

Statistical Discrimination or Prejudice? A Large Sample Field Experiment

Michael Ewens, Bryan Tomlin, and Liang Choon Wang⁺

Abstract

A model of racial discrimination provides testable implications for two features of statistical discriminators: differential treatment of signals by race and heterogeneous experience that shapes perception. We construct an experiment in the U.S. apartment rental market that distinguishes statistical discrimination from taste-based models of discrimination. Responses from over 14,000 rental inquiries with varying applicant quality show that landlords treat identical information from applicants with African-American and white sounding names differently. This differential treatment varies by neighborhood racial composition and signal type. The evidence indicates statistical discrimination by landlords and explains past findings of lower marginal return to credentials for minorities.

JEL Codes: J15, R3.

I. Introduction

Racial and ethnic discrimination continues to pervade many markets in the US. Roughly half of the annual discriminatory cases reported by federal agencies involve race or ethnicity, and the number of new incidents outpaced population growth over the past 10 years.¹ The economics literature posits two major sources of racial discrimination: taste-based and statistical. Racial prejudice produces taste-based discrimination, while statistical discrimination occurs in an environment of imperfect information where agents form expectations based on limited signals that correlate with race.² The result of both types of discrimination, however, is the same: similar individuals who differ only by their race experience different outcomes. A simple examination of differential treatment sheds little light on the source of discrimination and potentially explains why few studies are able to find conclusive evidence of statistical discrimination.

Employing an email correspondence experiment in the apartment rental market, this paper tests whether statistical discrimination explains differential treatment by race. We extend the Aigner and Cain (1977) model of statistical discrimination to provide testable implications for two features of statistical discriminators: differential treatment of signals by race and heterogeneous experience that shapes perception. The model guides our research design and isolates parameters that map to statistical discrimination. Using vacancy listings on Craigslist.org (Craigslist) across 34 U.S. cities and roughly 5,000 neighborhoods (census tracts), we send emails with two key components to 14,000 landlords. We use the common racial-sounding first names of Bertrand and Mullainathan (2004) to associate applicants with race, and the email contains differing – but limited – pieces of information: positive, negative and no signals beyond race. The dependent variable codes landlords' responses to capture an invitation to the fictional inquiry for future contact. Although the outcome reflects only a positive response during the

initial inquiry phase of a screening process, any differential treatment in screening will likely influence final outcomes in the same direction.

An ideal research environment to test for statistical discrimination has distant communication between agents with imperfect information and also avoids the confounding factors inherent in audit studies.³ Email correspondence for apartment rental inquiries via Craigslist gives such an environment. Craigslist is the dominant source of online classifieds in the U.S., especially for apartment listings, and is frequented by one-third of the white and black U.S. adult population.⁴ The website provides control of information and an ability to manipulate signals available to agents. Further, Craigslist allows us to accurately track responses and scale the experiment to thousands of heterogeneous neighborhoods. Since residential locations are closely tied to characteristics associated with welfare, such as the type of job held, crime levels, and school quality, our focus on the apartment rental market is policy relevant. The growing prevalence of online interactions in real estate, employment, lending, and auctions, suggest the results extend beyond the apartment rental market.

The experiment provides four major results. First, we present emails to landlords with racial sounding names as the only signal and confirm that applicants with African-American sounding names are 16 percent less likely to receive a positive response from a landlord than those with white sounding names. The finding conforms to a model where landlords use race to approximate tenant quality. This test, though simple, is unique in the correspondence literature and provides a base case for the model. The observed differential treatment can also be explained by racial prejudice, so we next test the additional model implications.

Second, the model posits that landlords may differ in their perceptions of signals due to past experience in the screening and rental process and in turn, incorporate race and signals into

decisions differently. To test this hypothesis, we introduce additional information in some emails. In the “positive information” inquiry, the fictional applicant says her name and informs the landlord she is a non-smoker with a respectable (and paying) job. In the “negative information” inquiry, the applicant states her name and tells the landlord she has below average credit rating and smokes. Sending negative signals may be unusual, however, applicants could find it advantageous to disclose such information upfront to avoid paying for a likely-to-fail credit check requested by most landlords or to avoid getting turned down for being a smoker after incurring the time cost of viewing the apartment. Using a difference-in-difference estimator, we show that the average landlord weights the same signal relatively more when it comes from an applicant with a white sounding name than one with an African-American sounding name. The coefficient estimate identifies the parameter that translates signal into quality assessment and potentially explains why past studies often show lower marginal return to credentials for minority groups.

Third, the model also defines a notion of “surprise,” where the base case acts as a benchmark for uninformed expectations and a means to quantify surprise relative to the better-than-expected (positive) and worse-than-expected (negative) information. This notion of surprise is particularly difficult to introduce in a job application setting where resumes are required, as it is impossible to provide “zero” information about education or experience in a resume. In the presence of differential weighting of signals by race, the model predicts that a positive surprise will not necessarily shrink the racial gap, but a negative surprise will. Our empirical results are consistent with these predictions.

Finally, we exploit neighborhood sorting to examine a source of heterogeneity consistent with the shaping of signal perception. The model shows that identification of statistical

discrimination requires finding distinct patterns in the weighting parameter (i.e., signal perception) by experience across racial groups. By allowing a signal's noise to depend on race, the model presents another testable hypothesis: a landlord's relative experience with a given race increases the relative weight she places on the signal from that group. The apartment rental market is an ideal setting for this test, since a landlord's past experience is closely tied to the neighborhood characteristics in which she is renting. We find that as the share of black residents in a neighborhood increases, a positive surprise closes the racial gap observed in the base case, while a negative surprise does little to close it.

These findings are difficult to reconcile with a simple theory of preference-based discrimination. A preference-based model would need to explain why landlords who own rental properties in predominantly white neighborhoods and exhibit distaste for minority applicants decide to treat positive information from both races equally, and yet respond more negatively to negative information presented by white applicants. A taste-based model would also need to explain the persistent gap between applicants with white and African American sounding names across all types of neighborhoods in the base case. We believe that it is difficult to build a tractable model of preference-based discrimination that is consistent with these patterns.

This paper fits into the large body of research on racial discrimination. With the exception of List (2004) and to some extent Levitt (2004), past evidence of statistical discrimination is inconclusive. For example, Altonji and Pierret (2001) and Bertrand and Mullainathan (2004) have found significant racial gaps in wages and job interview callback, respectively, but weak support for statistical discrimination.⁵ Scant evidence of statistical discrimination stems from the lack of objectively distinct information types, weak treatment effects from signals of quality, and/or the inability to identify differential perceptions in the data.

The new contribution is a research design within a difference-in-difference framework that can identify whether the observed racial gap is consistent with statistical discrimination. Our range of signals, large sample size, and diverse set of neighborhoods show that significant treatment effects and a difference-in-difference estimator are an important prerequisite to test for statistical discrimination and might explain a lack of such evidence in Bertrand and Mullainathan (2004). The evidence of statistical discrimination yields important policy implications that may differ from those used to address racial prejudice.

II. A Model of Discrimination in Screening

We extend the Aigner and Cain (1977) model of statistical discrimination to explain differential screening outcomes by race with an application to the apartment rental market. The model applies to other situations of semi-formal screening. Consider the following five-stage process of matching potential tenants to apartments:

1. A landlord posts details of an available unit on a public forum inviting inquiries.
2. Potential tenants select units to send costless inquiries to which include a signal, X , sent to the landlord. X can include both what the applicant says and any other observable features such as gender or race.
3. The landlord receives signals from potential tenants over time and uses content X to decide which individuals are qualified for (potentially costly) face-to-face interviews.
4. Applicants who pass the initial screening reveal their true quality θ during face-to-face interviews at some cost c to the landlord per interview.⁶
5. The landlord offers their unit to the best applicant after face-to-face interviews.

We focus on discrimination occurring in stage 3 of the above matching process.

To illustrate how differential outcomes by race may occur, consider the simple process of how landlords invite applicants for interviews. Landlords predict applicant quality θ using observable signals and choose to respond positively ($R = 1$ in our empirical specification) to applicants whose expected quality $\hat{\theta}$ is greater than some reservation quality, $\underline{\theta}$. Suppose that signal X proxies quality θ noisily with a race-specific error ε_r :

$$X_r = \theta_r + \varepsilon_r, \quad (1)$$

where $\theta_r \sim N(\mu_r, \sigma_\theta^2)$, $E(\varepsilon_r | \theta_r) = 0$, $\text{var}(\varepsilon_r | \theta_r) = \sigma_{\varepsilon,r}^2$, $E(X_r) = \mu_r$, and $\text{var}(X_r) = \sigma_\theta^2 + \sigma_{\varepsilon,r}^2$.⁷

Landlords have a sample of initial inquiries X and applicant qualities θ acquired during past iterations of stage 4 of the screening process outlined above. Although landlords do not know the relationship between quality θ and signal X , or their distributions, they can use past experience to form predictions about applicant quality given a signal. Since applicants may implicitly reveal other signals such as race and gender, it can be rational to use these implicit signals in forming predictions if they correlate with other unobserved indicators of quality. We assume that statistically discriminating landlords use past observations of θ and X to estimate the following forecasting regression for each race r :

$$\hat{\theta}_r = \hat{\mu}_r^L + \hat{\gamma}_r X_r, \quad (2)$$

where $\hat{\mu}_r^L$ is the Ordinary Least Squares (OLS) estimator of the intercept term; $\hat{\gamma}_r$ is the estimator of the marginal effect of signal X_r ; and r is W for whites, and B for blacks. It is sensible for a risk-averse landlord to use the OLS estimator for prediction since it minimizes the variance of the forecast errors. Because different landlords have different past experience, estimates of $\hat{\mu}_r^L$ and $\hat{\gamma}_r$ vary across landlords.

Now suppose a landlord observes a new signal \tilde{X} and race r from an applicant in stage 3. The landlord will predict quality using (2):

$$\hat{\theta}_r = \hat{\mu}_r^L + \hat{\gamma}_r \tilde{X}. \quad (3)$$

The landlord uses this estimated quality to determine whether or not they will invite the applicant for an interview and to view the apartment. Equations (1), (2), and (3) correspond to Aigner and Cain's (1977) model of statistically discriminating employers, and we may call $\hat{\gamma}_r$ the *information weighting parameter* for race r – since it informs a landlord how much to weight a signal or how to perceive a signal from an applicant of race r .⁸

It is possible for an applicant to inquire about an apartment and only reveal her race.⁹ The landlord can infer quality using the average signal (\bar{X}_r) received from race r to forecast:¹⁰

$$\hat{\theta}_r = \hat{\mu}_r^L + \hat{\gamma}_r \bar{X}_r. \quad (4)$$

Equations (2), (3), and (4) define a statistical discriminator who rationally responds to imperfect information by incorporating past experience into an OLS estimator (i.e., the best linear predictor) of applicant quality. Statistically discriminating landlords are not interested in the causal relationship between quality and signal or whether OLS estimates of $\hat{\mu}_r^L$ and $\hat{\gamma}_r$ are consistent, but are only interested in predictions that yield the lowest variance.¹¹

To see how differential outcomes by race may arise, we can focus on the OLS estimators in equation (2):

$$\hat{\gamma}_r = \frac{\text{cov}(\theta_r, X_r)}{\text{var}(X_r)} \quad (5)$$

$$\hat{\mu}_r^L = \bar{\theta}_r - \hat{\gamma}_r \bar{X}_r \quad (6)$$

Here $\hat{c}ov(\theta_r, X_r)$ is the sample covariance between quality and signal, $\hat{v}ar(X_r)$ is the sample variance of the signal, and $\bar{\theta}_r$ is the sample average of quality. Equation (5) shows that, given $\hat{c}ov(\theta_w, X_w) = \hat{c}ov(\theta_B, X_B)$, any differences in noise of signals, $\hat{v}ar(X_r)$, can induce differences in the weight a landlord places on the same signal from different races. For example, landlords with $\hat{v}ar(X_B) > \hat{v}ar(X_w)$ will have $\hat{\gamma}_B < \hat{\gamma}_w$.¹² For applicants with objectively identical signals except race, these landlords will be more likely to invite white applicants than black applicants for interviews. Note, we have not placed any restrictions on the mean signals nor noise across race. Differences in mean signal or noise could stem from fundamental racial differences in signals such as income or credit scores of the population or simply from those observed by landlords in their samples.

Of course, differential outcomes by race may also arise from prejudice. If the landlord has “distaste” for a particular racial group, then the reservation quality will be higher for that group by some amount which we will denote k . In this case, the reservation quality is $\underline{\theta}$ for the favored group, but $\underline{\theta} + k$ for the disfavored group.¹³ It follows that each landlord compares $\hat{\theta}_r$ and $(\hat{\theta}_{-r} - k)$ against $\underline{\theta}$, respectively, where r denotes the favored group and $-r$ denotes the disfavored group.

III. Testable Implications

Our model of discrimination presents four predictions. Although each landlord’s estimates of the intercept terms ($\hat{\mu}_r^L$) and the information weighting parameters ($\hat{\gamma}_r$) are unobservable, we can experimentally manipulate signals sent by applicants to infer the average parameter values from landlord responses.

If we, the researchers, could observe each landlord's sample of θ_r and X_r , we could separately average the numerator and denominator of the information weighting parameter across the sample of landlords to obtain:

$$\gamma_r \equiv \frac{(1/n) \sum \hat{c}ov(\theta_r, X_r)}{(1/n) \sum \hat{v}ar(X_r)} \quad (7)$$

Similarly, we have the average of the intercept term:

$$\mu_r^L \equiv \sum (\bar{\theta}_r/n) - \gamma_r \sum (\bar{X}_r/n) \quad (8)$$

In a large sample, equations (7) and (8) yield the means.

A. Testable Implication 1

With random assignment of race $r = \{W, B\}$ to a fictional applicant and in the absence of both additional signals and taste-based discrimination (i.e. $k = 0$), we can test whether landlord responses are consistent with an average landlord having $\hat{\theta}_W > \hat{\theta}_B$, $\hat{\theta}_W = \hat{\theta}_B$, or $\hat{\theta}_W < \hat{\theta}_B$.

Equations (4) and (8) imply that:

$$E(\hat{\theta}_r) = (1 - \gamma_r)\mu_r + \gamma_r\mu_r = \mu_r \quad (9)$$

where $E(\hat{\theta}_r) \equiv E(\hat{\theta}_r | X = E(X_r))$.

Since we cannot rule out taste-based discrimination *a priori*, the mean difference in the predicted quality between white and black applicants must be adjusted by the mean of the taste parameter $E(k) \equiv K$:

$$E(\hat{\theta}_W) - E(\hat{\theta}_B) + E(k) = \mu_W - \mu_B + K \quad (10)$$

Equation (10) means that when no signal of quality other than race of the applicant is revealed, a racial gap favoring white applicants is consistent with both statistical discrimination and racial prejudice.

Since black applicants have lower socio-economic backgrounds – as evident in most nationally representative surveys – and are more likely to suffer from racial prejudice according to numerous studies such as those referenced in section I, this leads to:

Hypothesis 1 – A white applicant is more likely to receive a positive response than a black applicant.

B. Testable Implication 2

If we randomly assign a negative signal $-\tilde{X}^- < 0$ or a positive signal $\tilde{X}^+ > 0$, and race to different applicants and present them to randomly selected landlords, we can use a difference-in-difference approach to make inferences independent of the simple taste-based explanation and test whether $\hat{\gamma}_W > \hat{\gamma}_B$, $\hat{\gamma}_W = \hat{\gamma}_B$, or $\hat{\gamma}_W < \hat{\gamma}_B$ on average.

Equation (2) and the possibility of preference-based discrimination imply that the mean difference between black and white applicants sending a positive signal is:

$$E(\hat{\theta}_B | \tilde{X}^+) - E(\hat{\theta}_W | \tilde{X}^+) - E(k) = (\mu_B^L - \mu_W^L) - (\gamma_W - \gamma_B)\tilde{X}^+ - K \quad (11)$$

Similarly, the mean difference between black and white applicants sending a negative signal adjusted for potential preference-based discrimination is:

$$E(\hat{\theta}_B | -\tilde{X}^-) - E(\hat{\theta}_W | -\tilde{X}^-) - E(k) = (\mu_B^L - \mu_W^L) + (\gamma_W - \gamma_B)\tilde{X}^- - K \quad (12)$$

Taking the difference of equations (12) and (11) yields:

$$E(\hat{\theta}_B - \hat{\theta}_W | -\tilde{X}^-) - E(\hat{\theta}_B - \hat{\theta}_W | \tilde{X}^+) = (\gamma_W - \gamma_B)(\tilde{X}^+ + \tilde{X}^-) \quad (13)$$

The extent of dependence across signals of tenant quality which landlords obtained through their past experience influences the average sample variance of signal $\hat{\text{var}}(X_r)$.¹⁴ Massey and Denton (1987) and Iceland et al.’s (2002) description of residential segregation and neighborhood sorting implies that signals are positively correlated within a racial group. Landlords renting in neighborhoods that are predominantly white are relatively more experienced with white tenants than with black tenants. These landlords’ average sample variance of signals from white applicants will be smaller than that from black applicants because of neighborhood sorting.¹⁵ Since the average landlord in a nationally representative sample rents in a predominantly white neighborhood, we expect $\hat{\gamma}_w > \hat{\gamma}_B$ for her.¹⁶ The predicted differences in information weighting parameters are also consistent with previous studies (e.g., Bertrand and Mullainathan [2004]) that find smaller marginal returns to a positive treatment for blacks than for whites. Testable implication 2 yields:

Hypothesis 2 – On average, the information weighting parameter is larger for white applicants than is for black applicants.

C. Testable Implication 3

If we randomly assign each applicant a race and a negative signal that is below the mean, $-\tilde{X}^- < E(\bar{X}_r)$, or a positive signal that is above the mean, $\tilde{X}^+ > E(\bar{X}_r)$, then present them to randomly selected landlords, we can validate whether signals lead to differences in responses that are consistent with the sign of $(\gamma_w - \gamma_B)$ established in testable implication 2.

Call the difference between the signal a landlord observes and her expected signal for the no-signal base case a “surprise”: $\tilde{X}^+ - E(\bar{X}_r)$. With an identical positive signal for black and

white applicants and $E(\bar{X}_W) > E(\bar{X}_B)$, we have $[(\tilde{X}^+ - E(\bar{X}_B))] > [\tilde{X}^+ - E(\bar{X}_W)]$. However, the experimentally manipulated negative information will be a greater surprise for whites than for blacks: $-(\tilde{X}^- + E(\bar{X}_W)) < -[\tilde{X}^- + E(\bar{X}_B)]$. Since the preference parameter K is a constant, it is cancelled out. Depending on the relative size of γ_W and γ_B , we can get three different patterns of responses from landlords.

In case 1, where $\gamma_W = \gamma_B$, a surprising signal, whether positive or negative, will be weighted equally for blacks and whites. Since $[\tilde{X}^+ - E(\bar{X}_B)] > [\tilde{X}^+ - E(\bar{X}_W)]$, the positive treatment benefits black applicants more than white applicants. Similarly, because $-(\tilde{X}^- + E(\bar{X}_B)) > -[\tilde{X}^- + E(\bar{X}_W)]$, the negative treatment hurts white applicants more than black applicants. Hence, it follows that:

$$\gamma_B [\tilde{X}^+ - E(\bar{X}_B)] > \gamma_W [\tilde{X}^+ - E(\bar{X}_W)] \quad (14)$$

$$-\gamma_B [\tilde{X}^- + E(\bar{X}_B)] > -\gamma_W [\tilde{X}^- + E(\bar{X}_W)] \quad (15)$$

Expressions (14) and (15) imply that when compared to the (no information) base case, the gap in expected quality between the two racial groups closes in the presence of either positive or negative information (Case 1 of Figure 1).

In case 2, where $\gamma_W > \gamma_B$, expression (15) is unambiguously satisfied, but the relationship in (14) may not be true. Thus, when $\gamma_W > \gamma_B$, negative information will shrink the gap in expected quality between blacks and whites, but positive information will not necessarily narrow the gap (Case 2 of Figure 1).

Finally, in case 3, where $\gamma_W < \gamma_B$, expression (14) will be satisfied, but expression (15) will not necessarily be. In this case, the positive treatment will narrow the racial gap, but the negative treatment may not (Case 3 of Figure 1).

Therefore, given Hypothesis 2, we have:

Hypothesis 3 – On average, negative information will shrink the racial gap observed in the base case, but positive information will have an ambiguous effect on the racial gap observed in the base case.

D. Testable Implication 4

Our model and assumptions show that whether $\hat{\gamma}_W > \hat{\gamma}_B$, $\hat{\gamma}_W = \hat{\gamma}_B$, or $\hat{\gamma}_W < \hat{\gamma}_B$ for an average landlord depends on whether $\hat{v}ar(X_W) < \hat{v}ar(X_B)$, $\hat{v}ar(X_W) = \hat{v}ar(X_B)$, or $\hat{v}ar(X_W) > \hat{v}ar(X_B)$ on average. If the relative size of $\hat{\gamma}_r$ varies with $\hat{v}ar(X_r)$ in the direction predicted by the model, it suggests that landlords' behaviors are consistent with our model of statistical discrimination.

Given neighborhood sorting and positive covariance of signals, as the share of black residents in a neighborhood, S_B , increases, we expect the variance of signal for blacks, that a landlord obtains to decrease on average, and thus γ_B to increase. As S_B increases from 0 to 1, it is increasingly likely that $\gamma_B \geq \gamma_W$. The positive relationship between γ_B and S_B implies that the relationship between a surprising signal and shrinkage in the racial gap (testable implication 3) will also vary with S_B . As $S_B \rightarrow 1$, the effect of a surprising positive signal in narrowing the racial gap in positive response rates will become more evident (case 3 in Figure 1). Therefore, we have:

Hypothesis 4 – Positive treatment should shrink the racial gap in positive response rates relatively more in predominantly black neighborhoods. Conversely, negative treatment will shrink the racial gap in predominantly white neighborhoods, but not necessarily so in predominantly black neighborhoods.

IV. Experimental Design and Econometric Specifications

To examine the four hypotheses listed in the previous section we experimentally manipulate race and signals presented in emails to landlords who listed rental apartments on Craigslist.¹⁷ Email is an excellent vehicle to test the model implications and Craigslist serves as an ideal experimental platform with its focus on email communication. First, Hypothesis 1 requires limiting the information to agents to just race, which is straightforward in email correspondence but difficult in audit studies or other correspondence experiments. Next, Hypotheses 2 and 3 demand clear signals that are also unambiguously distinct (i.e. positive vs. negative), which can be flexibly introduced in emails. Finally, the low cost of email and the prevalence of Craigslist in most major U.S. cities allows a researcher to send to a large, representative sample of agents.¹⁸

A. Experimental Subjects and Rental Market Data

We use landlords who posted listings on Craigslist, an online classified ad website of enormous popularity, particularly amongst apartment seekers, as our experimental subjects. As of 2009, 40 million unique internet visitors view Craigslist each month and the site is often considered one of the principal factors responsible for the sharp fall of newspaper classified ad revenues.¹⁹ According to a study by internet research firm Hitwise, about 95% of visits to online classified websites are to Craigslist.²⁰ Data from Pew Internet & American Life Project also reveal that

roughly 44% of black and 49% of white adult internet users have at some point used online classified ads like Craigslist (Table 1). These Craigslist users, whether black or white, represent roughly one-third of the adult population in the U.S. They are slightly younger and more educated than non-Craigslist users and non-internet users, and they are more likely to be high-income earners, employed full-time, and apartment renters. Furthermore, Table 1 also indicates that black Craigslist users are younger and less educated, and more likely to be low-income earners, single, and renting apartments than white Craigslist users. Therefore, findings based on Craigslist will be relevant for a large fraction of black and white adults, especially those who use internet and online classified ads.

Our apartment selection algorithm attempted to eliminate scams, misplaced listings, repeated listings, and listings posted by individuals with “non-landlord” incentives. Those with non-landlord incentives include employees of large corporations managing dozens of apartments and private “apartment finders” who make a living as middlemen between landlords and renters. Sampled apartments include only one-bedroom and studio listings so as to avoid concerns of roommates, children, etc., and ensure comparable rents between any two units within an area. Only one inquiry was sent in response to a given listing, and numerous precautions were taken to avoid sending multiple inquires to the same landlord and/or the same listing.²¹ The search algorithm checked whether a new posting had a phone number, email or web address already encountered. Within each city, the sample excludes units with rents below the 20th and above the 90th percentile to avoid sending emails to storage lockers, weekly rentals or homes for sale. Finally, emails were sent over three intervals of time after the listing was posted with an upper limit of 48 hours.

Table 2 lists the cities we surveyed, the number of emails we sent in each city, the number of neighborhoods (census tracts) from which postings were sourced, the share of black population in each city, and the average rents of apartments that we sent email inquiries.²² Comparing column (3) and column (4) reveals that the average share of black residents across the sampled neighborhoods is fairly similar to the actual share of black population in the greater metropolitan area as shown in Census 2000.

B. Email Generation and Experimental Treatments

As this study focuses solely on email correspondence, the only mechanism for signaling the race and gender of the applicant is a stated name. To maximize the probability that landlords will observe this signal, the full name of the fictitious applicant is presented three times in every email: first in the email address, which is always of the form “first.last<random number>@domain.com,” second in the introductory sentence of the email text, and third in the closing signature of the email. First names chosen are those utilized by Bertrand and Mullainathan (2004) in their correspondence study, combined with surnames sourced from the U.S. Census 2000 Family Name Survey. Names resulting from this combination include: Allison Bauer, Ebony Washington, Matthew Klein, and Darnell Booker.²³

Each email text was generated by randomly selecting the text for each of the five elements numerated in the sample emails in Illustration 1. With the exception of the statement of quality, all text was pulled from the same pools. (1) is an introductory *hello* statement. (2) is a *statement of interest* in the apartment which always includes the rent of the unit being applied to (to avoid confusion in case the landlord has posted multiple listings). (3) is a *statement of quality* which is randomly included (or not included) to define our treatments. (4) is an *inquiry statement*

regarding the availability of the unit (e.g. “is this apartment still available?”). This gives the landlord a specific question to respond to, allowing us to identify automated responses and test for differences in positive responses between groups. (5) is a *closing* which thanks the landlord and is always followed with the applicant’s full name.

Element 3, the statement of quality, is included in approximately two-thirds of all emails. The email texts that do not include a statement of quality are said to belong to the “baseline treatment” or “base case.” In this treatment, landlords know nothing of the applicant except their name and their interest in renting their apartment. The model detailed above assumes that landlords simply take the average signals by race as a proxy signal in this baseline scenario. Landlords may expect our fictional black applicants to be less desirable than our fictional white applicants in the base case since a typical black Craigslist user has lower socio-economic status than a typical white Craigslist user (Table 1). When the statement of quality is included, it discloses either “positive” or “negative” information. Positive information always informs the landlord that the applicant has a good job and does not smoke. Negative information always informs the landlord that the applicant smokes and has a bad credit rating. These particular pieces of information were selected because they are unambiguously positive or negative. The purpose of this methodology is not to determine how any specific piece of information affects outcomes, but instead to test how positive or negative information, in general, affects outcomes.²⁴ It is difficult to imagine a scenario in which a landlord would benefit from a tenant who smokes or has bad credit. Likewise, it is difficult to imagine a landlord being harmed because a tenant has a good job or does not smoke. Landlords typically verify characteristics such as credit worthiness and smoking habits in the interview stage and commonly ask applicants to pay for credit rating checks. An applicant with a low credit rating or a smoking habit may thus

preemptively reveal such information to avoid paying for a likely-to-fail credit check or rejection after showing up for an apartment viewing.²⁵ Furthermore, our focus on how landlords treat negative signals differently by race ensures that any peculiarity in the sending of negative signals is differenced out. Last, given the average characteristics of online classified ad users reported in Table 1, the negative information is likely surprising to landlords, providing a strong treatment effect.

This process of inquiry generation comes with a significant benefit. By pulling all texts randomly from the same pools and defining the treatments entirely by the statement of quality alone we assure that any differences in landlord responses are the result of our treatments and not the introduction of inquiry texts. Table 3 summarizes the number of emails sent by each applicant type as defined by their race, gender, and treatment.

C. Categorization of Outcome

The simplest response characteristic is whether or not a given inquiry receives a response. Responses were further classified into one of several categories. To avoid experimenter bias in this categorization, all instances of applicant names (first and last, as well as email address) and original bodies of text sent were automatically removed from view during categorization. Broadly, responses are classified as either positive or not-positive. Positive responses state that the unit is available and invite future contact in some manner. Non-positive responses include the non-response emails and those either stating that the unit is not-available, or stating that the unit is available, but in a discouraging manner. Each inquiry sent ended with a question such as “Is the apartment available?” Some 95% of landlords that answered “Yes” to that question also asked for further contact information (coded as a positive response). An email response that

simply read “Yes” lacks any direct contact information or interest and likely meant the landlord was not encouraging the applicant for future viewing of the unit, and was classified as “Disinterested.” The categorization process is crucial to our results as we are primarily interested in the differential treatment of black and white applicants by landlords. Differences in the likelihood of simply receiving a response may be misleading, since one group may receive a larger share of negative responses than the other. Our careful reading and categorization of all landlord responses takes into account the importance of the contents of a response.

D. Econometric Specifications

Given our experimental design, we estimate four regression equations that correspond to the testable implications of the model.

First, the empirical specification for equation (10) is:

$$R_i = \alpha_w + \alpha_B B_i + u_i \quad (16)$$

R is 1 if the landlord owning apartment i responded positively, 0 otherwise; B is 1 for an applicant with an African-American sounding name, 0 otherwise; and u_i is an error term.

Hypothesis 1 states that either taste-based or statistical discrimination predicts $\alpha_B < 0$. The possibility of the former implies that we cannot infer expected tenant quality from $\hat{\alpha}_B$. Because almost all our empirical specifications involve dichotomous regressors, we use linear probability model (OLS) to estimate all our empirical specifications. Simply, the issues of predicted probabilities outside the unit interval or incorrect marginal inference disappear with majority dummy independent variables (see Wooldridge (2003, pp 456-7) for further discussion).

Second, the difference-in-difference specification of equation (13) that tests hypothesis 2 ($\gamma_w > \gamma_B$) is:

$$R_i = \alpha_{PW} + \alpha_{PB}(B_i) + \alpha_{NW}(N_i) + \alpha_{NB}(N_i \times B_i) + u_i. \quad (17)$$

N takes the value of 1 if negative information is presented, 0 if positive information is presented. The omitted category is positive information for whites. If $\gamma_w > \gamma_b$, then the average landlord weights white applicant signals more heavily than those from black applicants. Its econometric analog α_{NB} should be positive, resulting in a greater marginal return to credentials for white applicants.

To verify the prediction that shrinkage in the racial gap is consistent with the relative size of the information weighting parameters, we estimate the following difference-in-difference regression:

$$R_i = \beta_w + \beta_B(B_i) + \beta_{PW}(P_i) + \beta_{PB}(P_i \times B_i) + \beta_{NW}(N_i) + \beta_{NB}(N_i \times B_i) + u_i. \quad (18)$$

The coefficients β_{PB} and β_{NB} measure the extent of shrinkage in the racial gap of positive response rates in the presence of a (surprising) positive and negative signal. According to hypothesis 3, we expect that $\beta_{NB} > 0$ and the sign of β_{PB} ambiguous.

Finally, the following empirical specification tests hypothesis 4:

$$\begin{aligned} R_i = & \delta_w + \delta_{SW}(S_{Bi}) + \delta_B(B_i) + \delta_{SB}(S_{Bi} \times B_i) + \delta_{PW}(P_i) + \delta_{SPW}(S_{Bi} \times P_i) + \delta_{PB}(P_i \times B_i) \\ & + \delta_{SPB}(S_{Bi} \times P_i \times B_i) + \delta_{NW}(N_i) + \delta_{SNW}(S_{Bi} \times N_i) + \delta_{NB}(N_i \times B_i) \\ & + \delta_{SNB}(S_{Bi} \times N_i \times B_i) + u_i \end{aligned} \quad (19)$$

S_{Bi} measures the fraction of black residents in the neighborhood (%Black) in which apartment i is listed and it ranges between 0 and 1. The terms S_{Bi} and $S_{Bi} \times B_i$ allow the (unobserved) expected value of X to vary across different types of neighborhoods and race. If landlords' experiences with black applicants increase the size of their information weighting parameter,

then we expect $\delta_{SPB} > 0$. Note that the preference parameter K is subsumed in the dummy variable for black applicants, as well as the interaction term between the share of black residents in the neighborhood and the race indicator. If landlords renting in a predominantly black neighborhood exhibit a preference for black residents relative to landlords renting in neighborhoods with a lesser share of black residents, a typical assumption made about preference toward an agent's own racial group, then we would expect δ_{SB} to be positive.

V. Results

Table 4 presents summary statistics of variables generated in our experiment, characteristics of the listed apartments, and responses from landlords. Of the 14,237 inquiries sent, 9,229 (65 percent) received a response. Of these responses 6,597 (46.3 percent) were positive as defined in section IV. Figure 2 shows the distribution of the shares of black residents in census tracts of listed apartments (S_{Bi} in equation (19)). The measure ranges from 100 percent white to 98.45 percent black with a mean of 12.4 percent black residents. Table 5 verifies that the characteristics of our fictitious white applicants and black applicants are statistically similar and not correlated with characteristics of listed apartments by treatment types.

A. Effective Informational Treatments

Table 6 reports response rates for positive and negative treatment relative to the baseline of no signal, pooling all applicants.

Positive response rate is a less noisy measure of expected applicant quality than response rate. Comparing the intercept terms in column (1) and column (3), which measure response and positive response rates respectively in the baseline treatment, reveals that roughly 18 percent of

responses in the baseline were negative in some way. This means that the simple rate of response is likely to misrepresent whether landlords encouraged future contact. Furthermore, the estimates in column (1) and column (3) show that landlords are equally likely to reply an email inquiry whether or not the applicant has revealed something positive about herself or not, but landlords are more likely to reply with a rejection if the tenant revealed nothing about her quality. Therefore, considering a “no” response as equivalent to a “yes” response is likely to invite error into our data interpretation.

As shown in column (3) of Table 6, applicants in the positive treatment group receive a significantly higher positive response rate (56% vs. 53%), receiving 107 positive responses for every 100 received by the baseline applicants. The effect of positive treatment is slightly higher for females (4%) than for males (3%). On the other hand, column (4) shows that applicants in our negative treatment receive a statistically significantly lower positive response rate (approximately 32% vs. 53.0%), receiving only 41 positive responses for every 100 received by baseline applicants. The significant differences illustrate that the treatment effectively manipulated landlord interest in the fictional applicants. Finally, the insignificant differences in response rates across gender and independent of race imply that differential treatment by race is negligible, so hereafter, we pool genders.

B. Hypothesis 1: Black Applicants receive Lower Response Rate

Column (1) in Table 7 confirms hypothesis 1 that landlords, on average, are more likely to respond to applicants with white sounding names than applicants with African-American sounding names when no other signal of quality is disclosed in an email inquiry. The estimated coefficient on Black of -0.093 is highly significant and confirms previous findings of

discrimination against African Americans or persons with African-American sounding names. Combined with the intercept estimate of 0.581, applicants with African-American sounding names receive about 84 positive responses for every 100 received by applicants with white sounding names. The estimates are consistent with the interpretation that average landlords expect the average black tenants to have lower quality than the average white tenants. This finding does not rule out preference-based discrimination as the racial gap is also consistent with racial prejudice (i.e., $K > 0$).

C. Hypothesis 2: Differential Information Weighting Parameters by Race

Column (2) in Table 7 presents the difference-in-difference estimates for equation (17). It shows that the estimated coefficient of the difference-in-difference effect of negative treatment for black is significantly positive (0.042). It confirms hypothesis 2 that on average, $\gamma_W > \gamma_B$, so signals from white applicants receive relatively more weight in an average landlord's estimate of quality. The difference-in-difference specification ensures that the estimates are not driven by taste-based discrimination. This finding potentially explains why some audit and correspondence studies show a lower margin return to credentials for blacks.

D. Hypothesis 3: Asymmetric Shrinkages by Race of Applicants

The experiment's effective treatment and the finding that that the average landlord's information weighting parameter for whites is greater than that for blacks suggests that the racial gap in the average landlord's positive responses will close with negative information, but not necessarily with positive information (hypothesis 3).

Table 7 column (3) shows that the marginal return to signaling a respectable occupation

and non-smoking behavior increases positive response rate by 6.7 percent (0.039/0.581) for whites. However, the coefficient on the interaction term “Black x Positive Information” is statistically indistinguishable from zero. Both black and white applicants benefit from the inclusion of positive information, but the information does not widen or narrow the racial gap observed in the base case. This result confirms hypothesis 3 and conforms to the model prediction $\gamma_W > \gamma_B$: the average landlord weights identical information more from white applicants. Although the positive information may also correlate with greater surprise to the average landlord about applicants with African-American sounding names, the relationship $\gamma_W > \gamma_B$ attenuates any improvement in treatment. Simply, when $\gamma_W > \gamma_B$, the model predicts a signal “I have a paying job” will disproportionately benefit white applicants and – in the presence of baseline differential treatment – potentially worsen relative outcomes.

The statistically negative coefficient on negative information and positive coefficient on “Negative Information x Black” shows that disclosing negative information about an applicant’s quality leads to a greater reduction in a white applicant’s probability of receiving a positive response. In particular, the negative information almost halves the racial gap observed in the baseline treatment. The model predicts that this change occurs when the negative signal from white applicants is more surprising to the average landlord and such signals receive relatively more weight from these applicants. The estimates are consistent with hypothesis 3.

E. Hypothesis 4: Differential Information Weighting Parameters across Neighborhoods

The final model prediction states that the difference $\gamma_W - \gamma_B$ is negatively related to the share of blacks in a rental property’s neighborhood. A landlord in a predominantly black neighborhood presumably has more experience screening and renting to black residents, lowering the relative

noise of signals from black applicants. So, the weighting parameters for black and white applicants approach each other as the share of black and white residents equalize.

Table 7 column (4) presents evidence that the information weighting parameters vary with the racial composition of an apartment's neighborhood (%Black) in a manner consistent with the model's predictions. The statistically significant and positive coefficient on the interaction term "Positive Information x Black x %Black" indicates that as the share of black residents in a neighborhood increases, the surprising positive signal becomes more effective in closing the racial gap in positive response rates between white and black applicants. In contrast, the insignificant positive coefficient on the interaction term "Negative Information x Black x %Black" indicates that in a predominantly black neighborhood, the negative signal does not significantly decrease the racial gap between black and white applicants when compared with the baseline treatment case. Thus, when we examine a non-average neighborhood, where residents are predominantly black, a surprising positive signal leads to significantly greater improvements in positive response rates. Although landlords in such a neighborhood are more likely to be black, it does not require that the landlord is black, but only that the landlord is more familiar with black applicants. Our estimates confirm the key testable implication of statistical discrimination (hypothesis 4): landlords' relative past experiences with different racial groups shape their information parameters.

Table 7 column (4) also reveals that the racial gap between white and black applicants in the base case does not vary across different neighborhoods. The insignificant estimate echoes Bertrand and Mullainathan's (2004) finding that job interview callback rates do not vary significantly across types of occupations, industries, and applicants' residential neighborhoods. However, based on the insight gained from our difference-in-difference specification and the

significant coefficient on “Positive Information x Black x %Black”, we arrive at different conclusions about statistical discrimination. With the significant treatment effect of positive information, our estimator can identify differences in γ across race and in turn illustrate that the lack of experience with a particular race influences the behavior of agents.

VI. Ruling Out Alternative Explanations

A. Preference-based Discrimination

Are our experimental results consistent with preference-based discrimination? It is possible that the observed differential outcomes by race, especially in the baseline treatment (Table 7 column 1), are the result of preference-based agent or customer prejudice (Yinger 1986). Agent prejudice is present when landlords are prejudiced against minority applicants, while customer prejudice exists if landlords reject black applicants for fear of offending prejudiced white tenants. Without characteristics of landlords it is difficult to separately test whether our empirical findings are consistent with agent prejudice or customer prejudice. Nevertheless, variation in response rates across neighborhoods with differing racial compositions can potentially eliminate preference-based explanations.

If landlord behavior is fully motivated by prejudice – be it their own or for the sake of their customers/tenants – we would expect landlords renting in predominantly white neighborhoods to be more likely to discriminate against black applicants, regardless of the informational signals presented. What we find, however, is that landlords who rent in predominantly white neighborhoods (i.e., the average neighborhoods) react more negatively to negative information from white applicants than from black applicants (estimates in Table 7 column 3). This clearly rejects such a preference-based argument and indicates that preference

based discrimination alone cannot explain our results. Similarly, the negative (though insignificant) coefficient of Black \times %Black reported in column (4) of Table 7 shows that landlords renting in predominantly black neighborhoods are no more likely to respond to black applicants, which is inconsistent with the prediction of preference-based customer prejudice. Moreover, as landlords renting in predominantly black neighborhoods are more likely to be black, the estimates also do not support the prediction of agent prejudice.

B. Lexicographic Search

Bertrand and Mullainathan (2004) argued that a lexicographic search model could explain their findings which show that job applicants with African-American sounding names and relatively high quality resumes did not obtain higher interview callback rates when compared to those with relatively low quality resumes. The ineffective treatment effects for minority applicants may imply that employers use the “read-no-further” rule when reading an African-American sounding name to screen applicants. Our results in columns (2), (3), and (4) of Table 7 show that positive information significantly increases the likelihood that applicants with African-American sounding names receive positive responses, ruling out the presence of lexicographic search.

VII. Potential Limitations

The paper and research design have several limitations. First, we could not collect information on landlord characteristics. The apartment listings rarely detail the landlord’s gender, age or location while Craigslist’s communication system also adds a layer of anonymity. We could potentially gather some characteristics of landlords among those who responded, but the information is at best limited to gender (as indicated by their names, should they choose to

provide one). Second, the above analysis also ignores most apartment-level characteristics beyond location and number of bedrooms. Coding up features such as the building size or amenities is as difficult as identifying landlord characteristics. Fortunately, our model does not require detailed landlord characteristics or apartment hedonics, as we simply ask whether behavior is consistent with the model. Third, the focus of our model and research design centers on discrimination in screening but not final offer. This means that our findings may not directly inform the extent of and reasons for differential final outcomes by race in the apartment rental market. Nonetheless, passing the initial screening stage is necessary for getting an interview and subsequently securing the final offer. So our model and findings yield important insights to our understanding of how statistical discrimination may lead to differential final outcomes by race.

VIII. Robustness Checks

A. Alternative Measures of Positive Responses

Table 8 presents results using other definitions of positive responses as the dependent variable. We chose to present estimates based on specification equation (19), as it is general and effectively illustrates the robustness of our results to alternative measures of positive responses. The estimates in columns (1) – (3) show that the estimated coefficients vary little across all three different measures of positive responses.

B. Name Choice

Column (4) of Table 8 presents estimation results of equation (19) without a set of less common first names: Hakim, Rasheed, and Tremayne, each have a frequency of less than 0.005% in Census 2000. Column (4) shows that all earlier conclusions about the testable implications

remain. Despite the reduction in sample size, the parameter estimate for the hypothesis about the racial composition in a neighborhood and information weighting remains similar and statistically significant at the 10% level. Based on unreported estimates, the results are also robust to the exclusion of four Muslim sounding first names: Hakim, Jamal, Karim, and Rasheed.

C. Social Background and Responses

If the names chosen in our study convey an applicant's social background beyond race, then discrimination observed is not necessarily racial. If differences in the social background associated with names explain the observed differences, then names associated with better social background should receive more positive responses. We follow Bertrand and Mullainathan's (2004) approach to examine whether average positive response rates are correlated with social background of each name within each race-gender group, using the fraction of mothers of babies born with the names who have at least a high-school diploma as a proxy. The within race-gender rank-order correlation test reported in Table 9 shows no evidence that positive response rate and social background are positively related. The p -values show that we cannot reject independence between positive response and average mother's education at the 5% level of significance, except in the case of black male applicants, where the correlation is significantly negative at the 5% level.

IX. Conclusion

Statistical discrimination can explain the differential outcomes in rental apartment inquiry screening by race. We detail a model of statistical discrimination that provides testable implications about a parameter that connects signals, expectations and race. The model also

implies a research design that distinguishes statistical discrimination from racial prejudice. We show that isolating behavior consistent with statistical discrimination requires an experiment with strong treatment effects provided by negative, neutral, and positive signals sent to subjects with different experience. Such a structure invites a difference-in-difference estimator that was not previously used in the literature on discrimination. Our experiment and estimation method produce results that contrast with those of earlier studies that lacked strong treatment effects, distinct signal types, and/or a setting free of confounding unobservable factors.

We show that, when no information other than race of an applicant is revealed to landlords, applicants with African-American sounding names receive 16 percent fewer positive responses than applicants with white sounding names. Our finding that the average landlords are more responsive to an identical signal from an applicant with a white sounding name than from one with an African-American sounding name potentially explains why previous studies found lower marginal returns to credentials for minorities. The variation in landlord response rates across neighborhood racial compositions conforms to the statistical discrimination model where agents use past experience to predict applicant quality by race. Racial prejudice or lexicographic search cannot easily explain the results. The findings provide justification for policies aiming to promote clear information dissemination and to improve communication between different racial groups, as well as for social programs designed to eliminate inequality across racial groups.

Environments of imperfect information where agents send limited signals along with their race should exhibit patterns found in this study. The applicability of our research design and initial screening environment suggests that such behavior might also exist in other situations such as bank loans, college applications and job search.

References

- Ahmed, Ali, and Mats Hammarstedt. 2008. "Discrimination in the Rental Housing Market: A field experiment on the internet." *Journal of Urban Economics*, 64: 362-372.
- Aigner, Dennis, and Glen G. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *Industrial and Labor Relations Review*, 30(2): 175-187.
- Altonji, Joseph, and Charles R. Pierret. 2001. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics*, 116(1): 313-350.
- Antonovics, Kate, Peter Arcidiacono, and Randall Walsh. 2005. "Games and Discrimination: Lessons from *The Weakest Link*." *Journal of Human Resources*, 40(4): 918-947.
- Arrow, Kenneth. 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets*, ed. Orley Ashenfelter and Albert Rees, Princeton University Press.
- Ayres, Ian, and Peter Siegelman. 1995. "Race and Gender Discrimination in Bargaining for a New Car." *American Economic Review*, 85(3): 304-321.
- Balsa, Ana I. and Thomas G. McGuire. 2001. "Statistical Discrimination in Health Care." *Journal of Health Economics*, 20: 881-907.
- Becker, Gary S. 1957. *The Economics of Discrimination*, University of Chicago Press.
- Bertrand, Marianne and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review*, 94(4): 991-1013.
- Bosch, Mariano, Lidia Farre, and Maria A. Carnero. 2010. "Information and Discrimination in the Rental Housing Market: Evidence from a Field Experiment." *Regional Science and Urban Economics*, 40(1): 11-19.

Carpusor, Arian G. and William E. Loges. 2006. "Rental Discrimination and Ethnicity in Names." *Journal of Applied Social Psychology*, 36(4): 934-952.

Doleac, Jennifer L. and Luke Stein. 2010. "The Visible Hand: Race and Online Market Outcomes." SIEPR Discussion Paper 09-015.

Hanson, Andrew and Zachary Hawley. 2010. "Do Landlords Discriminate in the Rental Housing Market? Evidence from an Internet Field Experiment in U.S. Cities." Unpublished.

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

Heckman, James and Peter Siegelman. 1992. "The Urban Institute Audit Studies: Their methods and findings." *Urban Institute Press*, 187-258.

Iceland, John, Daniel H. Weinberg, and Erica Steinmetz. 2002. *Racial and Ethnic Residential Segregation in the United States, 1980-2000*. U.S. Census Bureau, Special Report Series, CENSR #3.

List, John A. 2004. "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field." *Quarterly Journal of Economics*, 119(1): 49-89.

Levitt, Steven D. 2004. "Testing Theories of Discrimination: Evidence from the Weakest Link." *Journal of Law and Economics*, 47(Oct): 431-452.

Massey, Douglas S. and Nancy A. Denton. 1987. "Trends in the Residential Segregation of Blacks, Hispanics, and Asians: 1970-1980", *American Sociology Review*, 52(6): 802-825.

Ondrich, Jan, Stephen Ross, and John Yinger. 2003. "Now you see it now you don't: Why do real estate agents withhold available houses from black customer?" *Review of Economics and Statistics*, 85(4): 854-873.

Ondrich, Jan, Alex Stricker, and John Yinger. 1998. "Do real estate brokers choose to discriminate? Evidence from the 1989 Housing Discrimination Study." *Southern Economics Journal*, 64(4): 880-901.

Ondrich, Jan, Alex Stricker, and John Yinger. 1999. "Do Landlords Discriminate? The incidence and causes of racial discrimination in rental housing markets." *Journal of Housing Economics*, 8(3): 185-204.

Page, Marianne. 1995. "Racial and Ethnic Discrimination in Urban Housing Markets: Evidence from a recent audit study." *Journal of Urban Economics*, 38(2): 183-206.

Phelps, Edmund S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review*, 62(4):659-661.

Roychudry, Goodman. 1996. "Evidence of Racial Discrimination in Different Dimensions of Owner-Occupied Housing Search," *Real Estate Economics*, 24(2): 161-178.

Ruggles, Stephen, J., Trent Alexander, Katie Genadek, Ronald Green, Matthew B. Schroeder, and Matthew Sobek. 2010. *Integrated Public Use Microdata Series: Version 5.0* [Machine-readable database], University of Minnesota.

Siddique, Zahra. 2008. "Caste Based Discrimination: Evidence and Policy." IZA Discussion Paper No.3737.

U.S. Census Bureau, Census 2000 Summary File 3.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

Yinger, John. 1986. "Measuring Racial Discrimination with Fair Housing Audits: Caught in the Act." *American Economic Review*, 76(5): 881-893

Illustration 1: Representative Email Samples

<i>Positive Treatment</i>	<i>Negative Treatment</i>
(1) Hello, (2) My name is [Full Name], and I am writing in response to your listing for an apartment for [apartment rent]/month. (3) In case you're interested, I do not smoke and I work full time as an architect. (4) Is this unit still available? (5) Thank you for your time, [Full Name]	(1) Hi, (2) My name is [Full Name], I am responding to your craigslist posting for an apartment listed at [apartment rent]/month. (3) Just so you know, I am a smoker and my credit rating is below average. (4) I realize places go fast sometimes, is this unit still available? (5) Thanks, [Full Name]

Table 1: Comparison of Craigslist Users and Non-Craigslist Users

Mean	<i>Full Sample</i>		<i>Craigslist^(g)</i>		<i>Non-Craigslist</i>		<i>Internet^(f)</i>		<i>Non-internet</i>	
	Black	White	Black	White	Black	White	Black	White	Black	White
Age (years)	42.81	47.94	35.72	40.85	38.88	47.84	37.49	44.39	53.87	61.82
Male	0.46	0.48	0.47	0.54	0.51	0.45	0.49	0.50	0.38	0.43
College ^(a)	0.40	0.55	0.57	0.66	0.52	0.60	0.54	0.63	0.11	0.24
Low income ^(b)	0.63	0.40	0.52	0.32	0.70	0.35	0.62	0.33	0.65	0.65
Renter ^(c)	0.55	0.21	0.57	0.23	0.53	0.17	0.55	0.20	0.57	0.24
Single ^(d)	0.40	0.19	0.41	0.21	0.43	0.21	0.42	0.21	0.36	0.13
Full-time job ^(e)	0.41	0.45	0.63	0.57	0.43	0.49	0.52	0.53	0.20	0.16
Internet user ^(f)	0.67	0.79	--	--	--	--	--	--	--	--
Craigslist user ^(g)	0.30	0.39	--	--	--	--	0.44	0.49	--	--
Sample size	684	4311	202	1683	258	1739	460	3422	224	889

Notes: Authors' own calculation based on Pew Internet & American Life Project's "April 2009 – Economy" survey data of adult population. Only the sample of non-Hispanic whites and blacks are included. (a) Respondents with at least some college education; (b) persons earning less than \$50,000 per year; (c) persons renting apartments/houses; (d) never married or single persons; (e) persons employed full time; (f) persons who at least use the internet occasionally; (g) internet users who responded yes to "used online classified ads or sites like Craigslist."

Table 2: Cities Surveyed

<i>City</i>	<i>#Obs.</i>	<i>#Neighborhoods</i>	<i>Mean %Black across Neighborhoods</i>	<i>%Black in Metro</i>	<i>Mean Monthly Rent</i>
Atlanta	304	115	27.4%	29.4%	757.43
Austin	198	61	6.8%	7.4%	719.55
Baltimore	499	177	28.1%	27.4%	848.16
Boston	1324	413	6.1%	6.6%	1062.74
Charlotte	241	72	24.1%	20.3%	725.55
Chicago	596	216	15.9%	18.7%	1087.60
Cleveland	372	151	15.7%	18.2%	561.24
Dallas	150	51	11.9%	15.0%	873.59
Denver	744	230	6.2%	5.6%	728.71
Detroit	461	189	16.0%	22.6%	596.73
District of Columbia	1179	326	24.3%	26.2%	1353.51
Houston	313	99	13.7%	17.4%	794.60
Indianapolis	158	82	18.5%	14.0%	543.92
Jacksonville	126	51	21.0%	21.5%	672.43
Kansas City	276	117	14.6%	12.8%	589.39
Los Angeles	1029	482	7.4%	9.4%	1186.57
Louisville	239	63	14.9%	15.2%	549.50
Memphis	112	34	36.3%	44.1%	662.22
Milwaukee	219	89	11.9%	15.2%	621.27
Minneapolis	761	271	7.9%	5.3%	761.04
Nashville	181	66	20.1%	15.4%	794.80
Oklahoma City	179	76	12.0%	11.3%	492.27
Philadelphia	554	203	21.1%	19.5%	914.97
Phoenix	273	115	3.6%	3.4%	607.15
Portland	303	124	4.0%	2.7%	770.77
Raleigh	255	90	22.2%	22.1%	645.14
San Diego	793	273	5.1%	5.4%	1045.49
San Francisco	427	132	5.3%	5.2%	1471.80
San Antonio	86	36	5.4%	6.4%	612.77
San Jose	255	112	2.6%	2.6%	1171.98
Santa Barbara	164	40	1.9%	2.3%	1336.64
Seattle	448	182	4.6%	4.2%	935.66
Tampa	667	220	10.6%	9.8%	677.56
Tucson	351	78	2.8%	2.7%	532.82
Total	14237	5036	12.4%	12.9%	905.51

Note: (a) a neighborhood is a Census tract if cross-street information of the posting is available; otherwise it is a metropolitan statistical area; (b) %Black is defined as the number of non-Hispanic blacks divided by all population in census tract; the mean is obtained by averaging %Black across neighborhoods within the same city; (c) %Black in metropolitan statistical area based on the 5% public use Micro sample; (d) mean rent is calculated using the rents of units we surveyed. Population data sourced from Census 2000 Summary File 1 and the Integrated Public Use Microdata Series Census 2000 5% sample (Ruggles et al. 2010).

Table 3: Count of Observations by Race, Gender, and Treatment

	<i>Treatment</i>	<i>Male</i>	<i>Female</i>	<i>Pooled</i>
Black	Negative Info.	1004	956	1960
	Baseline	1061	1036	2097
	Positive Info.	1538	1501	3039
White	Negative Info.	984	1010	1994
	Baseline	1031	1098	2129
	Positive Info.	1446	1572	3018
Total		7064	7173	14237

Notes: Black is an applicant with an African-American sounding name. White is an applicant with a white sounding name. Male is an applicant with a male sounding name. Female is an applicant with a female sounding name. Baseline treatment refers to email text containing no information about credit rating, smoking, or occupation of an applicant. Negative treatment adds negative information about bad credit rating and smoking behavior to baseline email text. Positive treatment adds positive information about occupation and non-smoking behavior to baseline email text.

Table 4: Summary Statistics

<i>Variable</i>	<i>Obs.</i>	<i>Mean</i>	<i>Std. Dev.</i>	<i>Min</i>	<i>Max</i>
Sent on weekend	14237	0.272	0.445	0	1
Monthly rent	14237	905.5	323.68	350	2000
Negative information	14237	0.278	0.448	0	1
Baseline treatment	14237	0.297	0.457	0	1
Positive information	14237	0.425	0.494	0	1
Male	14237	0.496	0.500	0	1
Black	14237	0.498	0.500	0	1
% male in neighborhood	14237	0.497	0.041	0.25	1
% black in neighborhood	14237	0.124	0.162	0	0.984
Response	14237	0.648	0.478	0	1
Positive Response	14237	0.463	0.499	0	1

Notes: See definitions of monthly rent, % blacks in neighborhood, and neighborhood in notes of Table 2. See definitions of male, female, black, and white in notes of Table 3. Neighborhood demographic characteristics are sourced from Census 2000. Response indicates whether a landlord responded and positive response indicates whether a landlord responded positively to the inquiry. Positive response includes “Available” and “Available + if”. See Appendix B for response categories.

Table 5: Verification of Random Assignment

	<i>Baseline Treatment</i>			<i>Positive Information</i>			<i>Negative Information</i>		
	Black	White	Diff.	Black	White	Diff.	Black	White	Diff.
Pooled Gender									
Sent on weekend	0.259	0.266	-0.006 (.012)	0.263	0.278	-0.012 (0.012)	0.282	0.280	0.002 (0.015)
Monthly rent	895.13	908.95	-13.82 (10.76)	906.75	919.35	-12.60 (8.31)	895.31	899.96	-4.65 (9.44)
% black in neighborhoods	0.122	0.128	-0.006 (.005)	0.124	0.120	0.004 (0.004)	0.127	0.123	0.004 (0.005)
% male in neighborhoods	0.4976	0.496	0.0016 (.0014)	0.4973	0.4971	0.0003 (0.001)	0.499	0.498	-0.001 (0.001)

Notes: See definitions of variables in notes of Table 1, Table 2, and Table 3. Robust standard errors clustered by neighborhoods reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 6: Overall Treatment Effects on Response Rate and Positive Response Rate

	(1)	(2)	(3)	(4)
	----- Response -----		----- Positive Response -----	
Genders Pooled				
Positive Information	-0.005 (0.011)		0.039*** (0.011)	
Negative Information		-0.215*** (0.012)		-0.315*** (0.013)
Constant	0.710*** (0.010)	0.710*** (0.010)	0.534*** (0.011)	0.534*** (0.011)
Observations	10283	8180	10283	8180
R-squared	0.000	0.048	0.001	0.105
Males				
Positive Information	-0.009 (0.014)		0.034** (0.015)	
Negative Information		-0.217*** (0.017)		-0.312*** (0.017)
Constant	0.707*** (0.013)	0.707*** (0.013)	0.525*** (0.014)	0.525*** (0.014)
Observations	5076	4080	5076	4080
R-squared	0.000	0.049	0.001	0.104
Females				
Positive Information	-0.002 (0.013)		0.043*** (0.014)	
Negative Information		-0.213*** (0.016)		-0.318*** (0.016)
Constant	0.713*** (0.012)	0.713*** (0.012)	0.544*** (0.012)	0.544*** (0.012)
Observations	5207	4100	5207	4100
R-squared	0.000	0.048	0.002	0.106

Notes: The omitted category is the baseline (no-information) treatment. All samples pooled white and black applicants. See definitions of variables in notes of Table 1, Table 2, and Table 3. Robust standard errors clustered by neighborhoods reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 7: Differential Treatment by Race and Informational Signals

	(1)	(2)	(3)	(4)
Black	-0.093*** (0.015)	-0.092*** (0.012)	-0.093*** (0.015)	-0.084*** (0.019)
Positive Information			0.039*** (0.013)	0.053*** (0.017)
Positive Information x Black			0.001 (0.019)	-0.032 (0.025)
Negative Information		-0.377*** (0.013)	-0.338*** (0.016)	-0.347*** (0.018)
Negative Information x Black		0.044** (0.018)	0.045** (0.020)	0.044* (0.026)
% Black				0.014 (0.067)
Black x %Black				-0.077 (0.099)
Positive Information x %Black				-0.118 (0.082)
Positive Information x Black x %Black				0.267** (0.125)
Negative Information x %Black				0.078 (0.093)
Negative Information x Black x %Black				0.009 (0.130)
Constant	0.581*** (0.012)	0.619*** (0.009)	0.581*** (0.012)	0.579*** (0.014)
Omitted category	White Baseline	White Pos. Info.	White Baseline	White Baseline
Observations	4226	10011	14237	14237
R-squared	0.009	0.128	0.100	0.101

Notes: See definitions of variables in notes of Table 1, Table 2, and Table 3. Robust standard errors clustered by neighborhoods reported in parentheses. Columns (1), (2), (3), and (4) correspond to testable implications 1, 2, 3, and 4, respectively. *** p<0.01, ** p<0.05, * p<0.1, + p<0.15

Table 8: Alternative Measures of Positive Response and Excluding Rare First Names

	(1)	(2)	(3)	(4)
	<i>Alternative Measures of Positive Response</i>			Rare
	<i>Available + Ambiguously leaning yes</i>	<i>Available + Ambiguously leaning yes</i>	<i>Available + Available if + Ambiguously leaning yes + Available & more info</i>	First Names Excluded
Black	-0.072*** (0.019)	-0.084*** (0.019)	-0.093*** (0.019)	-0.074*** (0.020)
Positive Information	0.064*** (0.017)	0.054*** (0.017)	0.041** (0.017)	0.053*** (0.017)
Positive Information x Black	-0.046* (0.024)	-0.029 (0.024)	-0.016 (0.024)	-0.035 (0.025)
Negative Information	-0.318*** (0.019)	-0.342*** (0.019)	-0.296*** (0.020)	-0.347*** (0.018)
Negative Information x Black	0.029 (0.026)	0.041 (0.026)	0.026 (0.027)	0.040 (0.027)
% Blacks	0.017 (0.067)	0.021 (0.066)	-0.015 (0.066)	0.014 (0.067)
Black x %Blacks	-0.094 (0.099)	-0.097 (0.099)	-0.086 (0.099)	-0.076 (0.103)
Positive Information x %Blacks	-0.109 (0.082)	-0.125 (0.081)	-0.075 (0.082)	-0.118 (0.082)
Positive Information x Black x %Blacks	0.273** (0.124)	0.280** (0.124)	0.230* (0.124)	0.237* (0.133)
Negative Information x %Blacks	0.096 (0.093)	0.080 (0.094)	0.052 (0.096)	0.078 (0.093)
Negative Information x Black x %Blacks	0.020 (0.130)	0.029 (0.131)	0.094 (0.129)	0.005 (0.137)
Constant	0.556*** (0.015)	0.587*** (0.014)	0.619*** (0.014)	0.579*** (0.014)
Observations	14237	14237	14237	13007
R-squared	0.090	0.099	0.078	0.101

Notes: The omitted category is the baseline (no information) treatment for white. See definitions of variables in notes of Table 1, Table 2, and Table 3. Column (4) excludes three less common first names, Hakim, Rasheed, and Tremayne, which have within-gender frequencies below 0.005% in Census 1990. The results are similar if we exclude three Muslim sounding first names: Hakim, Karim, and Rasheed. Robust standard errors clustered by neighborhoods reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

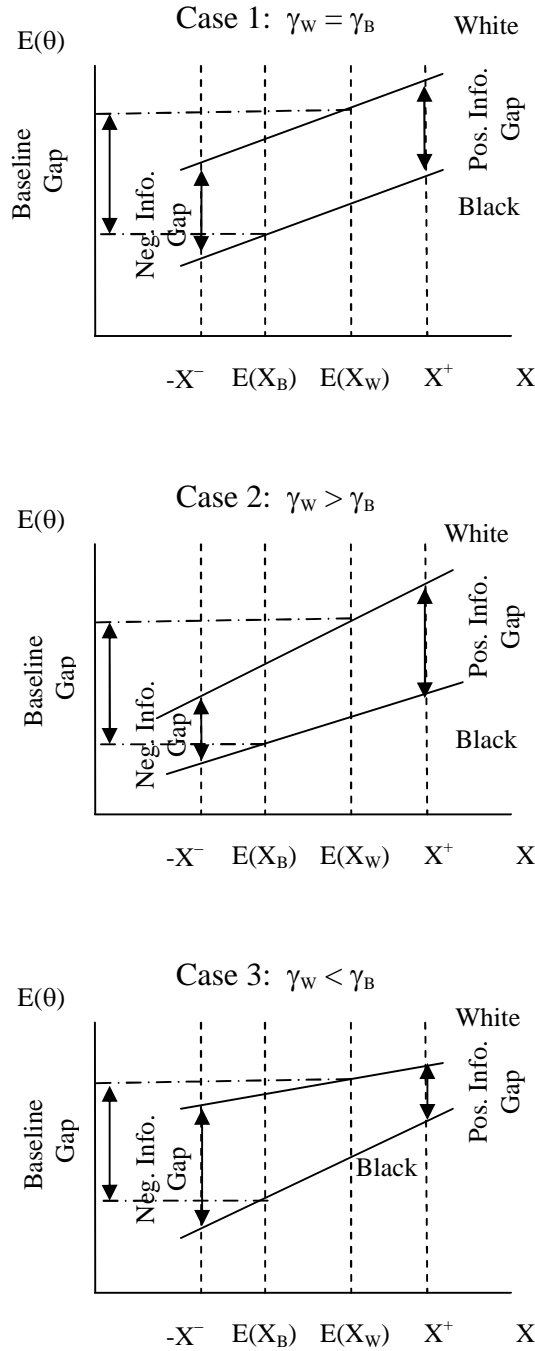
Table 9: Positive Response Rate and Mother's Education by First Name

White Female			White Male		
Name	% Positive Response	Mother Education	Name	% Positive Response	Mother Education
Jill	50.2	92.3	Todd	45.2	87.7
Carrie	50.3	80.7	Greg	45.5	88.3
Emily	50.3	96.6	Geoffrey	47.1	96.0
Kristen	50.4	93.4	Brett	47.5	93.9
Laurie	50.9	93.4	Matthew	49.7	93.1
Meredith	51.4	81.8	Brendan	50.8	96.7
Anne	51.6	93.1	Brad	51.0	90.5
Sarah	52.8	97.9	Neil	52.2	85.7
Allison	54.6	95.7	Jay	52.7	85.4
Correlation	0.477	($p = 0.194$)	Correlation	-0.300	($p = 0.433$)

Black Female			Black Male		
Name	% Positive Response	Mother Education	Name	% Positive Response	Mother Education
Latoya	37.0	55.5	Jamal	37.3	73.9
Tanisha	37.8	64.0	Tremayne	38.7	--
Ebony	42.6	65.6	Rasheed	40.2	77.3
Aisha	43.7	77.2	Hakim	40.5	73.7
Tamika	43.9	61.5	Kareem	41.4	67.4
Keisha	45.3	68.8	Leroy	41.4	53.3
Latonya	45.4	31.3	Tyrone	41.9	64.0
Lakisha	47.4	55.6	Jermaine	45.4	57.5
Kenya	47.7	70.2	Darnell	45.7	66.1
Correlation	0.100	($p = 0.798$)	Correlation	-0.762	($p = 0.028$)

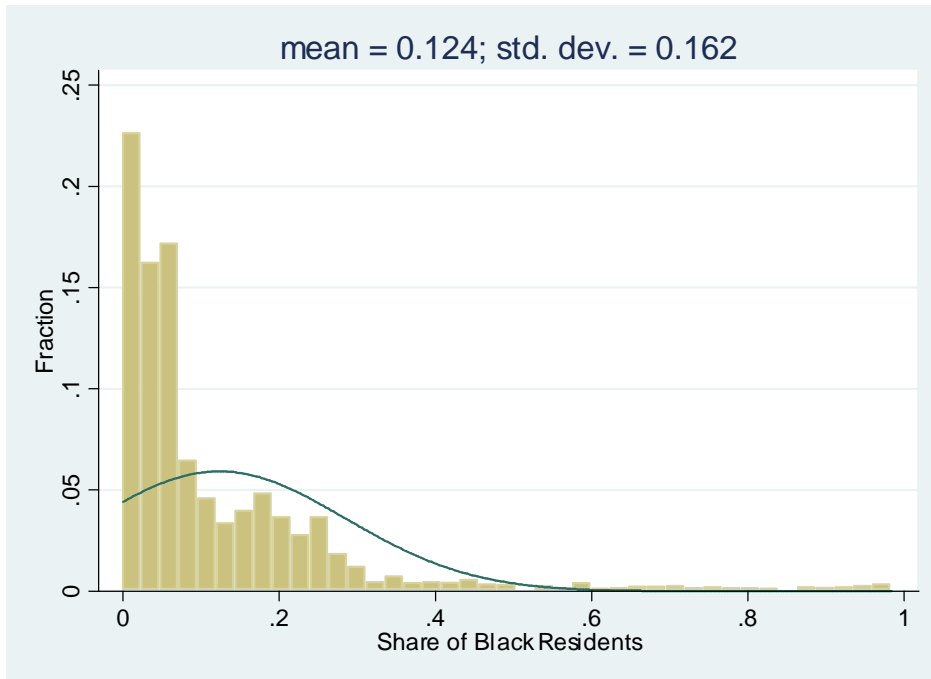
Notes: First names and mother education are sourced from Bertrand and Mullainathan (2004). Mother education is defined as the percent of babies born with that name in Massachusetts between 1970 and 1986 whose mother had at least completed a high school degree. "Correlation" reports the Spearman rank order correlation between positive response rate and mother education within each race-gender group, as well as the p -value for the test of independence (null hypothesis).

Figure 1 Shrinkage in Absolute Racial Gap and Information Weighting Parameters



Notes: We assume the preference parameter $K = 0$ to simplify the illustration. Case 1 shows shrinkages in racial gap for both positive and negative signals, comparing with the baseline treatment. Case 2 shows shrinkage in racial gap for negative signal only, comparing with the baseline treatment. Case 3 shows shrinkage in racial gap for positive signal only, comparing with the baseline treatment. The forecast equations for white applicants are arbitrarily placed above the forecast equation for black applicants to match stylized facts.

Figure 2: The Distribution of Shares of Black Residents across Census Tracts



Notes: Share of black residents is the number of non-Hispanic black persons divided all population resided in the census tract. For postings with missing addresses, we use metropolitan population figures. Data sourced from Census 2000.

Appendix A – Derivation of the Expected Value of Sample Variance of Signal

Equation (7) states that the denominator of the average information weighting parameter across a large sample of landlords is $(1/n) \sum \hat{\text{var}}(X_r)$. For a particular landlord, the sample variance of signal for racial group r is $\hat{\text{var}}(X_r)$. The mean of this landlord's sample variance of signal is:

$$E[\hat{\text{var}}(X_r)] = \text{var}(\theta) - \sum_i \sum_j \text{cov}(\theta_i, \theta_j) + \text{var}(\varepsilon_r) - \sum_k \sum_l \text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$$

$\text{cov}(\theta_i, \theta_j)$ is the pair-wise covariance of quality between individual tenant i and j for all $i \neq j$; $\text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ is the pair-wise covariance of the noise of signal between individual tenant k and l all $k \neq l$ in racial group r . If individuals are mutually independent, then $\text{cov}(\theta_i, \theta_j)$ and $\text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ are zero. However, neighborhood sorting means that landlords are likely to meet similarly individuals in neighborhoods in which they own properties and $\text{cov}(\theta_i, \theta_j)$ and $\text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ are not zero. In particular, $\sum_k \sum_l \text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ is positive and large for r if r is the majority group in the neighborhood, as there are more covariance terms. Thus, majority group of a neighborhood will have smaller $E[\hat{\text{var}}(X_r)]$. Whether $\text{var}(\varepsilon_r)$ is small, large, or constant across r is not really crucial to the relationship between neighborhood sorting, majority group, and the information weighting parameters.

Appendix B – Response Categories

Table A1: Response Categories

<i>Category</i>	<i>Description</i>
Available	The apartment is unambiguously stated as being available <i>and</i> future interaction is encouraged, i.e. a showing time is proposed or requested, they ask for future emails/phone-calls, etc.
Not Available	The apartment is said to be not available (unavailable), but no reason is provided as to why.
Not Available + reason	The apartment is said to be unavailable and a reason is given. The most common reason is that the apartment has already been rented.
Ambiguous leaning Yes	It is not clearly stated whether the apartment is available, but the language seems to indicate it is. i.e. “Thank you for your email. Feel free to call me whenever you like.”
Ambiguous leaning No	It is not stated whether or not the apartment is available, but the language seems to indicate it is not. i.e. “We may have other properties you are interested in become available.”
Disinterested	The landlord states the apartment is available but does not attempt to promote future contact/interaction. i.e. [Start of email] “The apartment is available.” [End of email].
Available + requirements	If any of the requirements were discussed/restated, such as: income, credit score, single resident only, no pets, full deposit, lease restrictions, etc.
Available + if	The unit is technically available, but an application has been submitted and the unit will only be available if this application falls through.
Available + more info	If the landlord requested more information concerning the quality of the tenant (i.e. not simply for their phone number): income, credit, number of residents, type of job, pets, etc.
Scam	A response which is clearly an attempt to obtain money or valuable information from the applicant.
Auto-reply	An automated response or “out of the office” reply that cannot be interpreted as any human response.
Blank	A response without anything in the body, which is likely an error due to email server.

Notes: Our preferred measure of positive response is “Available” & “Available + if”. Scams were all dropped from the sample.

⁺ Michael Ewens (mewens@cmu.edu), Tepper School of Business, Carnegie Mellon University; Bryan Tomlin (Brytomlin@gmail.com), Loyola University, Department of Economics; Liang Choon Wang (wchoon@gmail.com), Monash University, Department of Economics, and World Bank’s Development Research Group. We benefited tremendously from the suggestions and comments of Kate Antonovics, Gordon Dahl, Pushkar Maitra, Zahra Siddique, and two

anonymous referees. We also thank Eli Berman, Vince Crawford, Julie Cullen, Jacob LaRiviere, Valerie Ramey, and participants at UC San Diego seminars and the Western Economic Association International Annual Conference. We acknowledge the funding support from the Institute for Applied Economics at the UC San Diego. The findings, interpretations, and conclusions expressed in this paper are those of the authors and do not necessarily represent the views of the World Bank, its Executive Directors, or the governments it represents.

¹ For statistics on discrimination charges reported by the U.S. Equal Employment Opportunity Commission, see <http://www.eeoc.gov/eeoc/statistics/enforcement>. Statistics from the U.S. Department of Urban and Housing Development, for example, are available in various annual reports on fair housing at <http://www.hud.gov>

² See Arrow (1973) and Phelps (1972) for early discussions of statistical discrimination and Becker (1957) for taste-based discrimination.

³ See Heckman and Siegelman (1993) and Heckman (1998) for a critique on audit experiments that employ actors.

⁴ For more statistics about Craigslist popularity and users, see the experimental design section.

⁵ Several other papers show differential outcomes by race in the real estate and housing markets (Yinger 1986; Page 1995; Roychoudry and Goodman 1996; Ondrich et al. 1998, 1999; Ondrich et al. 2003; Ahmed and Hammarstedt 2008; Bosch et al. 2010; Carpusor and Loges 2006; and Hanson and Hawley 2010), automobile sales bargaining (Ayers and Siegelman 1995), labor market (Siddique 2008), TV game shows (Antonovics et al.'s 2005 and Levitt 2004), sportscard auctions (List 2004), and sales on eBay (Doleac and Stein 2010).

⁶ In practice, a landlord may actively search for other relevant signals of quality during face-to-face interviews (see Balsa and McGuire (2001) for a health care example).

⁷ We may assume that $X = r + \rho\theta + \varepsilon$, but it does not change the model predictions.

⁸ This is not strictly a “parameter”, but a landlord’s estimator of the parameter of the forecasting regression model.

⁹ See the section on research design for an example of such an inquiry.

¹⁰ This is equivalent to the landlord using some average θ for each race to form a prediction.

¹¹ The model can be generalized to a Bayesian framework, where priors about parameters of the forecasting regressions are updated with new experience. Furthermore, landlords may also update their prior about an applicant’s quality as new signals arrive, as in Altonji and Pierret’s (2001) example of employer learning. Since we focus on the initial stage of the screening process, we do not model how landlords update estimates and predicted quality about applicants over time.

¹² The assumption that the variance of quality (σ_θ^2) is the same across race is crucial for this interpretation. Note that $\text{cov}(\theta_r, X_r) = \text{var}(\theta_r) = \sigma_\theta^2$, given the assumption that $\theta_r \sim N(\mu_r, \sigma_\theta^2)$.

¹³ If $k = 0$, the landlord has no racial prejudice.

¹⁴ The fact that signals are not iid has no bearing on the landlord’s decision to use OLS, since the landlord cares only about getting the best linear prediction.

¹⁵ Appendix A provides a derivation of how differences in the variance of signals across different racial groups may arise.

¹⁶ Alternatively, the true variance of signals is larger for blacks than for whites.

¹⁷ After a pilot in June 2009, the experiment was conducted between 9/2009 and 10/2009.

¹⁸ The full (detailed) experimental design is available upon request.

¹⁹ Unique internet visitors are defined by a unique Internet Protocol (I.P.) address. A June 2009 report by AIMGroup shows a fall of newspaper classified ad revenue from \$16 billion in 2005 to

\$5 billion in 2009. Craigslist revenue grew from \$18 million to just over \$100 million over the same period.

²⁰ Approximately 2.5% (and growing) of all U.S. internet visits are to Craigslist, while other classified websites combined account for only 0.14% of U.S. internet visits.

²¹ The Institutional Review Board requires one inquiry per landlord so as to reduce potential harm and minimizes the likelihood of exposing the experiment to the landlord. Since treatments are randomly assigned, landlords are on average identical across groups.

²² Roughly one-third of postings do not contain cross-street information. These apartments are treated as located in the greater metropolitan area.

²³ White female, black female, white male, and black male, respectively. A full list of first names sourced from Bertrand and Mullainathan (2004) is listed in Table 9. The white surnames used in this study are Bauer, Becker, Erickson, Klein, Kramer, Mueller, Schmidt, Schneider, Schroeder, and Schwartz. These are surnames with highest fraction of whites among the top 500 most common surnames in Census 2000. The black surnames utilized are Washington, Jefferson, Booker, Banks, and Mosley, because these names more commonly belong to blacks than to other races among the top 1000 most common last names in Census 2000.

²⁴ We pooled two pieces of information together to increase the treatment effect.

²⁵ In contrast, revealing hard-to-verify characteristics such as social habit and cleanliness in advance is less realistic.