborhoods. This difference is statistically significant at the 5% level. Thus, we can conclude that employers discriminate by address. An employer is less likely to call back a job applicant from the poor and distant neighborhoods of far Southeast DC than an applicant who lives in a nearby and affluent neighborhood. The present study confirms the result of Bertrand and Mullainathan (2004) that employers discriminate by residential location. Column (2) includes job fixed effects and applicant controls (racial name dummies, year of listed work experience, age, length of work gap, and job category-work history profile dummies). As expected, since address affluence and distance to the job are randomly assigned, this does not change the measured coefficients significantly but improves statistical precision slightly.

However, the goal remains to distinguish whether employers discriminate by commuting distance or by neighborhood affluence. The difference between far poor and near affluent addresses could include both of these effects. In principle the estimated coefficients on far affluent and near poor addresses could parse this distinction, and the point estimates suggest that distance matters more than neighborhood poverty. However, the ordering of these coefficients is not statistically significant. So, in column (3) I pool the two treatments into overall near-far and rich-poor comparisons using dummies for far addresses (FP or FA) and poor addresses (NP or

FP). The results indicate that having an address distant from the job and having an address in a less affluent neighborhood both yield lower callback rates. However, commute distance has a larger effect, a 2.7 percentage point drop in callbacks versus 1.0 percentage point for neighborhood poverty. Without any controls the effect of commute distance is statistically significant at the 5% level. As before, adding control variables does not change this main story significantly. The coefficient changes very little and the p-value drops slightly to 0.06 in column (4) when I control for both job fixed effects and applicant characteristics. The effect of having an address in a poor neighborhood remains statistically insignificant throughout. Thus, pooling the two near-far comparisons provides stronger statistical evidence that employers do discriminate by commuting distance. The evidence for discrimination by affluence remains weaker.

For comparison, I also display the coefficient on a dummy for having a black name. The coefficient of -0.060 indicates that individuals with black names receive 6.0 percentage points fewer callbacks than those with ambiguous or white names. Since there is no difference between white and ambiguous name callback rates, this can also be interpreted as the standard white/black difference. Thus, an applicant living 2.6 miles further from the job receives at least 2.4 percentage points fewer callbacks, equal to 40% of the level of discrimination against those with stereotypically black names.

4.2. Controlling for Fixed, Address-Specific Attributes

As noted above, imperfect matching neighborhood attributes across addresses could lead me to overestimate the effect of distance on callback rates. Though the experimental design mostly breaks the negative correlation between commute distance and measured neighborhood affluence, a small negative correlation remains. I can address this concern by including applicant address fixed effects in my regression specification. Address fixed effects also alleviate a second potential source of bias in the matching process. In matching addresses by neighborhood affluence, I must quantify both distance and neighborhood affluence and then match addresses according to these characteristics (see Appendix for details). Measuring and matching by distance is relatively straightforward since distance can easily be reduced to a scalar. However, matching by neighborhood affluence presents a greater challenge. Employers may discriminate based on many different correlated but distinct fixed neighborhood attributes. Holding affluence constant and measuring purely discrimination by commute distance, as in the comparison of NP and FP in Figure 4, requires holding constant all attributes of the address other than distance. In this experiment, I hold affluence constant using propensity score matching based on median income, fraction white, and fraction with at least a bachelor's degree. The resulting index provides a single measure of neighborhood affluence and weights variables according to the extent to which employers value that attribute in calling back fictional job applicants. However, this strategy assumes that employers do not discriminate between a far, poor and a near, poor address due to variables other than income, racial composition, or education level. If employers do discriminate against some other fixed neighborhood attribute which is correlated with distance, then a simple comparison of callback rates may overestimate discrimination purely against commuting distance.

Thus, the experimental design leaves open two sources of bias. First, available addresses may not allow for perfect matches. Second, matching based on affluence cannot be completed perfectly given the multi-dimensionality of affluence. However, the design of the experiment allows me to tackle both of these issues convincingly using applicant address fixed effects. Due to variation in the location of the employer, the same applicant address may sometimes be classified as “far” and other times as “near.” The fixed effect will absorb small differences in neighborhood income, education, and racial composition due to imperfect matching. More importantly, address fixed effects control for differences in unobserved aspects of affluence across different addresses. Using only variation in commute distance within the same listed address, I can measure a pure distance effect separately from all fixed attributes of the address.

Column (5) of Table 3 displays the results for the regression including address fixed affects. The results allay concerns that measured distance discrimination has been driven by imperfect matching, either due to observed or unobserved variables. The measured effect actually gets much stronger. Even with a much larger standard error, it remains statistically significant at the 5% level. Given the large standard error resulting from using address fixed effects, I conservatively reference the specifications without address fixed effects as the main results. However, the results including address fixed effects indicate that my main results if anything underestimate rather than overestimate discrimination by commute distance.

The results in column (5) for affluence provide less helpful information. While assignment to far versus near which can vary with the location of the job leading to substantial within address variation in treatment assignment, very few addresses are assigned as poor sometimes and affluent at others. Exploiting the limited within address variation in the far vs. poor treatment generates a large *positive* estimate of 0.092 but also an extremely large standard error. More useful results can be obtained by an intermediate step between no address controls and address fixed effects. Washington, DC addresses are divided into four quadrants (defined by location relative to the US Capitol Building) with differing reputations regarding neighborhood affluence. I remove the address fixed effects and instead control for quadrant fixed effects, testing how controlling for neighborhood attributes at the level of the quadrant affects the results.

As shown in column (6), discrimination by neighborhood affluence disappears. The effect of being assigned a poor address is actually positive but now with a smaller standard error. This more precise zero suggests that any discrimination by neighborhood affluence that does exist appears to be very broad according to large regions of the city. On the other hand, discrimination by distance remains negative and statistically significant. Employers do appear to discriminate against job applicants from distant neighborhoods (even conditional on quadrant of the city) while the evidence for discrimination by neighborhood affluence is both statistically weak and not evident after controlling for broad regions of the city.

4.3. Continuous Measures of Distance and Affluence

To draw a more direct comparison with the previous literature, I can replace the 2x2 design and treatment dummies with continuous measures of distance and affluence. Table 4 provides these results. Each column tests the relationship between a continuous measure of treatment and a callback dummy. Each regression includes applicant controls and job fixed effects but omits address fixed effects for the sake of statistical precision. The first column demonstrates again that employers respond negatively to distant applicants. The callback rate falls by 0.9 percentage points per mile. Extrapolating, the result implies that a person moving 6-7 miles further from the job results in a similar impediment to their job applications as having a stereotypically black name.

As expected from previous results, the median income of the census tract of the applicant's address correlates positively but not significantly with employer responses. On the other hand, education levels and racial composition of the census tract do correlate with callback rates. Because the affluence index places greater weight on median income (see Appendix), the index of neighborhood affluence less significantly correlates with callback rates. These results could provide support for the conclusion that employers do discriminate according to neighborhood affluence in a rather imprecise manner, negatively responding to applicants from the far Southeast of the city in a manner that correlates with racial and educational composition. The coefficient of 0.04 on fraction college educated indicates that switching a neighborhood with no college graduates to one with all graduates would lead to an increase in callback rates of 4 percentage points, which is of similar magnitude to what Bertrand and Mullainathan (2004) find without controlling for distance. However, these results also conflict with the conclusions I obtain using the ex-ante 2x2 research design. While evidence regarding discrimination according to commute distance remains strong, evidence regarding discrimination by neighborhood affluence is mixed.

5. Conclusion

In this study, I have demonstrated that employers discriminate against job applicants who list more distant residential addresses. When presented with otherwise similar fictional résumés, hiring managers for actual low-wage job vacancies call back applicants living further away 14 percent less often. This effect is large. Living on average 2.6 miles further away decreases callback rates by an amount equal to 40% of the discount experienced by applicants with 'black names' relative to those with 'white names.' Since commuting distance and neighborhood poverty are correlated, such discrimination by distance accounts for most of the previously documented discount in callback rates experience by applicants from poor neighborhoods. On the other hand, the evidence regarding whether the affluence of the listed address's neighborhood directly affects employer behavior is mixed.

These results provide support for the spatial mismatch hypothesis, the idea that living far from employment opportunities has a direct negative impact on labor market prospects of the urban poor. I document that living far from the job leads to lower callback rates. While I cannot observe actual employment or wage offers, a drawback of all résumé studies, standard random search models predict that a lower arrival rate of contacts with employers should result in lower employment rates and lower wages in equilibrium (McCall, 1970). In a labor market with frictions, this evidence of employer discrimination provides a causal mechanism running between living in a neighborhood far from jobs and poor labor market outcomes. While a large non-experimental literature has tested this theory, to my knowledge very few experiments have directly tested such causal mechanisms of spatial mismatch. Phillips (2014) provides evidence for a search cost mechanism on the worker side, demonstrating that subsidizing job search can increase search intensity for those living far from job vacancies. The present study confirms that a mechanism on the employer side could also contribute to spatial mismatch effects.

Understanding the mechanisms behind spatial mismatch helps to guide policy responses. The Moving to Opportunity (MTO) project prominently found that providing housing vouchers to public housing residents did not improve their labor market outcomes (Kling, et. al., 2007; Ludwig, et. al. 2012). However, while MTO participants moved to neighborhoods with lower poverty rates, Quigley, et. al. (2008) point out that the voucher recipients tended to move to neighborhoods with similarly poor geographic access to jobs. If neighborhood effects operate through spatial mechanisms relating specifically to commuting distance, then it would not be surprising for housing vouchers to demonstrate no improvement in labor market outcomes. For instance, a housing voucher to move to a less poor but equally distant neighborhood would not affect the employer discrimination observed in the present study. Eliminating the negative employment effects of living in poor neighborhoods would instead require housing interventions moving residents further from home and closer to jobs. Perhaps more practically, public transit could be improved. Employer discrimination by commuting distance exists. The present study demonstrates that it is a proven mechanism by which spatial mismatch can operate. This fact can help interpret previous attempts to address spatial mismatch and inform future public policy responses.

References

Aliprantis, D. and F. Richter (2012) “Evidence of Neighborhood Effects from MTO: LATEs of Neighborhood Quality.” *Federal Reserve Bank of Cleveland Working Paper,* No. 12-08.

Andersson, F., J.C. Haltiwanger, M.J. Kutzbach, H.O. Pollakowski, and D.H. Weinberg (2014) “Job Displacement and the Duration of Joblessness: The Role of Spatial Mismatch.” *NBER Working Paper No.* 20066.

Arceo-Gomez, E. O., and R.M. Campos-Vazquez (2014) ‘Race and Marriage in the Labor Market: A Discrimination Correspondence Study in a Developing Country.’ *The American Economic Review: Papers and Proceedings, 104*(5).

Åslund, O., Östh, J., & Zenou, Y. (2010). How important is access to jobs? Old question—improved answer. *Journal of Economic Geography*, *10*(3), 389-422.

Bertrand, M. and S. Mullainathan (2004) ‘Are Emily and Greg More Employable Than Lakisha and Jamal.” *American Economic Review,* 94(4).

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014) “Where if the Land of Opportunity? The Geography of Intergenerational Mobility in the United States.” *Quarterly Journal of Economics,* forthcoming.

Duguet, E., N. Leandri, Y. L’Horty, and P. Petit (2010) “Are Young French Jobseekers of Ethnic Immigrant Origin Discriminated Against? A Controlled Experiment in the Paris Area.” *Annals of Economics and Statistics,* No. 99/100.

Eriksson, S., & Rooth, D. O. (2014). Do employers use unemployment as a sorting criterion when hiring? Evidence from a field experiment. *The American Economic Review*, *104*(3), 1014-1039.

Ewens, M., B. Tomlin, and L.C. Wang (2014) "Statistical Discrimination or Prejudice? A Large Sample Field Experiment," *The Review of Economics and Statistics*, 96(1).

Franklin, S. (2014) “Location, Search Costs and Youth Unemployment: The Impact of a Randomized Transport Subsidy in Urban Ethiopia.” Unpublished Working Paper.

Gobillon, L., Selod, H., and Zenou, Y. (2007). The mechanisms of spatial mismatch. *Urban Studies*, *44*(12), 2401-2427.

Hellerstien, J., D. Neumark, and M. McInerney (2008) ‘Spatial Mismatch or Racial Mismatch?’ *Journal of Urban Economics,* 64(2)

Holzer, H. J., & Reaser, J. (2000). Black applicants, black employees, and urban labor market policy. *Journal of Urban Economics*, *48*(3), 365-387.

Kain, J. (1968) “Housing Segregation, Negro Employment, and Metropolitan Decentralization.” *Quarterly Journal of Economics*, 82.

Kling, J., J. Liebman, and L. Katz (2007) “Experimental Analysis of Neighborhood Effects.” *Econometrica*, 75.

Kneebone, E. (2014) “The Growth and Spread of Concentrated Poverty, 2000 to 2008-2012.” *Brookings Institution Research Brief*.

Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment. *The Quarterly Journal of Economics*, *128*(3).

Lahey, J. N. (2008). Age, Women, and Hiring An Experimental Study. *Journal of Human Resources*, *43*(1), 30-56.

Ludwig, J., et. al. (2012) ‘Neighborhood Effects on the Long-Term Well-Being of Low-Income Adults.” *Science*, 337(September 21).

Manning, A., and B. Petrongolo (2013) ‘How local are labor markets? Evidence from a spatial job search model.’ Working Paper.

Marinescu, I., Rathelot, R. (2013) ‘The geography of job search and mismatch unemployment.’ Working Paper.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 113-126.

Mulalic, I., Van Ommeren, J. N. and Pilegaard, N. (2014) “Wages and Commuting: Quasi-natural Experiments' Evidence from Firms that Relocate.” *The Economic Journal*, 124.

NYC Open Data (2009) “Most Popular Baby Names by Sex and Mother’s Ethnic Group.” https://nycopendata.socrata.com/ Accesssed: 4/24/2014

Oreopoulos, P. (2011) "Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes." *American Economic Journal: Economic Policy*, 3(4): 148-71.

Phelps, E.S. (1972) "The statistical theory of racism and sexism." *The American Economic Review*.

Phillips, D.C. (2014) “Getting to work: Experimental evidence on job search and transportation costs” *Labour Economics*, 29.

Quigley, J.M., S. Raphael, L. Sanbonmatsu, and B.A. Weinberg (2008) “Neighborhoods, Economic Self-Sufficiency, and the MTO Program” *Brookings-Wharton Papers on Urban Affairs*.

Sanchis-Guarner, R. (2012) “Driving Up Wages: The Effects of Road Construction in Great Britain.” *SERC Discussion Papers*, No. 120.

Tunstall, R., A. Green, R. Lupton, S. Watmough, and K. Bates (2014) “Does Poor Neighbourhood Reputations Create a Neighbourhood Effect on Employment? The Results of a Field Experiment in the UK.” *Urban Studies*, 51(4).

Wilson, W.J. (1997) *When Work Disappears: The World of the New Urban Poor*. Vintage Books.

Zax, J. and J. Kain (1996) “Moving to the Suburbs: Do Relocating Companies Leave Their Black Employees Behind?” *Journal of Labor Economics*, 14.

Zenou, Y. and N. Boccard (2000) “Labor discrimination and redlining in cities.” *Journal of Urban Economics*, 48.

Zenou, Y. (2002) “How do firms redline workers?” *Journal of Urban Economics*, 52.

Zenou, Y. (2009) *Urban Labor Economics*. Cambridge University Press.

Figure 1. Poverty Rates across Washington, DC Area Zip Codes

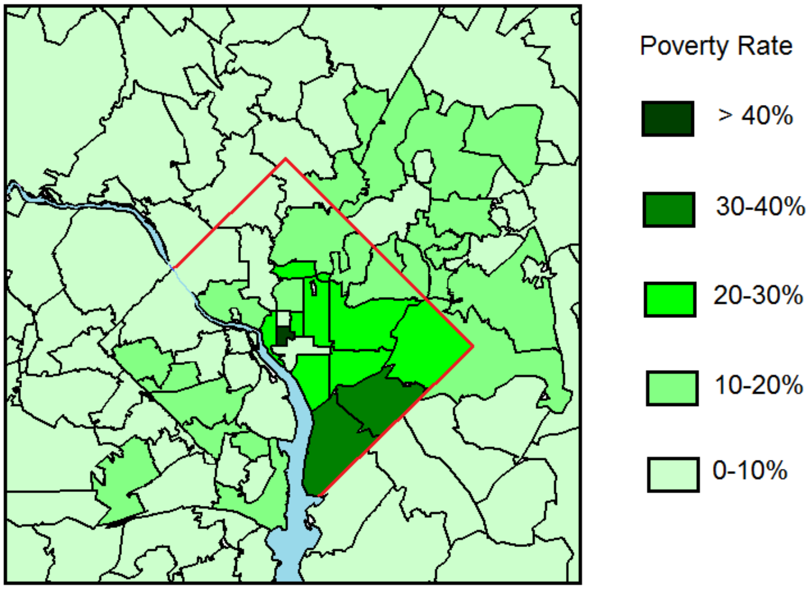

Source: US Census of Population, 2000

Figure 2. Location of Job Vacancies Included in Experiment

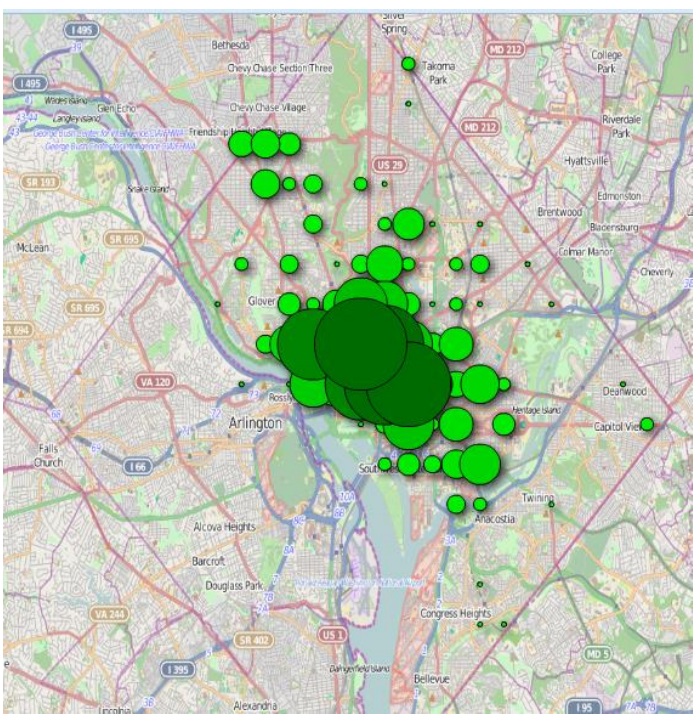

Source: Data from experiment. Larger/darker circles indicate more job vacancies in that location.

Figure 3. Job Access and Education Levels in Washington, DC Census Tracts

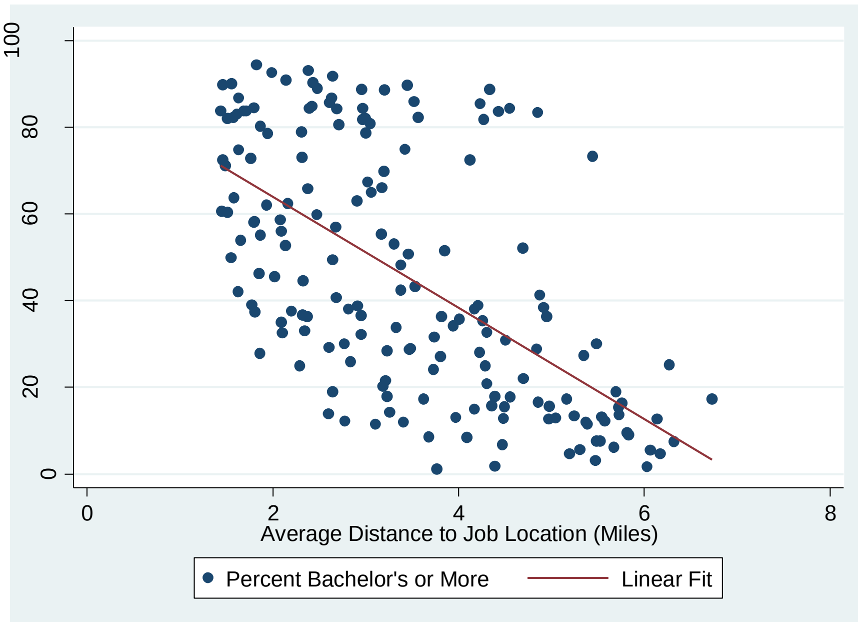

Source: Average distance to jobs is computed as average great circle distance from the tract centroid to the job vacancies used in the audit experiment. Fraction of the population with a Bachelor's or More comes from the American Community Survey 2011 5-Year estimates.

Figure 4. Identification Strategy

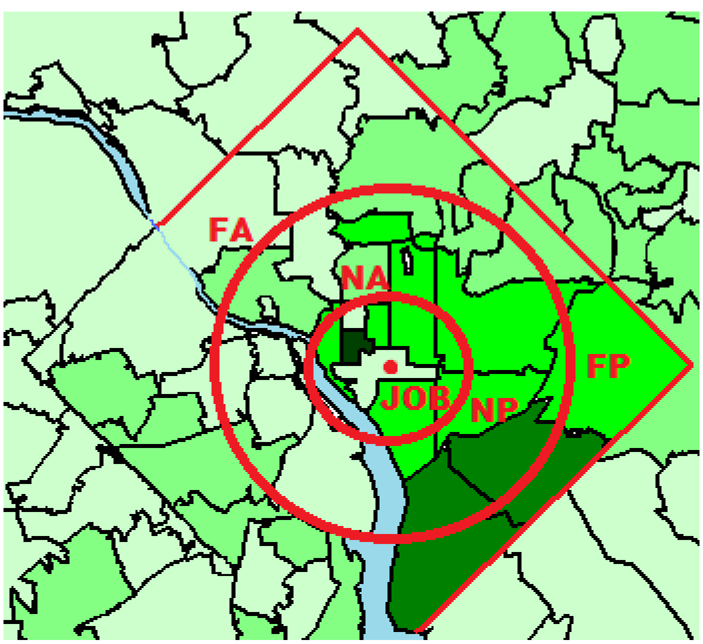

Source: US Census of Population, 2000

Table 1. Summary Statistics and Baseline Balance

	Near, Affluent	Far, Affluent	Near, Poor	Far, Poor	Poor - Affluent	Far - Near
A. Applicant Address Characteristics						
Distance to Job (miles)	3.0	5.3	3.0	5.8	0.2	2.6
					[0.001]	[0.000]
Tract Median Income ($)	101,698	106,371	32,429	27,812	-73,914	28
					[0.000]	[0.98]
Tract Percent Bachelor's or Higher	70	62	20	11	-50	-9
					[0.000]	[0.000]
Tract Percent White	60	31	9	1	-40	-19
					[0.000]	[0.000]
Tract Affluence Index	0.146	0.142	0.083	0.077	-0.064	-0.005
					[0.000]	[0.000]
B. Job Address Characteristics						
Tract Median Income ($)	95,672	95,672	95,672	95,672	0	0
					[1.00]	[1.00]
Tract Percent Bachelor's or Higher	74	74	74	74	0	0
					[1.00]	[1.00]
Tract Percent White	65	65	65	65	0	0
					[1.00]	[1.00]
C. Applicant Characteristics						
White	0.35	0.35	0.40	0.34	0.02	-0.03
					[0.42]	[0.20]
Black	0.36	0.42	0.35	0.41	-0.01	0.06
					[0.65]	[0.01]
Age	41	41	41	41	-1	0
					[0.27]	[0.56]
Work Gap (days)	88	91	89	95	2	5
					[0.31]	[0.04]
Work Experience (years)	8.1	7.7	8.0	8.2	0.2	0.0
					[0.32]	[0.82]
Female	0.49	0.49	0.54	0.48	0.02	-0.03
					[0.31]	[0.11]
Overall Quality (Predicted Callback Rate)	0.186	0.183	0.192	0.188	0.005	-0.003
					[0.053]	[0.22]
Sample Size	565	565	565	565		

The first four columns display means for each characteristic by treatment group. The final two columns measure differences in characteristics by regressing the variable of interest on a dummy variable for a poor address or a dummy variable for a far address, respectively. P-values are reported in brackets. Standard errors are clustered at the job level. The overall quality variable predicts a callback dummy using a female name dummy, racial name dummies, age, years of listed work experience, length of work gap, and job profileXjob category dummies.

Table 2. Response Rates by Treatment Type

	Affluent	Poor	Difference
Near	0.207 [565]	0.195 [565]	0.012
Far	0.177 [565]	0.170 [565]	0.007
Difference	0.030	0.025	0.005

All values are fraction of applications called back by employers; total numbers of applications sent are in brackets

Table 3. Effect of Address Category on Employer Response

Dependent Variable:	Callback Dummy					
	(1)	(2)	(3)	(4)	(5)	(6)
Far, Poor	-0.037**	-0.039**	--	--	--	--
	(0.019)	(0.018)				
	[0.05]	[0.04]				
Far, Affluent	-0.030*	-0.027	--	--	--	--
	(0.018)	(0.018)				
	[0.09]	[0.14]				
Near, Poor	-0.012	-0.018	--	--	--	--
	(0.024)	(0.018)				
	[0.60]	[0.32]				
Far	--	--	-0.027**	-0.024*	-0.074**	-0.030**
			(0.013)	(0.013)	(0.030)	(0.014)
			[0.03]	[0.06]	[0.02]	[0.03]
Poor	--	--	-0.010	-0.015	0.092	0.006
			(0.013)	(0.013)	(0.065)	(0.019)
			[0.46]	[0.24]	[0.16]	[0.74]
Black	--	-0.060***	--	-0.060***	-0.064***	-0.061***
		(0.016)		(0.016)	(0.018)	(0.016)
		[0.00]		[0.00]	[0.00]	[0.00]
Applicant Controls	N	Y	N	Y	Y	Y
Job Fixed Effects	N	Y	N	Y	Y	Y
Applicant Address Controls	N	N	N	N	Address Fixed Effects	Quadrant Dummies
Sample size	2,260	2,260	2,260	2,260	2,260	2,260

Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and * respectively. Applicant controls include a female name dummy, racial name dummies, age, years of listed work experience, length of work gap, and job profileXjob category dummies. Standard errors are clustered at the job vacancy level. Selected p-values are in brackets.

Table 4. Continuous Measures of Distance and Affluence

Dependent Variable:	Callback Dummy				
	(1)	(2)	(3)	(4)	(5)
Distance to Job (Miles)	-0.009*	--	--	--	--
	(0.005)				
	[0.06]				
Log Median Income	--	0.013	--	--	--
		(0.009)			
		[0.19]			
Fraction Bachelor's or More	--	--	0.05*	--	--
			(0.02)		
			[0.053]		
Fraction White	--	--	--	0.05**	--
				(0.02)	
				[0.03]	
Affluence Index	--	--	--	--	0.30
					(0.19)
					[0.11]
Applicant Controls	Y	Y	Y	Y	Y
Job Fixed Effects	Y	Y	Y	Y	Y
Address Controls	N	N	N	N	N
Sample size	2,260	2,260	2,260	2,260	2,260

Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and * respectively. Applicant controls include a female name dummy, racial name dummies, age, years of listed work experience, length of work gap, and job profileXjob category dummies. Standard errors are clustered at the job vacancy level. Selected p-values are in brackets.

Appendix

A.1. Measuring Distance and Affluence

The 2x2 research design described above requires measuring both commuting distance from an applicant's address to a job location and measuring an index of affluence for any address. I measure distance using great circle distance in miles. This can be easily measured by geo-coding the address of the job vacancy and the address listed on the job application. To measure affluence, I draw on publicly available data from the American Community Survey (2011 5-Year Estimates) and previous work by Bertrand and Mullainathan (2004). The challenge is to summarize all fixed (i.e. not dependent on the location of the employer, such as distance) neighborhood attributes such as poverty, racial composition and educational attainment into an index describing employer perception of that neighborhood. I use propensity-score matching techniques to this end. Using the Bertrand and Mullainathan (2004) experimental data, I can estimate the following probit regression:

$$\Pr[C_i = 1] = \Phi(\beta_0 + \beta_1 Inc_i + \beta_2 FracWhite_i + \beta_3 FracCol_i)$$

C_i is a indicator of whether applicant i received a callback; Inc_i is the log median income of the census tract of the address listed on i's résumé; $FracWhite_i$ is the fraction of census tract residents who are white; $FracCol_i$ is the fraction of the census tract with at least a bachelor's degree; $\Phi(\cdot)$ is the normal distribution. Appendix Table A.1. presents the results of estimating this equation with data from Bertrand and Mullainathan (2004) data. The three variables are jointly significant (F-test p-value of 0.005).

I extrapolate these results to the new setting in Washington, DC by combining the results with the ACS data for Washington, DC. I calculate expected callback rates for any census tract in DC as:

$$\text{Index of Affluence} = \Phi\left(\hat{\beta}_0 + \hat{\beta}_1 Inc_i + \hat{\beta}_2 FracWhite_i + \hat{\beta}_3 FracCol_i\right)$$

This is my measure of affluence. This process combines census tract income, racial composition, and educational attainment into one measure where different attributes are weighted depending on the observed importance placed on these characteristics by employers in the Bertrand and Mullainathan (2004) data. More specifically, the index is the propensity score that can then be used to match census tracts by how their characteristics are viewed by employers. Two addresses with similar propensity scores should be treated similarly by employers if neighborhood income, racial composition, and educational attainment sufficiently characterize the information contained in an address.

A.2. Choosing Addresses

To choose specific addresses, I list addresses in an 18x18 equally spaced grid with borders formed by the points of the Washington, DC diamond. I then eliminate points outside of Washington, DC. I also eliminate addresses in census tracts dominated by universities (at least 30% college students) or military bases (at least 30% in armed forces). Each remaining point on this grid is paired with an address on the nearest "main street" (defined as streets shown as white or yellow on Google Maps at a particular level of zoom). I use main streets because a public survey (described above) indicated that respondents can more accurately identify characteristics of addresses on such main streets. This alternative performed better than solely manipulating the quadrant of the address or using addresses whose locations are communicated by the alphabetical/numeric system of streets in DC. The result is a grid of addresses across Washington, DC entirely composed of addresses on main streets. The distance of each address to the location of a particular job vacancy can be measured easily. I also attach ACS data, and thus an affluence index, to each address according to its census tract.

Given the location of the job vacancy, I first define "Near, Poor" addresses by requiring that they be below the 10th percentile of the affluence index among all addresses on the grid. Then, I select addresses that are no more than 1 mile further from the job than the closest such address. From this group of potential addresses, I choose one at random. I require that the "Near, Affluent" address be the same distance from the job as "Near, Poor" ($\pm$ 0.15 $miles$) and select one address at random from those that have an index above median affluence.[8] For "Far, Poor" addresses I choose an address at random from among those that have the same affluence as the NP address ($\pm$ 0.01; or 0.3 s.d.) and are at least two miles further away from the job than the NP address.

Choosing the "Far, Affluent" address is the most difficult as it requires matching both the affluence of the NA address and the distance of the FP address. Sometimes these two goals trade off against each other. In practice, I balance these two concerns by choosing the address that minimizes the following:

$$\left(\frac{Dist_{FA}-\mu_{dist}}{\sigma_{dist}}-\frac{Dist_{FP}-\mu_{dist}}{\sigma_{dist}}\right)^2+\left(\frac{Affluence_{FA}-\mu_{aff}}{\sigma_{aff}}-\frac{Affluence_{NA}-\mu_{aff}}{\sigma_{aff}}\right)^2$$

where μ's are means, σ's are standard deviations, $Dist$ is distance to job, and $Affluence$ is the affluence index. In words, I translate measurements of the affluence index and distance into z-scores, calculate the squared difference of the FA type z-score from the one it should match (FP

---

[8] If there is no such address, I choose the most affluent address.

for distance; NA for affluence), and then add the two squared differences together. In the ideal, this calculation would result in zero, indicating that the FA matches the affluence of the NA and the distance of the FP exactly. In practice I come close to this ideal, as demonstrated in Table 1, though the tradeoff between matching distance and matching affluence leads the FA addresses to be somewhat nearer than FP addresses and somewhat less affluent than NA addresses. As discussed above, though, I can control for imperfect matching using address fixed effects.

Appendix Table 1. Probit Predicting Callback Dummy Using Bertrand and Mullainathan (2004) Data

	(1)
Tract Median Income ($)	0.17
	(0.12)
Tract Percent Bachelor's or Higher	0.25
	(0.20)
Tract Percent White	0.01
	(0.11)
Sample size	4,784

Probit index coefficients reported. Standard errors in parentheses. Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and * respectively.