

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. rental apartment market that distinguishes statistical discrimination from taste-based 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 in a manner consistent with statistical discrimination and in contrast to patterns predicted by a model of taste-based discrimination.

JEL Codes: J15, J70, J71, R3.

* Affiliations: Tepper School of Business, Carnegie Mellon University; NERA Economic Consulting; Department of Economics, Monash University.

We benefited tremendously from the comments of Kate Antonovics, Gordon Dahl, Gordon Hanson, Pushkar Maitra, Birendra Rai, Zahra Siddique, and several anonymous referees. We also thank Eli Berman, Martin Bog, Vince Crawford, Julie Cullen, Ben Gillen, Jacob LaRiviere, Mike Price, Valerie Ramey, and Kunal Sengupta. We acknowledge the funding support from the Institute for Applied Economics at the UC San Diego.

I. Introduction

Racial and ethnic discrimination pervades many markets in the U.S. Roughly half of the discriminatory cases reported by federal agencies involve race or ethnicity, and the number of reported new incidents outpaced population growth over the past 10 years.¹ The 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 in which 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 outcomes sheds little light on the source of discrimination.

Employing an email correspondence experiment in the U.S. rental apartment market, we test whether statistical or taste-based discrimination can explain differential outcomes between white and African-American rental applicants. We extend Aigner and Cain's (1977) and Morgan and Vardy's (2009) models of statistical discrimination to test the key feature of statistical discriminators: heterogeneous experience.³ The model posits that landlords differ in their perceptions of signals due to past experience in the screening and rental process and place a greater weight on signals from the familiar group than the unfamiliar group in making decisions. We contrast the predictions with those of taste-based discrimination, in which prejudiced landlords use information independent of race to predict expected tenant quality but derive lower marginal utility from renting to the out-group. We show that lower marginal return to signal of quality for minority groups is consistent with both statistical and taste-based discrimination,

¹ For statistics on discrimination charges reported by the U.S. Equal Employment Opportunity Commission, see <http://www.eeoc.gov/eeoc/statistics/enforcement>.

² Arrow (1973) and Phelps (1972) first discuss statistical discrimination and Becker (1957) details prejudice.

³ Also see Cornell and Welch (1996) for a model of statistical discrimination where agents can better interpret signals from their own culture, race, or ethnicity.

highlighting the empirical hurdle involved in attempting to separate the two explanations. The model guides our experimental design.

Using vacancy listings on Craigslist.org (Craigslist), an online classified ad website, across 34 U.S. cities, we send inquiry emails with two key components to 14,000 landlords. We use the common, racial-sounding first names employed by Bertrand and Mullainathan (2004) to associate applicants with race, and the inquiry email contains differing – but limited – pieces of information about the applicants: positive, negative, and no signals beyond race. In the “no-signal” inquiry, landlords receive emails with racial-sounding names as the only signal. In the “positive information” inquiry, the fictional applicant informs the landlord that she is a non-smoker with a respectable job. In the “negative information” inquiry, the applicant tells the landlord she has a below-average credit rating and smokes. The dependent variable codes landlords’ responses to capture an invitation to the 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 influence final outcomes in the same direction. Since residential locations are tied closely to characteristics associated with welfare, such as the type of job held, crime levels, and school quality, our focus on the rental apartment market is policy relevant. As the dominant source of online classifieds for apartment listings in the U.S., Craigslist is frequented by one-third of the white and black U.S. adult population. The growing prevalence of online interactions in real estate, employment, finance, and auctions suggests the results extend beyond the rental apartment market.

The experiment yields four major results. First, when no-signal inquiry emails are sent, applicants with African-American sounding names have a 9.3 percentage point lower positive response than applicants with white sounding names. Second, using a difference-in-differences

(hereafter DiD) estimator, we show that the racial gap in response rates widens in the switch from negative to positive information. Both findings are consistent with statistical discrimination and prejudice, as they follow from landlords placing more weight on signals sent by white applicants than black applicants or landlords receiving a greater marginal utility of renting from white applicants than from black applicants. Thus, the research design requires further refinement to separate the explanations.

Third, the model defines the notion of a “surprising” signal, where the no-signal base case acts as a benchmark for uninformed expectations and as a means to quantify surprise relative to the better-than-expected (positive) and worse-than-expected (negative) signals. This notion of a surprising signal is difficult to introduce in a job application setting in which resumes are required, as it is impossible to provide “zero” information about education or experience in a resume. With differential weighting of signals by race, statistical discrimination predicts that a surprising positive signal will not necessarily shrink the racial gap in the base case, but that a negative surprise will. In contrast, the taste-based discrimination model shows that a surprising positive signal will widen the racial gap. Our empirical results are consistent with statistical discrimination.

Finally, we exploit neighborhood racial composition as a source of heterogeneity in landlord experiences with, or preferences for, different racial groups. By allowing a signal’s noise to depend on race, the statistical discrimination model presents another testable hypothesis: a landlord’s relative experience with a given race increases the relative weight she places on the signals from that racial group. Conversely, taste-based discrimination predicts that if landlords exhibit out-group prejudice and that if their race and racial preferences are correlated with neighborhood racial composition, as the share of blacks in a neighborhood increases,

discrimination against black applicants in the base-case will attenuate. Surprising negative information, in contrast, will hurt black applicants more. As the share of blacks in a neighborhood increases, a surprising positive signal closes the racial gap observed relative to the base-case, while a surprising negative signal does little to close it. More importantly, the base-case racial gap persists across all types of neighborhoods and contradicts the prediction of taste-based discrimination.

This paper extends the large body of research on racial discrimination. With the exception of List (2004) and Levitt (2004), past evidence of statistical discrimination is inconclusive. Altonji and Pierret (2001) and Bertrand and Mullainathan (2004) find significant racial gaps in wages and job interview callback rates, respectively, but weak support for statistical discrimination.⁴ On the other hand, the related literature on racial profiling in the context of police searches, such as Knowles et al. (2001) and Antonovics and Knight (2009), shows mixed evidence of racial prejudice. Audit studies, such as Yinger (1986) and Page (1995), show discrimination against minorities in U.S. housing markets. These results may suffer from the confounding factors inherent to in-person audit studies, as they rely on actors who often differ in many dimensions.⁵ In contrast, Carpusor and Loges (2006) and Ahmed and Hammarstedt (2008) pioneered the use of email correspondence design to study ethnic discrimination in rental housing markets.

Our approach differs from those of previous studies and contributes to the literature in several ways. We model explicitly how and why statistical and taste-based discrimination in the screening process can each predict lower marginal return to signal for the discriminated-against group, which potentially resolves some contradictory hypotheses and findings in the

⁴ For other environments, see Ayers and Siegelman (1995) (automobile sales), Siddique (2011) (labor market), Antonovics et al. (2005) (game shows), and Doleac and Stein (2010) (online sales).

⁵ See Pager and Shepherd (2008) for more examples and Heckman (1998) for a critique on audit experiments.

discrimination literature. We extend Aigner and Cain's (1977) and Morgan and Vardy's (2009) non-prejudice discrimination framework and juxtapose it with a prejudice discrimination framework to permit a research design that separates statistical and taste-based explanations empirically. Our framework also allows for landlord risk aversion and our findings are robust to this specification. In other empirical work, Ahmed et al. (2010) and Bosch et al. (2010) argue that statistical discrimination implies reduced discrimination against minorities with increased positive information. They show, however, an unchanging racial gap in the likelihood of positive responses when positive information is introduced in rental email inquiries and argue that the finding supports taste-based discrimination. Bertrand and Mullainathan (2004), in contrast, argue that the finding of lower or equal marginal returns to credentials for minorities is inconsistent with taste-based discrimination. On the other hand, Hanson and Hawley (2011) argue that decreased reply differences between white and African-American rental applicants from improving the prose quality of emails suggest statistical discrimination. Because reply rates decrease with class for white applicants, however, their results may be driven by the higher fraction of negative replies to low-class emails. Our novel modeling and estimation framework demonstrates the difficulty of separating statistical discrimination from taste-based discrimination in most field experiments.

II. Discrimination in Screening and Testable Implications

We present a model to guide our research design and to distinguish between statistical and taste-based discrimination in the rental apartment market.⁶ A landlord seeks to maximize the expected utility of interviewing each applicant, subject to a capacity constraint of M interviews

⁶ Although specific to the rental market, the model should extend to other situations of semi-formal screening.

and a constant marginal cost of interview c (i.e., a budget constraint of cM). The expected utility derived from each applicant depends on the stream of future rental income (i.e., tenant quality) from renting the apartment successfully. This quality is summarized by θ . Although the rent is preannounced, θ may still vary as a result of default, lease renewal, and so on. Hence, the landlord forms a predicted quality $\hat{\theta}_i$ (a random variable) for each applicant and maximizes the expected utility $E[U(\hat{\theta}_i)]$.

Consider the following four-stage process of matching potential tenants to apartments:

1. Inquiry: An applicant with quality θ selects publicly posted rental units to send costless inquiries with signal x to landlords.
2. Screening: Given signals $X_T = \{x_1, \dots, x_T\}$ received from T independent applicants, the landlord forms a set of predicted qualities $\Theta_T = \{\hat{\theta}_1, \dots, \hat{\theta}_T\}$ and responds to n applicants.
3. Interview: Interviews, which include credit and reference checking, reveal the true quality θ and cost of applicants and landlords m and c , respectively.
4. Decision: The candidate with the highest true quality θ is offered the apartment.

This setting is similar to that of Morgan and Vardy (2009), in which employers have one vacancy, only want “competent” workers, and receive signals about the workers’ types. That model has an additional stage in which true type is not revealed until the hire is complete. Our experiment can only distinguish statistical and taste-based discrimination occurring at stage 2 of the above process, and the results are consistent with Morgan and Vardy’s (2009) partial equilibrium, one-sided search approach. Nevertheless, we consider stages 2 through 4 in the model and discuss the implications of strategic signaling by applicants.

A. Statistical Discrimination

Suppose that signal x proxies the true quality θ noisily with a race-specific error ε_r :⁷

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

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 $\text{var}(x_r) = \sigma_\theta^2 + \sigma_{\varepsilon,r}^2$.

The assumptions mostly correspond to those in Aigner and Cain's (1977) model, except that we permit the signal mean to vary by race.

Landlords have a sample of inquiries x and applicants' true quality θ acquired during past iterations of stages 2 and 3 outlined above. Using this sample, a landlord can 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. Estimates of $\hat{\mu}_r^L$ and $\hat{\gamma}_r$ will vary across landlords because of their differing experiences. Equation (2) is similar to the posterior belief in Morgan and Vardy's (2009) or Balsa and McGuire's (2001) model. We assume, however, that landlords do not have prior beliefs about the true means and variances of quality and signal but rather form posterior beliefs using OLS estimates obtained via past samples. Indeed, Bayes's rule yields the same $\hat{\mu}_r^L$ and $\hat{\gamma}_r$ under a joint normality assumption.⁸

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

⁸ As landlords are not interested in the causal relationship between quality and signal, but only predictions that yield the lowest variance, OLS is sensible. It can be generalized to a Bayesian framework. For instance, landlords have prior beliefs that yield the intercept and slope of equation (2) and update them as they observe more realizations of quality and signal, as in Altonji and Pierret's (2001) example of employer learning. Our focus on the initial screening stage means that this generalization is not necessary.

A landlord observes signal \tilde{x} and race r from an applicant in stage 2 and predicts the quality by plugging \tilde{x} into (2):

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

We call $\hat{\gamma}_r$ the *information weighting parameter* for race r since it informs a landlord how much to weight a signal from an applicant of race r .⁹

Some applicants may reveal only their race in an email inquiry.¹⁰ In this case, the landlord infers quality using the average signal (\bar{x}_r) observed among race r in past email inquiries in the following forecasting regression:

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

Equation (4) is equivalent to the landlord using some average θ among r to form a prediction.¹¹

A.1 *Statistical Discrimination under Risk Neutrality and its Implications*

For a risk-neutral landlord with race-invariant utility, if the total number of applicants is $T \leq M$, the landlord responds positively to n of the T independent applicants, where each yields $E[U(\hat{\theta}_i)] \geq c$. If applications exceed capacity ($T > M$), then the landlord will sort all applicants by $E[U(\hat{\theta}_i)]$ and will invite the top M . As utility and cost are race-invariant, the decision rule is in line with Morgan and Vardy's (2009) color blind threshold. Here, the decision rule will be some $\underline{\theta}$.

As statistical discrimination influences a landlord's decision through $\hat{\theta}_r$, differential outcomes by race arise through the OLS estimators in equation (2):

⁹ This is not strictly a "parameter", but a landlord's estimator of the parameter of the regression model.

¹⁰ See section III for an example of such an inquiry.

¹¹ See subsection B for the implications when applicants strategically reveal no information.

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

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

Here, $\hat{\text{cov}}(\theta_r, x_r)$ is the sample covariance between quality and signal, $\hat{\text{var}}(x_r)$ is the sample variance of the signal, and $\bar{\theta}_r$ and \bar{x}_r are the sample average of quality and signal of applicants, respectively. Despite unobservability of $\hat{\mu}_r^L$ and $\hat{\gamma}_r$, we can, experimentally, manipulate race of applicants and signals sent by applicants to examine whether landlord responses are consistent with the model's predictions. If each landlord's sample of θ_r and x_r were observable, we could average the numerator and denominator of the information-weighting parameter across the sample of landlords (n_l) to obtain:

$$\gamma_r \equiv \frac{(1/n_l) \sum \hat{\text{cov}}(\theta_r, x_r)}{(1/n_l) \sum \hat{\text{var}}(x_r)} \quad (7)$$

Similarly, we have the average of the intercept term:

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

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

Given the assumption $E[\hat{\text{cov}}(\theta_r, x_r)] = E[\hat{\text{cov}}(\theta_{-r}, x_{-r})] = \sigma_\theta^2$, equation (7) shows that any differences in noise of signals, $E[\hat{\text{var}}(x_r)]$, can induce differences in the weight placed on the same signal from different races. For example, landlords with $E[\hat{\text{var}}(x_B)] > E[\hat{\text{var}}(x_W)]$ will have $\gamma_B < \gamma_W$.¹² For applicants with objectively identical signals except race, the predicted quality is greater for white than for black applicants. Note that we have not placed any restrictions on the mean signals or on noise across race. Differences in mean signal or noise could stem from

¹² Here, assuming quality variance (σ_θ^2) is the same across race is crucial. Note that $\text{cov}(\theta_r, x_r) = \text{var}(\theta_r) = \sigma_\theta^2$, given the assumption that $\theta_r \sim N(\mu_r, \sigma_\theta^2)$. We discuss the robustness of our results to this assumption in section V.

fundamental racial differences in signals such as income or credit scores of the population or simply from those observed by landlords in their samples.

With random assignment of race $r = \{W, B\}$ and in the absence of additional signals, we can examine whether landlord responses are consistent with an average landlord having $\bar{\theta}_W > \bar{\theta}_B$, $\bar{\theta}_W = \bar{\theta}_B$, or $\bar{\theta}_W < \bar{\theta}_B$. Equations (4) and (8) imply that:

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

Given past findings of discrimination against black applicants, we expect:

Hypothesis 1 – On average, a white applicant is more likely to receive a positive response than a black applicant in the no-signal base case.

If we assign a negative signal $-\tilde{x}^- < 0$ or a positive signal $\tilde{x}^+ > 0$ and race randomly to different applicants and present them to randomly selected landlords, we can use a DiD approach to test whether $\gamma_W > \gamma_B$, $\gamma_W = \gamma_B$, or $\gamma_W < \gamma_B$.

Equation (2) implies that the mean difference between black and white applicants sending a positive signal is:

$$E(\hat{\theta}_B | \tilde{x}^+) - E(\hat{\theta}_W | \tilde{x}^+) = (\mu_B^L - \mu_W^L) - (\gamma_W - \gamma_B)\tilde{x}^+ \quad (10)$$

Similarly, the mean difference between black and white applicants sending a negative signal is:

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

Taking the difference of equations (11) and (10) yields:

$$E(\hat{\theta}_B - \hat{\theta}_W | -\tilde{x}^-) - E(\hat{\theta}_B - \hat{\theta}_W | \tilde{x}^+) = (\gamma_W - \gamma_B)(\tilde{x}^+ + \tilde{x}^-) \quad (12)$$

The extent of dependence across signals of tenant quality which landlords obtained through their past experience influences the average sample variance of signal $\hat{\text{var}}(x_r)$.¹³ Massey and Denton's (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 dealing with white tenants than black tenants. These landlords' average sample variance of signals from white applicants will be smaller than that from black applicants because of neighborhood sorting.¹⁴ Since the average landlord in a national sample rents in a predominantly white neighborhood, we expect $\gamma_w > \gamma_B$. This prediction is also consistent with previous studies (e.g., Bertrand and Mullainathan [2004] and Ahmed et al. [2010]) that find smaller marginal returns to signal for blacks than for whites. So we have:¹⁵

Hypothesis 2 – On average, the positive response gap between white and black applicants is larger with a positive signal sent than with a negative signal sent.

If we (randomly) send a negative signal that is below the mean signal of race r , $-\tilde{x}^- < E(\bar{x}_r)$, or a positive signal that is above the mean, $\tilde{x}^+ > E(\bar{x}_r)$, to landlords, we can validate whether signals lead to differences in responses that are consistent with $\gamma_w > \gamma_B$.

Call the difference between the signal a landlord observes and her expected signal for the no-signal base case a “surprise in signal.” With an identical positive signal for black and white applicants and $E(\bar{x}_w) > E(\bar{x}_B)$,¹⁶ we have $[\tilde{x}^+ - E(\bar{x}_B)] > [\tilde{x}^+ - E(\bar{x}_w)]$. The experimentally

¹³ 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.

¹⁴ Appendix A shows how differences in the variance of signals across racial groups may arise.

¹⁵ We assume that applicants are truthful and that there is a separating equilibrium in signals by type.

¹⁶ As evident in past studies showing African-Americans having lower average social-economic backgrounds than white Americans (Harris, 2010).

manipulated negative information will be a greater surprising signal for whites than for blacks: $-\tilde{x}^- + E(\bar{x}_W) < -[\tilde{x}^- + E(\bar{x}_B)]$. Depending on the relative sizes of γ_W and γ_B , the patterns of landlord responses will differ.

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

$$\gamma_B[\tilde{x}^+ - E(\bar{x}_B)] > \gamma_W[\tilde{x}^+ - E(\bar{x}_W)] \quad (13)$$

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

That is, compared to the no-signal 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 (14) is unambiguously satisfied, but the relationship in (13) 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 (13) will be satisfied, but expression (14) will not necessarily be satisfied. 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.

The model implies that whether $\gamma_W > \gamma_B$, $\gamma_W = \gamma_B$ or $\gamma_W < \gamma_B$ for an average landlord depends on whether $\hat{\text{var}}(x_W) < \hat{\text{var}}(x_B)$, $\hat{\text{var}}(x_W) = \hat{\text{var}}(x_B)$, or $\hat{\text{var}}(x_W) > \hat{\text{var}}(x_B)$ on average. If the relative size of γ_r varies with $E[\hat{\text{var}}(x_r)]$ 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 $\hat{\text{var}}(x_B)$ to decrease and $\hat{\text{var}}(x_W)$ to increase on average. 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 (hypothesis 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). Thus:

Hypothesis 4 – Positive treatment should shrink the racial gap in positive responses 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.¹⁷

A.2 *Risk Aversion and Implications*

To assess the implications of landlord risk aversion, assume that the landlord's expected utility takes a mean-variance form (e.g., exponential utility). Here, the variance of the predicted

¹⁷ Previous discriminatory behavior might contribute to segregation and variation in neighborhood racial composition. If this discrimination is statistical, then using neighborhood racial composition as the proxy for landlords' relative experiences with different racial groups will lead to strong results, because landlords' experiences reinforced themselves in the samples observed.

quality generates another source of differential treatment by race. Based on equation (3), the variance of predicted quality conditional on a given signal \tilde{x} from a race- r applicant is:

$$\text{var}(\hat{\theta}_r | \tilde{x}) = \hat{\sigma}_\theta^2 \left(\frac{1}{N_r} + \frac{(\tilde{x} - \bar{x}_r)^2}{\sum_{j=1}^{N_r} (x_j - \bar{x}_r)^2} \right) = \hat{\sigma}_\theta^2 \left(\frac{1}{N_r} + \frac{(\tilde{x} - \bar{x}_r)^2}{N_r \hat{\text{var}}(x_r)} \right) \quad (15)$$

where N_r is the number of past race- r applicants that the landlord ever observed.

In the no-information base case, equation (15) shows that compared to a risk-neutral landlord, a risk-averse landlord further discounts applicant signals from the group with which she has relatively less experience. Risk aversion will thus widen the response gap between white and black applicants in Hypothesis 1.

When a black applicant sends a surprising positive (negative) signal, the applicant's conditional variance of predicted quality is larger (smaller) than that of a similar white applicant, holding all else equal, as $(\tilde{x} - \bar{x}_B)^2$ is larger (smaller) than $(\tilde{x} - \bar{x}_W)^2$ in equation (15). The larger is a positive (negative) surprising signal, the more (less) the black applicant is hurt from landlord risk aversion, making the effects of changing the signal on racial gap more pronounced. Thus, hypotheses 2 and 3 are similar with landlord risk aversion.

Since the effects of surprising signals on landlords' positive responses may vary depending on the extent of landlord risk aversion, neighborhood racial composition must only proxy landlords' relative experiences with different racial groups and not their degree of risk aversion. Otherwise, Hypothesis 4 will only examine the extent of landlords' risk aversion across neighborhoods under statistical discrimination, instead of identifying how relative experience with a particular race shapes the information-weighting parameter. Furthermore, as the share of black residents in a neighborhood (S_B) increases, N_B increases and $\hat{\text{var}}(x_B)$ decreases in equation (15). There is a possibility for positive responses to black applicants' surprising positive

signals to increase relatively less than positive responses to white applicants' surprising positive signals as $S_B \rightarrow 1$. Here, the extent of risk aversion and the effect from shrinking $\hat{\text{var}}(x_B)$ on the variance of predicted quality are large enough to offset the increase in the information-weighting parameter. We assess whether our results are sensitive to this possibility in section V.

B. Strategic Signaling and Truth-telling

The landlord's expected signal \bar{x}_r – assumed to be the average signal received in the landlord's history – depends crucially on the applicants' signal choices. If applicants instead send signals not directly tied to their type (e.g., strategically lie or not reveal), the interpretation of differences in response rates may need to be altered.

Let there be three applicant types: high, average, and low. Suppose in stage 3 of the rental process that applicants pay a credit check fee, landlords observe true applicant type, and reject any applicant who lied in the first stage. Here, applicants will never find it advantageous to lie about their type in the inquiry stage. Applicants who do reveal their type are still truthful, but may still decide to not signal type at all.¹⁸

Non-revelation of type mimics our “no information” case. High-type applicants will always find it advantageous to reveal their type to separate themselves, but low-type applicants can increase their response rate by non-revelation and by pooling with the average applicant type. Such a pooling equilibrium presents two problems for our analysis. First, the average signal \bar{x}_r is no longer the analogue of the average type, but rather the mean of low and average types. This lower average signal lowers the predicted response rate to no information and the interpretation of the gaps between it and other signals. Our predictions are robust to this scenario if landlords use the average of past no-information inquiries and interview results to forecast

¹⁸ The argument also applies to race revelation. If the average white applicant is of higher quality than the average black applicant, then the former will always reveal their race in inquiries.

quality. Second, under pooling it is no longer optimal for low-type applicants to send signals congruent with their type and thus the signals we send do not fit into the landlords' decision-making frameworks.

What does the low-type applicant's non-revelation strategy accomplish? With a higher expected response rate relative to full revelation, low-type applicants have increased the expected cost of interviewing. Suppose that some landlords would have rejected the low-type applicants if their signals matched their types. If there is a sufficient fraction of such landlords, low-type applicants will risk paying for interviews they are certain to fail. Sending a negative information email thus benefits the low-type applicant through more accurate inference of the landlord's likelihood of renting the apartment. Pooling adds significant noise and higher cost. We conclude that landlords treat negative information emails as coming from low-type applicants.

C. *Taste-Based Discrimination*

Differential outcomes by race may also arise from prejudice. Let a prejudiced landlord predict applicant quality based on a race-independent signal using a pooled OLS regression:

$$\hat{\theta}_i = \hat{\mu}_L + \hat{\gamma} \cdot x_i \quad (16)$$

The intercept and slope in equation (16) are race-invariant, which results in similar invariance for sample means and variances.¹⁹

Prejudiced landlords use the same decision rule for selecting applicants as above, but their utility is now race-dependent. Assume that the landlord exhibits out-group prejudice such that a prejudice parameter, k , discounts the (marginal) utility derived from an out-group applicant so that $E[U(\tilde{\theta}_r)] > E[U(\tilde{\theta}_{-r})]$ when $\tilde{\theta}_r = \tilde{\theta}_{-r}$. This framework contrasts with that of statistical

¹⁹ We thank an anonymous referee for this suggestion.

discrimination, where $E[U(\tilde{\theta}_r)] = E[U(\tilde{\theta}_{-r})]$ if $\tilde{\theta}_r = \tilde{\theta}_{-r}$. As the predicted quality has race-invariant variance, we focus solely on the implications of risk-neutral landlords.

We introduce a prejudice parameter into the landlord's utility function, which is similar to Knowles et al.'s (2001) race-varying marginal cost.²⁰ Although our pooled OLS assumption is more restrictive, it avoids additional assumptions about how prejudiced landlords incorporate information differently by race, while simultaneously not discriminating statistically.²¹ If a prejudiced landlord instead incorporates information differently for each race by penalizing the out-group with a smaller slope and a lower or equal intercept in equation (16), testable hypotheses similar to the pooled OLS regressions will still emerge.

C.1 Taste-based Discrimination Testable Implications

Given equation (16) and greater marginal utility for in-group applicants than for out-group applicants, if email inquiries reveal only the races of applicants, then:

Hypothesis 1A – On average, a white applicant is more likely to receive a positive response than a black applicant in the no-signal base case.

Next, greater marginal utility for white than for black applicants implies that an increase in signal benefits the white applicant more than the black applicant:

Hypothesis 2A – On average, the response gap between white and black applicants when a positive signal is sent is larger than the response gap between white and black applicants when a negative signal is sent.

²⁰ If we instead assume race-varying marginal cost, the testable hypotheses remain the same.

²¹ In addition to having a prejudice parameter on the marginal search (interview) cost, Knowles et al. (2001) allow the police's (landlords') beliefs about the probability of guilt (predicted quality) of motorists (applicants) with a certain characteristics (x) to differ by race, but they do not specify how landlords incorporate information differently by race, while simultaneously not statistically discriminating. We take the position that prejudice is a utility/preference assumption and incorporate it only into the utility function.

When a positive signal greater than the mean signal (\bar{x}) is sent, equation (16) combined with prejudice predicts that the average landlord's response gap between white and black applicants will widen relative to the base case. When a negative signal below the expected signal for applicants is sent, the landlord's response gap between white and black applicants will narrow. Therefore, the predictions differ slightly from those under statistical discrimination:

Hypothesis 3A – On average, negative information will unambiguously narrow the racial gap observed in the no-signal base case, but positive information will unambiguously widen the racial gap observed in the base case.

As the share of black residents in a neighborhood (S_B) increases, the probability that a landlord renting in that neighborhood is black also increases.²² This implies that the fraction of landlords having prejudice against black applicants decreases. Therefore, hypotheses 1A, 2A, and 3A will switch signs or directions as we move from a majority white neighborhood to a majority black neighborhood, leading to:

Hypothesis 4A – As the share of black residents in a neighborhood S_B increases, the response gap between white and black applicants in the base case decreases. In a majority black neighborhood, a surprising positive signal will unambiguously benefit a black applicant relatively more than a white applicant, while a surprising negative signal will unambiguously hurt a black applicant relatively more than a white applicant.

III. Experimental Design and Econometric Specifications

²² The correlation between the share of black homeowners and the share of black residents across public use micro areas is 0.96 in the American Community Survey 2009.

Craigslist serves as an ideal experimental platform to test the model implications because of its focus on email communication.²³ First, hypothesis 1(A) requires limiting the information available to agents to just race, which is straightforward in email correspondence but difficult in audit/in-person studies or other correspondence experiments. Next, hypotheses 2(A) and 3(A) demand clear signals that are also unambiguously distinct (i.e., positive vs. negative), which can be introduced in emails in a flexible way. Finally, the low cost of email and the popularity of Craigslist in the U.S. provide us with a large sample of agents.²⁴

A. *Experimental Subjects and Rental Market Data*

We use landlords who posted listings on Craigslist 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 in newspaper classified advertisement revenues.²⁵ Research firm Hitwise found that Craigslist receives 95% of visits to online classified websites.²⁶ Data from the Pew Internet & American Life Project reveal that roughly 44% of black and 49% of white Internet users have, at some point, used online classified advertisements like Craigslist (table 1). These Craigslist users, whether black or white, represent roughly one-third of the adult population in the U.S. Thus, findings based on Craigslist will be relevant for a large fraction of black and white adults, especially those using Internet and online classified advertisements. We eliminate scams, misplaced listings, repeated listings, and listings posted by individuals with “non-landlord” incentives. Those with non-landlord incentives

²³ The experiment was conducted between 9/2009 and 10/2009.

²⁴ The full (detailed) experimental design is available upon request.

²⁵ A report by AIMGroup shows newspaper classified advertisement sales fell from \$16 million to \$5million between 2005 and 2009, while Craigslist’s revenue grew from \$18 million to over \$100 million over the same period.

²⁶ Approximately 2.5% of all U.S. Internet visits are to Craigslist, while other classified websites combined account for only 0.14% of U.S. Internet visits.

include employees of large corporations managing dozens of apartments and private “apartment finders” who make a living as agents between landlords and renters.

Sampled apartments include only one-bedroom apartments and studios to avoid concerns about roommates, children, etc., and to ensure comparable rents between any two units within an area. Only one inquiry per listing was sent and numerous precautions were taken to avoid sending multiple inquiries to the same landlord and/or the same listing.²⁷ The sample excludes units with rents below the 20th and above the 90th percentile within each city to avoid sending emails to weekly rentals or homes for sale. Finally, emails were sent within 48 hours after the listing was posted. Table 2 lists the cities surveyed and their characteristics.²⁸ The average share of black residents across neighborhoods is similar to the actual share of the black population in the greater metropolitan area in Census 2000.

B. Email Generation and Experimental Treatments

To maximize the probability that landlords will observe the racial-sounding names, 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 used by Bertrand and Mullainathan (2004). Surnames are sourced from the U.S. Census 2000 Family Name Survey. Resulting name combinations include: Allison Bauer, Ebony Washington, Matthew Klein, and Darnell Booker.²⁹

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

²⁷ The Institutional Review Board requires one inquiry per landlord so as to reduce potential harm. Since treatments are randomly assigned, landlords are, on average, identical across groups.

²⁸ Roughly one-third of postings which lack cross-street information are placed in the greater metropolitan area.

²⁹ White female, black female, white male, and black male, respectively. A full list of first names and surnames used is in table 9.

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 (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 which to respond, allowing us to identify automated responses and to 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. Emails that do not include a statement of quality belong to the “baseline treatment” or “base case.” In this treatment, landlords only know the applicant’s name and her interest in renting the apartment. The model assumes that landlords take the average signals by race or independent of race as a proxy signal in the base case. When the statement of quality is included, it discloses either “positive” or “negative” information. Positive information informs the landlord that the applicant has a good job and does not smoke. Negative information informs the landlord that the applicant smokes and has a bad credit rating.³⁰ Both types are unambiguously positive or negative. This methodology does not aim to determine how any specific piece of information affects outcomes, but instead endeavors to test how positive or negative information, in general, affects outcomes.³¹ 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 creditworthiness and smoking habits in the interview stage and commonly ask applicants

³⁰ Revealing hard-to-verify characteristics such as habits and cleanliness is less realistic.

³¹ We pooled two pieces of information together to increase the treatment effect.

to pay for credit rating checks. As we discuss in section II.B, it is rational for some low-type applicants to reveal their type in the email inquiry because full revelation may reduce their cost of renting an apartment. 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 advertisement users reported in table 1, the negative information is likely surprising to landlords, providing a strong treatment effect.

C. Categorizing Outcomes

The simplest dependent variable identifies a landlord response. Responses were further classified into one of several categories, narrowed to either positive or negative.³² Positive responses state that the unit is available and invite future contact in some manner. Negative 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 who 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 being “disinterested.” Simple 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.

D. Econometric Specifications

We estimate four regression equations to test our hypotheses. First, the empirical specification to test hypotheses 1 and 1A is:

³² 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.

$$R_i = \alpha_W + \alpha_B B_i + u_i \quad (17)$$

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. We expect $\alpha_B < 0$.³³

Second, the following DiD specification tests hypotheses 2 and 2A:

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

N takes the value of 1 for negative information, and 0 for positive information. The omitted category is positive information for whites. If the average landlord weights signals from white applicants more heavily than signals from black applicants ($\gamma_W > \gamma_B$) or if the marginal utility for whites is greater than for blacks, we expect α_{NB} to be positive, resulting in a greater marginal return to signal for white applicants.

We estimate the following DiD regression to test hypotheses 3 and 3A:

$$R_i = \beta_W + \beta_B(B_i) + \beta_{PW}(P_i) + \beta_{PB}(P_i \times B_i) + \beta_{NW}(N_i) + \beta_{NB}(N_i \times B_i) + u_i. \quad (19)$$

The omitted category is no information for whites. 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. If hypothesis 3 is true, we expect that $\beta_{NB} > 0$ and that the sign of β_{PB} will be ambiguous. If hypothesis 3A is correct, we expect that $\beta_{NB} > 0$ and that $\beta_{PB} < 0$. Hence, failing to reject $\beta_{PB} \geq 0$ will cast doubt on taste-based discrimination.

Finally, the following empirical specification tests hypotheses 4 and 4A:

$$R_i = \delta_W + \delta_{SW}(S_{Bi}) + \delta_B(B_i) + \delta_{SB}(S_{Bi} \times B_i) + \delta_{PW}(P_i) + \delta_{SPW}(S_{Bi} \times P_i) + \delta_{PB}(P_i \times B_i)$$

³³ With all dichotomous regressors, we use a linear probability model (OLS) throughout (see Wooldridge (2003, pp. 456-7) for further discussion).

$$\begin{aligned}
& + \delta_{SPB}(S_{Bi} \times P_i \times B_i) + \delta_{NW}(N_i) + \delta_{SNW}(S_{Bi} \times N_i) + \delta_{NB}(N_i \times B_i) \\
& + \delta_{SNB}(S_{Bi} \times N_i \times B_i) + u_i
\end{aligned} \tag{20}$$

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) \bar{x}_r in the statistical discrimination model or the parameter k in the taste-based discrimination model to vary across different types of neighborhoods. If landlords' experiences with black applicants increase the size of their information-weighting parameter, then we expect δ_{SPB} and δ_{SNB} to be ambiguously signed. If landlords renting in a predominantly black neighborhood exhibit a preference for black residents relative to landlords renting in neighborhoods with a smaller share of black residents, then δ_{SB} should be positive. Moreover, if prejudice explains the observed racial differentials, the surprising positive information helps black applicants more and surprising negative information hurts black applicants more as S_{Bi} increases. So, $\delta_{SPB} > 0$ and $\delta_{SNB} < 0$.

IV. Results

Table 3 presents summary statistics of the experiment and data collected. Of the 14,237 inquiries sent, 9,229 (65%) received a response. Of these responses 6,597 (46%) were positive, as defined in section III. Figure 2 shows the distribution of the shares of black residents in census tracts of listed apartments (S_{Bi} in equation (20)). The measure ranges from 100% white to 98.45% black with a mean of 12.4% black residents. Table 4 verifies that the characteristics of our fictitious white and black applicants are statistically similar and not correlated with characteristics of listed apartments by treatment types.

A. Effective Informational Treatments

Table 5 reports response rates for the positive and negative treatments relative to the baseline of no signal, pooling all applicants. Comparing the intercept terms in column (1) and column (3) in table 5, which measure response and positive response rates respectively in the baseline treatment, reveals that roughly 18% 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) in table 5 show that landlords are equally likely to reply to an email inquiry, whether or not the applicant has revealed something positive about herself, but also that landlords are more likely to reply with a rejection if the applicant 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.

Column (3) in table 5 shows that applicants in the positive treatment group receive a significantly higher positive response rate than baseline applicants (0.57 vs. 0.53). The effect of positive treatment is slightly higher for females than for males (0.04 vs. 0.03). Column (4) in table 5 shows that applicants in our negative treatment receive a significantly lower positive response rate than baseline applicants (0.32 vs. 0.53). The differences illustrate that the treatments effectively manipulated landlord interest in the fictional applicants. Finally, the insignificant differences in response rates across gender and independent of race lead us to pool genders hereafter.

B. Hypotheses 1 and 1A: Black Applicants Receive Lower Response Rates

Column (1) in table 6 confirms hypotheses 1 and 1A that landlords, on average and without a signal of quality, are more likely to respond to applicants with white sounding names than applicants with African-American sounding names. The statistically significant coefficient on Black of -0.093 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 16% fewer positive responses.

C. Hypotheses 2 and 2A: Positive Information versus Negative Information

Column (2) in table 6 presents the DiD estimates for equation (18). It shows that the estimated effect of the negative treatment on black relative to white is significantly positive (0.042). It is consistent with 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 estimates also fail to reject hypothesis 2A that prejudice against black applicants can generate a larger response gap between white and black applicants as the signal switches from negative to positive.

D. Hypotheses 3 and 3A: The Effects of Surprising Signals

If statistical discrimination can explain differential outcomes for the hypothesized relative sign of the information-weighting parameter, then the surprising negative signal will close the racial gap in column (1) in table 6. Conversely, surprising positive information will not necessarily do so (hypothesis 3). If prejudice instead explains differential outcomes, then the racial gap will amplify with positive information and will close with negative information (hypothesis 3A).

In column (3) in table 6, the statistically negative coefficient on negative information and the statistically positive coefficient on "Negative Information x Black" are consistent with both hypotheses 3 and 3A. Column (3) in table 6 shows that disclosing negative information about an applicant's quality leads to a 50% greater reduction in a white applicant's probability of receiving a positive response.

Column (3) in table 6 also shows that the marginal return to signaling a respectable

occupation and non-smoking behavior increases the positive response rate by 6.7% (0.039/0.581) for whites. The coefficient on the interaction term “Black x Positive Information” is, however, statistically indistinguishable from zero. Both applicant types 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 is consistent with hypothesis 3, rather than hypothesis 3A, and conforms to the statistical discrimination model’s prediction that $\gamma_W > \gamma_B$. Although the positive information is much greater than the mean signal for black applicants, the relationship $\gamma_W > \gamma_B$ attenuates any improvement in signal. Taste-based discrimination cannot generate this finding as it predicts that a positive surprise will widen the racial gap unambiguously.

E. Hypothesis 4 and 4A: Differences by Information Types across Neighborhoods

The final prediction states that the difference $\gamma_W - \gamma_B$ is negatively related to the share of blacks (%Black) in a rental property’s neighborhood. An increase in this fraction equalizes the weighting parameters for black and white applicants. If prejudice instead explains the observed racial gap, then as the share of black residents and landlords increases, taste-based discrimination will lead to a shrinking racial gap in the baseline.

Column (4) in table 6 presents evidence that differential outcomes by race vary with the racial composition of an apartment’s neighborhood in a manner consistent with statistical discrimination. The 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. 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 decrease the racial gap between black and white applicants

significantly when compared with the baseline treatment case. Thus, in neighborhoods in which 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 hypothesis 4: landlords' relative past experiences with different racial groups shape their information parameters.

In contrast, column (4) in table 6 also reveals that the racial gap between white and black applicants in the base case does not vary across different neighborhoods. Furthermore, as the share of black residents increases, surprising negative information does not hurt black applicants more than white applicants, as indicated by the positive coefficient of “Negative Information \times Black \times %Black.” Our results contradict the predictions of taste-based discrimination.

Last, the results in columns (2)–(4) in table 6 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 posited in Bertrand and Mullainathan (2004). Overall, the DiD specifications and the significant coefficient on “Positive Information \times Black \times %Black” produces different conclusions about statistical discrimination than those of previous studies. With the significant treatment effect of positive information, our estimator can identify differences in γ across race and in turn illustrates that the lack of experience with a particular race influences the behavior of agents.

V. Robustness

The results are robust to several alternative specifications. First, our findings do not contradict the theoretical assumption of race-invariant quality variance nor do they support the

competing taste-based discrimination explanation. Second, we assume that landlords exhibit in-group favoritism and out-group animosity in the taste-based discrimination model. If black landlords are instead prejudiced against black applicants, then, as the share of black residents/landlords increases, the surprising positive signal should unambiguously widen the racial gap favoring white applicants. Column (4) in table 6 rejects this possibility.

Landlord risk aversion presents an ambiguity that could limit our ability to rule it out as the main source of the results. We exploit differences in neighborhood racial composition to separate statistical and taste-based explanations, so racial composition must not be correlated with landlord risk aversion. As risk aversion and wealth are highly correlated (Paravisini et al., 2010), neighborhood racial composition must not vary with property value. The high correlation (0.544) between average rent and property values of one-bedroom and studio non-farm units in the American Community Survey 2009 justifies rent as a proxy for value. The correlation between the share of black residents and the rent of the apartment is insignificant. In unreported estimates, we also find that the inclusion of both apartment rent (relative to the neighborhood average) and its polynomial have no significant effects on our estimates presented above. We conclude that risk aversion is not the main driver of our results.

Table 7 presents results using other definitions of positive responses. The estimates in columns (1)–(3) in table 7 show that the estimated coefficients vary little across all three measures of positive responses. The results are also robust to use of names. Column (4) in table 7 presents the estimation results of equation (20) without less common first names. All earlier conclusions about the testable implications remain.³⁴

Finally, the names chosen in our study could convey an applicant's social background

³⁴ The results are also robust to the exclusion of four Muslim sounding first names: Hakim, Jamal, Kareem and Rasheed.

beyond race. We follow Bertrand and Mullainathan's (2004) approach to examine whether average positive response rates are correlated with the 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 shows no evidence that positive response rate and social background are positively related (table 8).

VI. 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 and contrast the model's predictions with those of an alternative, taste-based model. We show some overlap between the implications of taste-based and statistical discrimination and illustrate that lower marginal return to signal or credentials for black applicants than for white applicants is consistent with both statistical and taste-based discrimination.

When no information other than the race of an applicant is revealed to landlords, applicants with African-American sounding names receive 9.3 percentage points fewer positive responses than applicants with white sounding names. The lack of a differential response to positive information casts doubt on taste-based discrimination as the dominant source of differential treatment. Landlord response rates across neighborhood racial compositions conform to the statistical discrimination model, in which agents use past experience to predict applicant quality by race. Racial prejudice or lexicographic search alone cannot explain these 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.

REFERENCES

- Ahmed, Ali, and Mats Hammarstedt, "Discrimination in the Rental Housing Market: A Field Experiment on the Internet," *Journal of Urban Economics* 64 (2008), 362–372.
- Ahmed, Ali M., Lina Andersson, and Mats Hammarstedt, "Can Discrimination in the Housing Market Be Reduced by Increasing the Information about the Applicants?" *Land Economics* 86 (2010), 79–90.
- Aigner, Dennis, and Glen G. Cain, "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review* 30 (1977), 175–187.
- Altonji, Joseph, and Charles R. Pierret, "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics* 116 (2001), 313–350.
- Antonovics, Kate, Peter Arcidiacono, and Randall Walsh, "Games and Discrimination: Lessons from *The Weakest Link*," *Journal of Human Resources* 40 (2005), 918–947.
- Antonovics, Kate, and Brian G. Knight, "A New Look at Racial Profiling: Evidence from the Boston Police Department," *Review of Economics and Statistics* 91 (2009), 163–175.
- Arrow, Kenneth, "The Theory of Discrimination," in Orley Ashenfelter and Albert Rees (Eds.), *Discrimination in Labor Markets* (Princeton, NJ: Princeton University Press, 1973).
- Ayres, Ian, and Peter Siegelman, "Race and Gender Discrimination in Bargaining for a New Car," *American Economic Review* 85 (1995), 304–321.
- Balsa, Ana I., and Thomas G. McGuire, "Statistical Discrimination in Health Care," *Journal of Health Economics* 20 (2001), 882–907.
- Becker, Gary S., *The Economics of Discrimination* (Chicago, IL: University of Chicago Press, 1957).

- Bertrand, Marianne, and Sendhil Mullainathan, “Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination,” *American Economics Review* 94 (2004), 991–1013.
- Bosch, Mariano, Lidia Farre, and Maria A. Carnero, “Information and Discrimination in the Rental Housing Market: Evidence from a Field Experiment,” *Regional Science and Urban Economics* 40 (2010), 11–19.
- Carpusor, Arian G., and William E. Loges, “Rental Discrimination and Ethnicity in Names,” *Journal of Applied Social Psychology* 36 (2006), 934–952.
- Cornell, Bradford, and Ivo Welch, “Culture, Information, and Screening Discrimination,” *Journal of Political Economy*, 104 (1996), 542–571.
- Doleac, Jennifer L., and Luke Stein, “The Visible Hand: Race and Online Market Outcomes,” SIEPR Discussion Paper 09-015 (2010).
- Hanson, Andrew, and Zachary Hawley, “Do Landlords Discriminate in the Rental Housing Market? Evidence from an Internet Field Experiment in U.S. Cities,” *Journal of Urban Economics* 70 (2011), 99–114.
- Harris, Angel L., “The Economic and Educational State of Black Americans in the 21st Century: Should We Be Optimistic or Concerned?” *Review of Black Political Economy* 37 (2010), 241–252.
- Heckman, James J., “Detecting Discrimination,” *Journal of Economics Perspectives* 12 (1998), 101–116.
- Iceland, John, Daniel H. Weinberg, and Erica Steinmetz, *Racial and Ethnic Residential Segregation in the United States, 1980-2000*. Census 2000 Special Report CENSR-3, US Census Bureau (2002).

- Knowles, John, Nicola Persico, and Petra Todd, "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy* 109 (2001), 203–229.
- List, John A., "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," *Quarterly Journal of Economics* 119 (2004), 49–89.
- Levitt, Steven D., "Testing Theories of Discrimination: Evidence from "Weakest Link"," *Journal of Law and Economics* 47 (2004), 431–452.
- Massey, Douglas S., and Nancy A. Denton, "Trends in the Residential Segregation of Blacks, Hispanics, and Asians: 1970-1980," *American Sociology Review* 52 (1987), 802–825.
- Morgan, John, and Felix Vardy, "Diversity in the Workplace," *American Economic Review* 99 (2009), 472–85.
- Pager, Devah, and Hana Shepherd, "The Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit, and Consumer Markets," *Annual Review of Sociology* 34 (2008), 181–209.
- Page, Marianne, "Racial and Ethnic Discrimination in Urban Housing Markets: Evidence from a Recent Audit Study," *Journal of Urban Economics* 38 (1995), 183–206.
- Paravisini, Daniel, Veronica Rappoport, and Enrichetta Ravina, "Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios," NBER working paper no. 16063 (2010).
- Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *American Economic Review* 62 (1972), 659–661.
- Ruggles, Stephen, J. Trent Alexander, Katie Genadek, Ronald Green, Matthew B. Schroeder, and Matthew Sobek, *Integrated Public Use Microdata Series: Version 5.0* [Machine-readable database] (Minneapolis, MN: University of Minnesota, 2010).

Siddique, Zahra, “Evidence on Caste-Based Discrimination,” *Labour Economics* 18 (2011),
S146–S159.

US Census Bureau, *Census 2000 Summary File 1* (2000).

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data* (Cambridge,
MA: MIT Press, 2002).

Yinger, John, “Measuring Racial Discrimination with Fair Housing Audits: Caught in the Act,”
American Economic Review 76 (1986), 881–893

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 calculations based on the Pew Internet & American Life Project's "April 2009 –Economy" survey data of the adult population. The sample includes non-Hispanic whites and blacks only. (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 use the Internet at least 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

Notes: (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 the population in the neighborhood; 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: 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 the notes of table 2. Black and male refer to applicants with African-American and male sounding names, respectively. 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 respectable occupation and non-smoking behavior to baseline email text. Neighborhood demographic characteristics are sourced from Census 2000. Response indicates whether a landlord responded and a positive response indicates whether a landlord responded positively to the inquiry. Positive response includes “available” and “available + if”. See online supplementary materials for response categories.

Table 4: 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.
<u>Pooled Gender</u>									
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 2 and table 3. Robust standard errors clustered by neighborhood reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 5: Overall Treatment Effects on Response Rate and Positive Response Rate

	(1)	(2)	(3)	(4)
	----- Response -----		----- Positive Response -----	
<u>Genders Pooled</u>				
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
<u>Males</u>				
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
<u>Females</u>				
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 2 and table 3. Robust standard errors clustered by neighborhood reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 6: 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 2 and table 3. Robust standard errors clustered by neighborhood reported in parentheses. Columns (1), (2), (3), and (4) correspond to hypotheses 1&1A, 2&2A, 3&3A, and 4&4A, respectively. *** p<0.01, ** p<0.05, * p<0.1

Table 7: 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 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 four Muslim sounding first names: Hakim, Jamal, Kareem, and Rasheed. Robust standard errors clustered by neighborhood reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 8: 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 percentage of babies born with that name in Massachusetts between 1970 and 1986 whose mother had at least completed a high school diploma. "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). The white surnames used in this study are Bauer, Becker, Erickson, Klein, Kramer, Mueller, Schmidt, Schneider, Schroeder, and Schwartz. These are surnames with the highest fraction of whites among the top 500 most common surnames in Census 2000. The black surnames used are Washington, Jefferson, Booker, Banks, and Mosley, because these names are more likely belong to blacks among the 1,000 most common surnames in Census 2000.

Illustration 1: Representative Email Samples

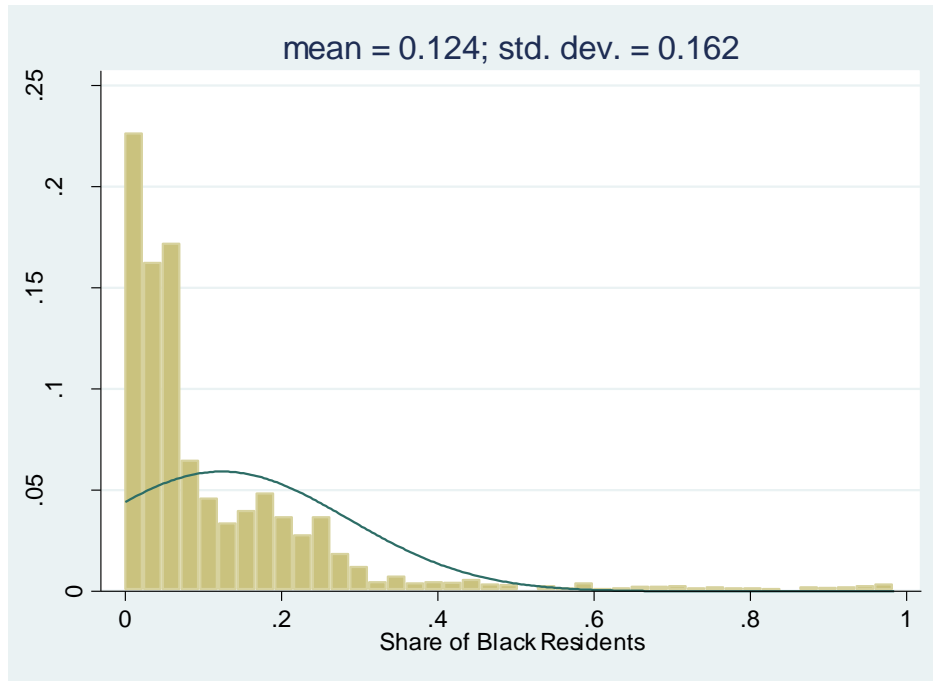
<i>Positive Treatment</i>	<i>Negative Treatment</i>
(1) Hello, (2) My name is [Full Name], and I am writing in response to your listing for an apartment for [apartment rent]/month. (3) In case you're interested, I do not smoke and I work full time as an architect. (4) Is this unit still available? (5) 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]

Figure 1 Shrinkage in Absolute Racial Gap and Information Weighting Parameters



Notes: We assume insignificant risk aversion to simplify the illustration. Case 1 shows shrinkages in the racial gap for both positive and negative signals, compared with the baseline treatment. Case 2 shows shrinkage in the racial gap for a negative signal only, compared with the baseline treatment. Case 3 shows shrinkage in the racial gap for a positive signal only, compared with the baseline treatment. The forecast equations for white applicants are placed arbitrarily 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 by all of the population residing 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 γ_r in a sample of landlords is $(1/n_r)\sum \hat{\text{var}}(x_r)$, where, for a particular landlord, the sample variance of signal for racial group r is $\hat{\text{var}}(x_r)$. In a large sample of landlords, the mean of the 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 the signal between the individual tenant k and l for all $k \neq l$ in racial group r . If individuals are mutually independent, then $\text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ is zero. Neighborhood sorting means, however, that landlords are likely to meet similar individuals in neighborhoods in which they own properties and $\text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ is not zero. $\sum_k \sum_l \text{cov}(\varepsilon_{rk}, \varepsilon_{rl})$ is positive and large for r if r is the neighborhood majority, as there are more covariance terms. Thus, the 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.