

12-1-2011

Subtle Discrimination in the Rental Housing Market: Evidence from E-Mail Correspondence of Landlords

Andrew Hanson

Marquette University, andrew.r.hanson@marquette.edu

Zackary Hawley

Georgia State University

Aryn Taylor

Ohio Wesleyan University

Accepted version. *Journal of Housing Economics*, Vol. 20, No. 4 (December 2011): 276-284. [DOI](#).

Andrew Hanson was affiliated with Georgia State University at time of publication.

Published under Creative Commons license [Attribution-NonCommercial-NoDerivatives 4.0 International](#).

Subtle Discrimination in The Rental Housing Market: Evidence from E-Mail Correspondence with Landlords

Andrew Hanson

*Department of Economics, Georgia State University,
Atlanta, GA*

Zackary Hawley

*Department of Economics, Georgia State University,
Atlanta, GA*

Aryn Taylor

*Department of Economics, Ohio Wesleyan University,
Delaware, OH*

Abstract: We find that landlords practice subtle discrimination in the rental housing market through the use of language associated with describing and viewing a unit, inviting further correspondence, making a formal greeting, and using polite language when replying to e-mail inquiries from a white name more often than to an African American name, they also send longer e-mails and respond quicker to white names.

Highlights:

- We test for subtle discrimination using matched-pair audit experiments.
- We examine e-mail text of landlord correspondence.
- Landlords use positive language to encourage whites and not African Americans.

- Landlords send longer e-mails and respond quicker to whites than African Americans.

Keywords: Discrimination; Field experiment; Housing; Race

I. Introduction

"work ref. rental ref. name address and ss#"

"its avail give me your # and I will have my daughter show it to you"

The quotes above are e-mail responses from the same landlord to separate inquiries for the same rental housing unit. Why the difference in response? They are replies to inquiries from (fictitious) men with different names – Tremayne Williams, and Brett Murphy, respectively – leading the landlord to believe the inquiry generating the first response came from an African American man, and the inquiry generating the second response came from a white man. The difference in these quotes represents unequal treatment of home-seekers that would not be uncovered in measures that only examine if a landlord responds to rental housing inquiries.

The quotes above suggest a form of subtle discrimination, defined by Charles (1999) as unequal treatment between groups that occurs but is difficult to quantify, and may not always be identifiable through common measures such as price differences. Subtle discrimination is particularly difficult to measure in housing market transactions where interaction (face to face, and telephone) cannot be practically measured on a large scale; nonetheless, landlords may use subtle discrimination to discourage minority clients from pursuing housing options. Traditional response/non-response measures in e-mail audit studies ignore subtle discrimination, and likely underestimate discrimination by mistaking treatment that appears equal for treatment that is different.¹

We contribute to the literature on racial discrimination in housing markets by using the text and timing of replies from an audit-style experiment conducted through e-mail contact with landlords to identify subtle forms of discrimination. Our work builds on work by

Hanson and Hawley (2011) that examines the response/non-response discrimination from landlords in the rental housing market.

We find that landlords reply faster, and with an e-mail that is longer to inquiries made from white names. We also find that landlords are more likely to use language associated with describing and viewing a unit, inviting further correspondence, use polite language more often, and make a formal greeting more often when replying to e-mail inquiries from a white home seeker. We find no statistically significant evidence (although individual cases sometimes show otherwise) that landlords use language associated with mentioning fees, asking for employment or rental history, or request background information as a means of discrimination.

The remainder of the paper begins by discussing the previous literature on housing market discrimination and places our work in context. Section 3 describes the experiment that generates the data we use for analysis. Section 4 details how we measure subtle discrimination through e-mail correspondence and Section 5 discusses the results. The final section of the paper concludes.

II. Previous literature on housing discrimination

Much of the previous literature on housing market discrimination focuses on quantifiable actions of real estate agents—showing the unit auditors originally inquired about, or a count of homes shown. Prominent examples in the literature include Yinger, 1986, Ondrich et al., 1989 and Zhao, 2005. These studies answer the question does discrimination occur in the housing market with an emphatic “yes.”

Many of these same studies also begin to answer the question of how discrimination occurs by taking a closer look at the interaction between real estate agents and potential clients. Ondrich, Stricker, and Yinger (1989) use the 1989 Housing Discrimination Study (HDS) to analyze discriminatory actions of real estate agents. The authors propose three main types of behaviors agents can use to discriminate against minority clientele: information given about which units are available, facilitating the sale of a unit, and the geographic location of housing units shown or recommended. They quantify the acts of discrimination by likelihood of brokers doing various actions—calling

back, asking about income, follow-up calls, asking about needs, financial assistance, invitations to inspect the advertised unit, and mentioning the advertised unit is available. The authors find real estate agents discriminate in all cases. These behaviors measure how discrimination occurs; however, there are still subtle differences in the way an agent can carry out those actions towards minorities that they do not quantify. For example, if a landlord makes a follow-up call to both white and minority auditors in the study, the agent may be more encouraging or use different tactics to entice the white client to visit again. The language used during these exchanges is not recorded, so it is impossible to say if there are subtle forms of discrimination happening in addition to the more overt discrimination.

Yinger (1998) discusses the evidence found on discriminatory actions in housing, automobile, and fast food markets. In the housing market, Yinger highlights significant discriminatory behavior in both the sales and rental markets. Agents exclude available units, show fewer units, make fewer call backs, signal less positive comments about the units, and are less likely to discuss financial incentives or help with minority home seekers. These results do indicate ways agents can discriminate, yet, within each action, an agent can still use subtle discrimination against minorities. The researcher has no way of knowing this type of discriminatory behavior since the actions of the agent are not recorded nor is the auditor asked to discuss these types of agent behaviors. Zhao (2005) uses the 2000 HDS to examine real estate brokers' choices of how many homes to show a potential client. He finds blacks and Hispanics are shown 30% and 10% fewer homes, respectively. This difference is mainly attributed to white customer's prejudice.

The research based on in-person audits is limited to what information is recorded at the time of the study by auditors. It is possible that the real estate agents in these studies also found more subtle ways to discourage their African American customers or encourage white customers through vocabulary, body language or other means. If subtle discrimination exists in addition to the quantifiable discrimination these studies find, they may underestimate the unequal treatment that actually occurs in housing market transactions. We extend the measure of unequal treatment by examining a precise record detailing the communication between

auditors and landlords. We use this record to test for differences in the language used to respond to each auditor, correspondence length (a measure of landlord effort), and the time elapsed between correspondences.

Previous housing market studies based on e-mail correspondence examine response/non-response discriminatory actions. Carpusor and Loges (2006) find Arab and African American sounding names are significantly less likely to receive a response from potential landlords. They briefly discuss how landlords choose to respond negatively to clients. In their study, when a negative response occurred, Arab-sounding names were statistically more likely to receive an overt response than whites—a response indicating the unit is not available. African Americans were also more likely to receive no response than whites.

Ahmed and Hammarstedt, 2008 and Ahmed et al., 2008, and Ahmed and Hammarstedt (2009) formulate e-mail correspondence studies to examine racial and sexual preference discrimination in the Swedish housing market. Each paper applies the standard response/non-response outcome as the discriminatory action in addition to a measure of the intensity of response, yet, they do not examine wording or amount of contact between landlords and auditors. Ahmed et al., 2010 and Bosch et al., 2010 employ e-mail correspondence studies where the information about the potential tenant is varied to examine if the additional information reduces discrimination. They find including positive information about the auditors reduces gross discrimination while net discrimination remains the same, but do not examine e-mail text. For a complete review on the use of field experiments to test for discrimination across all markets see Riach and Rich (2002).

III. Experiment design

The experiment is a matched-pair design or audit, where we send inquiries about available rental housing to landlords. Each landlord receives two e-mails- one sent from an e-mail address using a white name, the other sent from an e-mail address using an African American name. The experiment manipulates the racial group (African American or white) of home-seekers through the name associated with

an e-mail inquiry made to landlords who advertise rental housing on-line. Landlords are exposed to the name (race) of the home-seeker in three different ways: from the e-mail address, from the name-plate in the landlords' inbox, and in the signature of each e-mail.

The sample of names comes from Bertrand and Mullainathan (2004), who use Massachusetts birth certificate data from 1974 to 1979 to identify names that have a high likelihood of association with only one race. The first names used to represent white home-seekers are Brad, Brendan, Brett, Matthew, Neil, Geoffrey, Todd, Greg, and Jay. The first names used to represent African American home seekers are Darnell, Hakim, Jamal, Jermaine, Kareem, Leroy, Rasheed, Tremayne, and Tyrone. Last names also follow Bertrand and Mullainathan. The last names for white home-seekers are Davis, Ryan, Murphy, O'Brien, Baker, McCarthy, Young, Jones, and Wright. The last names used to represent African American home-seekers are Johnson, Washington, Robinson, Jackson, Hall, Parker, Williams, Jones, and Cooper.

To test for subtle discrimination we search the text of e-mail replies for keywords, measure the length (number of words) of replies, and the time elapsed between the inquiry and response. Table 1 details the keyword categories that we use to describe the content of each response. Our tests for subtle discrimination use the same experiment and data generated by Hanson and Hawley (2011) to test for response/non-response discrimination; we extend their analysis to examine the text of e-mail replies more carefully.

The venue for the experiment is the popular classified advertisement website Craigslist (<http://www.craigslist.org>); we use advertisements from Atlanta, Boston, Chicago, Dallas, Washington, D.C., Houston, Los Angeles, New York, Seattle, and San Francisco. Craigslist allows participants to place and reply to on-line advertisements specific to local markets for jobs, housing, companionship, and other goods and services. Landlords may create an advertisement for their property for no monetary cost, and home-seekers may reply to an unlimited number of advertisements for no monetary cost.² The audits use only listings pertaining to the rental housing market.

All e-mail inquiries are sent between 9 am and 12 pm on the day after a landlord posts an advertisement (always a Wednesday). Inquiries are sent from g-mail account addresses in the following format: *firstname.lastname.###@gmail.com*, where *###* is a three-digit number unique to each name. The experiment consists of 4728 audits, or 9456 e-mail inquiries. The overall response rate to e-mail inquiries is 53.9%, with 63.6% of landlords responding to at least one e-mail inquiry from a pair of e-mails. We limit our analysis to using 3153 audits from the original experiment that include one African American and one white name, the original experiment also includes audits with two white or two African American names with different e-mail quality.

IV. Measuring subtle discrimination from e-mail correspondence

An on-line venue provides a cost effective way to record the correspondence (or lack of correspondence) and timing of the correspondence landlords send in response to inquiries for rental housing. We use the actual text of e-mails sent by landlords in this paper as a way to test for subtle discrimination. This allows us to go beyond tradition response/non-response measures of discrimination and test for differences in how landlords respond across racial groups.

We know of no specific guide to measuring subtle discrimination, in fact the academic literature is so thin on this topic that researchers often rely on fictitious examples to define it (Rowe, 1990) or rely on self-reported survey data to measure its existence (Fox and Stallworth, 2005). There are a few studies of subtle discrimination in the labor market that offer some suggestions, including Levinson, 1975 and Bendick et al., 1991, and Bendick and Jackson (1994). These studies examine differences in the type of response and interactions actors have in matched-pair audits with potential employers on the phone and in-person. These studies measure subtle discrimination by the incidence of positive and negative comments made during a job interview, if interviewers made discouraging comments about the applicants chances of being hired, if interviewers were polite, the length (in minutes) of the interview, and the topics discussed during the interview. In the housing market,

Ondrich, Stricker, and Yinger (1989) examine qualitative face-to-face real estate agent and client interactions, broken into two broad categories: unit availability and sales facilitation, we incorporate each of these measures into our search for subtle discrimination.

We categorize our correspondence measures into three broad categories: content of response, time to response, and length of response. For time to response, we measure both the average length between inquiry and response and the percentage of landlords who take longer to respond to each race. For the length of response, we measure the number of words contained in the response, and the percentage of landlords who write a longer reply to each race.

Analogous to the labor market studies, we separate the content of landlord responses into positive and negative categories, and do keyword searches to determine the incidence of differential treatment. We further divide our searches into sub-categories to reflect both differential treatment associated with unit availability and sales facilitation as examined in Ondrich, Stricker, and Yinger. Our positive categories measure favorable landlord treatment along several dimensions including if landlords describe the unit in a positive way, if they mentioned other available units, if they invited further contact, and if they used generally friendly language. Negative search categories measure landlord treatment along dimensions that focus on responding unfavorably to one group, including if landlords are more likely to mention fees, or ask for more information from a prospective tenant such as questioning their employment or background.

To objectively measure the content of response, we examine the text of each e-mail by using automated searches for keywords. For the "descriptive" sub-category we searched for terms commonly used to describe the housing unit in positive terms. For instance, a landlord may describe "new" carpet, or a "clean" neighborhood. The "other units" sub-category intends to capture a landlord's willingness to provide information about additional rental units that may be available. The sub-category "view unit" combines searches for several terms associated with a landlord making an invitation to view the unit. "Email" and "Phone" capture the landlords inclusion of language associated with invitation to make further contact. The "greeting"

category includes search terms for friendly greetings, and the “polite” category includes search terms for friendly valedictions.

Searches for negative language focused on requests for more information or mentions of fees. The sub-category “fees” combines searches for words that mention application fees or a deposit amount or requirement. “Employment” is a sub-category search for terms that ask (or imply) the landlord is looking for some verification of employment. The “history” sub-category encompasses a landlord’s inquiry for more information about the rental or eviction history of the tenant. “Background” combines searches for terms that suggest the landlord is asking for a criminal background, or other verification.

Table 1 details the sub-categories for positive and negative language and describes the search terms in each sub-category. For all sub-category searches we count whether a particular e-mail includes any of the terms in each sub-category, we do not account for the use of multiple search terms in one e-mail. For keyword search tests, we measure differences at the landlord level. Differences at the landlord level measure differential treatment by the same person for the same property. This measures net discrimination, or the difference in the proportion of landlords that favor whites against those that favor African Americans. These results are equivalent to running Probit regressions with landlord fixed effects, but we present them here in terms proportions as it also shows the percentage of landlords who practice equal treatment. For the time elapsed and word count tests, we examine differences at both the e-mail level (measuring the difference between group means) and the landlord level (measuring the difference in proportion of landlords who respond longer/faster to the white name).

V. Results

Content of response

The positive keyword search terms, detailed in Table 1, are more likely to appear in replies to white names, although the majority of landlords treat the inquiries the same– they do not use the keywords at all, or they use them in both e-mails. Table 2 shows the

results of McNemar paired difference in proportions tests for our keyword searches.³ The largest difference occurs in the use of a greeting, where we find about 2.4% of landlords discriminate on net. We find that landlords are more likely to use polite language such as "Thanks" or "Please Call" when replying to whites than African Americans. The net level of discrimination for the use of polite language is about 2% of landlords.

We find that landlords are more likely to offer contact information in the form of a phone number or e-mail address (many landlords initially list apartments using an anonymous Craigslist e-mail address, so it is meaningful to list a different means of contact) when corresponding with white names. The magnitudes suggest that about 1.5% of landlords favor whites for these categories, statistically significant at less than the 1% level. Unequal treatment exists when offering to view the unit, as we find landlords are more likely to mention words like "view", "tour", or "show" to white tenants. Again, this difference is statistically significant at the 1% level.

We also find landlords are more likely to use descriptive language in e-mails to whites than they are to African Americans. This includes terms such as "new", "clean", and "quiet". The size of this difference suggests about 1% of landlords discriminate on net with this measure. We do not find any difference in suggestions that other units are available, perhaps because this could be viewed as either a positive (if the current unit has been rented) or negative (if steering the African American to another area), so that the net effect is zero.

While the positive search term results clearly show favorable treatment toward white names, the negative search term results do not suggest strong discrimination. Table 2 shows negative search term results (Fees, Employment, History, Background), broken down by the categories described in Table 1. Despite the egregious example in the introduction, we do not find statistically significant subtle discrimination through the use of asking about employment verification, rental history, or the tenant's background. We also find no evidence of unequal treatment when mentioning fees or deposits between replies sent to white and African American names.

The keyword searches reveal the majority of landlords respond using similar language to both inquiries; however our tests show that discrimination still exists in the text of many replies. Net discrimination, measured as the difference between landlords that favor whites and landlords that favor African Americans, ranges from about 1% (using descriptive language) to about 2.4% (using a greeting) – always in favor of the white names. The discrimination toward African American names is statistically significant in all cases except the “other units” category.

As a robustness check to the search terms listed in Table 1, we combined several alternative searches on the text of landlord replies and tested for differences between African American and white names. Table 3 shows the results of the alternative searches. They include searching for any negative or positive terms, combining the e-mail and phone categories, combining the other and view unit categories, and combining the greeting and polite categories. All of these searches yield similar results to the primary subtle discrimination tests in Table 2. We find no difference in the use of negative search terms, but some differences in the use of positive search terms. The magnitude of the difference for our robustness checks is larger than the magnitude in our original searches, suggesting that extending the search to include more positive terms reveals more landlords who potentially practice subtle discrimination.

We further extend the robustness checks to include a more extensive list of descriptive terms (including mentioning positive attributes of the surrounding neighborhood like schools, shops, and stores). These results are quite similar to the original descriptive search results, with only a slightly larger magnitude. We also search for terms that suggest the landlord is prompting action on the part of the potential tenant (quick, hurry, fast, come, visit, application, won't last), but find no statistically meaningful difference. Finally, we searched for discouraging language such as “already rented”, “no longer” “not available”, but found no differences across races. The notes of Table 3 detail the exact search terms we use for all robustness checks.

The “All Positive” and “All Negative” searches in Table 3 show the percentage of landlords who use any of the terms in our search

categories. The level of net discrimination for this category is 4.73% of landlords, or close to double the magnitude of the largest individual category. This is evidence that African Americans are more likely to encounter some form of subtle discrimination than they are any particular form. We can also determine if the same landlords discriminate in multiple ways (appearing in multiple categories) or different landlords discriminate in a few ways. Table 4 shows the overlap in categories for landlords who are net discriminators – those who reply using only positive language to whites, or only negative language to African Americans. We find that of the landlords using positive language only in e-mails to whites, the average number of categories is 2.14, and about 34% only appear in one category, while 66% appear in two or more. Landlords using negative language only in e-mails to African Americans appear in 1.3 negative categories on average, and about 72% of them appear in only one category.

Time to response

Table 5 shows the time to response results at the e-mail and landlord level. Row 1 of Table 5 shows that the average time elapsed between the inquiry and subsequent landlord reply was 7 h and 9 min for e-mails from white names, and 7 h and 48 min for e-mails from African American names – a difference of 39 min. Column 3 of Table 5 shows the results of a standard difference in means test that suggests this difference is statistically significant at just above the 5% level (p -value of .0656).⁴

The landlord level results also show that responses come quicker to e-mails sent from white names. These results, shown in rows 2 and 3 of Table 5, control for all individual landlord level characteristics and report the percentage of landlords who replied in the same amount of time (to the nearest minute), took longer to reply to the white name, or took longer to reply to the African American name. Row 2 of Table 5 reports results using all audits, while row 3 uses only audits where the landlord replied to both e-mails.

Using all audits (row 2 of Table 5), the landlord level results show that almost 37% of landlords take the same amount of time to reply to both e-mails. The large percentage of landlords who take equal time in replying is driven by the “zero” length of time for

landlords who do not reply to either e-mail.⁵ 28% of landlords take longer to reply to the white name, while almost 35% take longer to respond to the African American name for a net level of discrimination of about 7% of landlords. Column 7 shows that this difference is statistically significant at less than the one-percent level using a McNemar paired difference in proportions test.

Using only landlords who replied to both e-mails (row 3 of Table 5) shows that less than 1% of landlords actually took the same amount of time to reply to both inquiries (these are due to automatic reply e-mails and responses that came within the first minute after inquiries were sent). 47% of landlords who replied to both e-mails took longer to respond to white names, while over 52% took longer to respond to African American names, a net level of discrimination of just over 5% of landlords. As shown in column 7, this measure is statistically significant at less than the 1% level using the McNemar paired difference in proportions test.

Length of text

Table 6 shows word count results for the text of landlord replies to inquiries about advertised rental housing units. The first three columns in Row 1 of Table 6 shows average word counts from all e-mail replies, and include counting a zero length for non-responders. The average word count for white names is 25.86 words, as opposed to 23.70 for African American names- a difference of 2.16 words. Column 3 shows the results of a difference in means test that the difference is statistically significant at less than the 1% level. Row 3 of Table 6 shows that this difference increases to 3.14 words when we do not include e-mail responses that were the exactly the same length (mostly zero counts for non-response and some automatic reply e-mails), and that the difference remains statistically significant at less than the 1% level using standard difference in means z tests.

The last four columns of Table 6 shows word count results at the landlord level, examining the percent of landlords who make either an equal length reply or a longer reply to one race. Row 1 shows the results using all audits, while Row 2 shows results using only landlords who made a reply to both e-mails. The Row 1 results include counts of zero words for non-response and show that close to 47% of landlords

reply with the same number of words to both inquiries. Our results also show that almost 30% of landlords reply with more words to inquiries from white names, while only 23.5% reply with more words to inquiries from African American names. Net discrimination in length of reply occurs for about 6.5% of landlords, and this difference is statistically significant at less than the 1% level using the McNemar paired difference in proportions test.

Row 3 of Table 6 shows that the net measure of discrimination is smaller, about 4% of landlords, when only using landlords that actually replied to both inquiries. The difference is still statistically significant at less than the 5% level using the McNemar paired difference in proportions test. This measure excludes the landlords that already discriminated by not replying to one of the inquiries, and suggests even landlords that reply to both e-mails still practice discrimination. Despite dropping all of the non-responses, we still see that over 23% of landlords send an e-mail with the same number of words to both inquiries.

VI. Discussion and conclusion

We find that landlords favor whites by responding quicker, writing e-mails that are longer, and using more positive language when replying to inquiries about rental housing. On balance, the keyword searches show that landlords choose to discriminate by encouraging white tenants through positive language instead of discouraging African American tenants through negative language. Although the experiment does not offer insight as to why this may be the case, one reason may be the threat of legal action. The Fair Housing Act explicitly states that it is illegal for landlords to discriminate on the basis of terms and conditions of the housing unit (exemplified by the "fees" subcategory). If landlords fear that explicitly discouraging African Americans through mention of fees or background checks may cause them legal trouble, they may substitute toward encouraging white tenants with positive language.

To put our results in context, Hanson and Hawley (2011) find 4.5% of landlords discriminate against African Americans on net in terms of response/non-response. Depending on the measure we use for subtle discrimination, an additional 1.08–6.78% of landlords

practice subtle discrimination on net. One way of viewing these results is that the traditional response/non-response measure underestimates the share of landlords who practice unequal treatment, although they may do it in less egregious ways than not responding to potential minority tenants. Another way of viewing these results is that they represent a different dimension of discrimination that is more difficult to uncover. It seems unlikely that all types of subtle discrimination we uncover hinder housing search as much as a non-response; nevertheless, they can discourage potential renters and restrict choice, and should be considered in the literature moving forward.

The question remains: why do landlords discriminate against African Americans in the rental housing market? There is a “statistical” explanation, that landlords discriminate because it is in their best interest in terms of profit maximization, and they perceive whites to be more worthy tenants. There is also the taste-based view, that landlords have some personal preferences for tenants of one race, and they act on these feelings. Although the source of discrimination is important to understand, we believe our work is limited to describing how discrimination occurs, not why. For instance, landlords using a greeting only for white tenants could be perceived as practicing taste-based discrimination, because a greeting is a way of being friendly with someone and may be indicative of personal preference. The same result, however, could also be explained as statistical discrimination—landlords are friendlier with whites because they believe them to be better tenants, so they try to attract them to the unit.

Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant No. 0754791. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Ahmed and Hammarstedt, 2008. A.M. Ahmed, M. Hammarstedt. Discrimination in the rental housing market: a field experiment on the internet. *Journal of Urban Economics*, 64 (2) (2008), pp. 362–372

- Ahmed and Hammarstedt, 2009. A.M. Ahmed, M. Hammarstedt. Detecting discrimination against homosexuals: evidence from a field experiment on the internet. *Economica*, 76 (2009), pp. 588–597
- Ahmed et al., 2008. A. Ahmed, L. Andersson, M. Hammarstedt. Are lesbians discriminated against in the rental housing market? Evidence from a correspondence testing experiment. *Journal of Housing Economics*, 17 (3) (2008), pp. 234–238
- Ahmed et al., 2010. A. Ahmed, L. Andersson, M. Hammarstedt. Can discrimination in the housing market be reduced by increasing the information about the applicants? *Land Economics*, 86 (1) (2010), pp. 79–90
- Bendick and Jackson, 1994. M. Bendick, J. Charles. Measuring discrimination through controlled experiments. *Review of Black Political Economy*, 23 (1) (1994), pp. 25–48
- Bendick et al., 1991. M. Bendick, J. Charles, R. Victor, H. Laura. Discrimination against latino job applicants: a controlled experiment. *Human Resource Management*, 30 (4) (1991), pp. 469–484
- Bertrand and Mullainathan, 2004. M. Bertrand, S. Mullainathan. Are emily and greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94 (4) (2004), pp. 991–1013
- Bosch et al., 2010. M. Bosch, M.A. Carnero, L. Farré. Information and discrimination in the rental housing market: evidence from a field experiment. *Regional Science and Urban Economics*, 40 (1) (2010), pp. 11–19
- Carpusor and Loges, 2006. A.G. Carpusor, W.E. Loges. Rental discrimination and ethnicity in names. *Journal of Applied Social Psychology*, 36 (4) (2006), pp. 934–952
- Charles, 1999. Charles, K., 1999. *A Simple Model of Subtle Discrimination*. Unpublished manuscript.
- Ewens et al., 2009. Ewens, M., Bryan T., Choon W., 2009. *Statistical Discrimination in the US Apartment Rental Market: A Large Sample Field Study*.
<<http://dss.ucsd.edu/~btomlin/pdfs/ewenstomlinwang.pdf>>.
- Fox and Stallworth, 2005. S. Fox, S. Lamont. Racial/ethnic bullying: exploring links between bullying and racism in the US Workplace. *Journal of Vocational Behavior*, 66 (3) (2005), pp. 438–456
- Hanson and Hawley, 2011. A. Hanson, H. Zackary. Do landlords discriminate in the rental housing market? Evidence from an internet field experiment in US cities. *Journal of Urban Economics*, 70 (2–3) (2011), pp. 99–114

- Levinson, 1975. R. Levinson. Sex discrimination and employment practices: an experiment with unconventional job inquires. *Social Problems*, 22 (4) (1975), pp. 533–543
- Ondrich et al., 1989. J. Ondrich, A. Stricker, J. Yinger. Do real estate brokers choose to discriminate? Evidence from the 1989 housing discrimination study. *Southern Economic Journal*, 64 (4) (1989), pp. 880–901
- Riach and Rich, 2002. P.A. Riach, J. Rich. Field experiments of discrimination in the market place. *Economic Journal*, 112 (483) (2002), pp. 480–518
- Rowe, 1990. M. Rowe. Barriers to equality: The power of subtle discrimination to maintain unequal opportunity. *Employee Responsibilities and Rights Journal*, 3 (2) (1990), pp. 153–163
- Yinger, 1986. J. Yinger. Measuring racial discrimination with fair housing audits: caught in the act. *The American Economic Review*, 76 (5) (1986), pp. 881–893
- Yinger, 1998. J. Yinger. Evidence on discrimination in consumer markets. *The Journal of Economic Perspectives*, 12 (2) (1998), pp. 23–40
- Zhao, 2005. B. Zhao. Does the number of houses a broker shows depend on a homeseeker's race? *Journal of Urban Economics*, 57 (1) (2005), pp. 128–147

Notes.

- ¹ Carpusor and Loges, 2006, Ahmed and Hammarstedt, 2008, Ewens et al., 2009 and Bosch et al., 2010 and Hanson and Hawley (2011) are examples of e-mail correspondence studies that test for discrimination in housing markets using response/non-response differences.
- ² Craigslist puts a notice on every advertisement posted on the site warning that stating a discriminatory preference is illegal and allows users to easily notify them of any such advertisements.
- ³ For all statistical significance tests at the landlord level (Table 2 and Table 3, and columns (4–7) of Table 5 and Table 6) of the difference in proportion of landlords between white and African American tenants, we use the McNemar test. This test is designed for testing the difference in proportion of respondents for paired subjects, the test statistic is $\chi^2 = (N_{\text{only W}} - N_{\text{only AA}})^2 / (N_{\text{only W}} + N_{\text{only AA}})$, where N represents the number of landlords only using language for one group. The test statistic has a chi-squared distribution, and we calculate all p -values accordingly.
- ⁴ This test is at the e-mail level, average times do not include non-responses. E-mail level tests are not paired tests, they are standard difference in means z tests.
- ⁵ 1151 of the 1161 equal time responses were no reply or 'zero' length.

Table 1. Keyword groups used for e-mail correspondence searches.

Positive language categories

Descriptive	New, new	Clean, clean	Quiet, quiet	Nice, nice	Good, good	
Other units	Another, another	Second, second	Several, several			
View unit	View, view	Tour, tour	Show, show	Look, look	Stop by, Stop By, stop by	Come by, Come By, come, by
Email	@	Email, email	Contact, contact			
Phone	\ (?[0-9][0-9][0-9])\ (?[0-9][0-9][0-9])\ ?[-]*[0-9] [0-9] [0-9][-]*[0-9] [0-9] [0-9]					
Greeting	Hi, hi	Hello, hello	Hey, hey	Dear, dear		
Polite	Thank, thank	Thanks, thanks	Please call, please call	Sincerely, sincerely	Regards, regards	

Negative language categories

Fees	Application fee, application fee	Deposit, deposit	\$		
Employment	Employment, employment	Employer, employer	Income, income	pay stub, paystub	pay stubs, paystubs
History	Rental history, rental history	Eviction, eviction			
Background	Criminal, criminal	Background, background	Verification, verification	Verify, verify	

Table 2. Landlord response differences in e-mail content.

Keyword search	Present in neither	Present in both	White only	African-American only	Ho: $R_w - R_{AA} = 0$
<i>Total audits: 3153</i>					
Descriptive	91.53%	3.33%	3.11%	2.03%	1.08%
	[2886]	[105]	[98]	[64]	$p = 0.0076^{***}$
Other units	96.61%	0.79%	1.43%	1.17%	0.26%
	[3046]	[25]	[45]	[37]	$p = 0.4397$
View unit	73.10%	13.57%	7.42%	5.90%	1.52%
	[2305]	[428]	[234]	[186]	$p = 0.0235^{**}$
Email	89.88%	4.73%	3.43%	1.97%	1.46%
	[2834]	[149]	[108]	[62]	$p = 0.0005^{***}$
Phone number	77.77%	12.31%	5.74%	4.19%	1.55%
	[2452]	[388]	[181]	[132]	$p = 0.0066^{***}$
Greeting	59.31%	20.71%	11.20%	8.79%	2.41%
	[1870]	[653]	[353]	[277]	$p = 0.0028^{***}$
Polite	64.00%	19.00%	9.51%	7.48%	2.03%
	[2018]	[599]	[300]	[236]	$p = 0.0065^{***}$
Fees	93.63%	3.27%	1.33%	1.78%	-0.45%
	[2952]	[103]	[42]	[56]	$p = 0.1888$
Employment	96.96%	1.30%	0.79%	0.95%	-0.16%
	[3057]	[41]	[25]	[30]	$p = 0.5901$
History	98.70%	0.38%	0.44%	0.48%	-0.04%
	[3112]	[12]	[14]	[15]	$p = 1.0000$

Keyword search	Present in neither	Present in both	White only	African-American only	Ho: $R_w - R_{AA} = 0$
Background	97.24%	1.27%	0.82%	0.67%	0.15%
	[3066]	[40]	[26]	[21]	$p = 0.5601$

Notes: See Table 1 for contents of keyword searches. p -values are from McNemar paired difference in proportions tests. Number of Landlords in [].

**0.05 significance.

***0.01 significance.

Table 3. Alternative definitions of landlord response categories.

Keyword search	Present in neither	Present in both	White only	African-American only	Ho: $R_w - R_{AA} = 0$
<i>Total audits: 3153</i>					
All negative terms	90.29%	4.66%	2.70%	2.35%	0.35%
	[2847]	[147]	[85]	[74]	$p = 0.4278$
All positive terms	42.59%	35.81%	13.16%	8.44%	4.73%
	[1343]	[1129]	[415]	[266]	$p = 0.0000^{***}$
E-mail or phone	72.82%	14.62%	7.42%	5.14%	2.28%
	[2296]	[461]	[234]	[162]	$p = 0.0003^{***}$
Other or view unit	63.18%	14.11%	12.56%	10.15%	2.41%
	[1992]	[445]	[396]	[320]	$p = 0.0050^{***}$
Greeting or polite	49.03%	29.21%	12.72%	9.04%	3.68%
	[1546]	[921]	[401]	[285]	$p = 0.0000^{***}$
Extended descriptive	90.10%	4.00%	3.55%	2.35%	1.21%
	[2841]	[126]	[112]	[74]	$p = 0.0347^{**}$
Encourage action	82.02%	8.63%	4.95%	4.41%	0.54%
	[2586]	[272]	[156]	[139]	$p = 0.3516$
Discourage	95.62%	1.30%	1.68%	1.36%	0.32%
	[3015]	[41]	[53]	[43]	$p = 0.2837$

Notes: All negative terms is a search for any of the terms from all of the negative search categories. All positive terms is a search for any of the terms from all of the positive search categories. Extended descriptive includes all of the original words in the descriptive category plus the following: large, spacious, bright, neighbor, school, store, shop. The Encourage action category includes the keywords: fast, quick, hurry, fill out, drop, stop, come, visit, application, won't last, soon. The Discourage category includes the keywords: someone else, already rented, other, are you sure? afford, no longer, not available, sorry. p -values are from McNemar paired difference in proportions tests. Number of Landlords in [].

** 0.05 significance.

*** 0.01 significance.

Table 4. Discriminating landlords in multiple categories.

	Average	Number of categories						
		1	2	3	4	5	6	7
Positive language to white only (415 total)	2.14	33.98%	33.01%	21.20%	9.40%	1.20%	0.96%	0.24%
		[141]	[137]	[88]	[39]	[5]	[4]	[1]
Negative language to African American only (74 total)	1.32	71.62%	24.32%	4.05%	0.00%			
		[53]	[18]	[3]	[0]	-	-	-

Number of Landlords in [], total number of landlords in each row is the total number of landlords who replied using at least one form of positive (negative) language as measured by the categories in Table 2. Average is the average number of categories a landlord who is in at least one category appears in.

Table 5. Time elapsed between inquiry and landlord response.

Total audits: 3153	E-mail level			Landlord level			Ho: $R_w - R_{AA} = 0$
	White	African-American	(2)-(1)	Equal length	White longer	African-American longer	
Time elapsed all audits	7:09 (19:30)	7:48 (22:11)	0:39* $p = 0.065$				
Frequency of more time elapsed				36.82% [1161]	28.20% [889]	34.98% [1103]	-6.78% $p = 0.000^{***}$
Frequency of more time elapsed, reply to both races				0.74% [10]	47.02% [638]	52.25% [709]	-5.23% $p = 0.009^{***}$

Notes: Row 1 shows the average time elapsed between when an inquiry is sent and when a landlord reply is received, reported in h:mm format, these averages do not include e-mails where no reply was made, p -value for Row 1 results if from a standard difference in means z test. Row 2, reports the percentage of audits where landlords replied in the same amount of time (measured to the minute), took longer to reply to the white name, or took longer to reply to the African American name using all audits. Row 2 results include landlords that did not reply to one race, row 3 reports the same statistic as row 2 with the calculation made for only landlords who replied to both e-mails. p -values from Row 2 and 3 are from McNemar paired difference in proportions test. Standard Deviations are reported in (), number of landlords reported in [].

* 0.10 significance.

*** 0.01 significance.

Table 6. Word count of landlord response.

Total audits: 3153	E-mail level			Landlord level			
	White	African-American	(1)–(2)	Equal length	White longer	African-American longer	Ho: $R_w - R_{AA} = 0$
Words in all audits	25.86	23.7	2.16***	46.72%	29.78%	23.50%	6.28%
	(51.39)	(51.32)	$p = 0.002$	[1473]	[939]	[741]	$p = 0.000$ ***
Words, reply to both races				23.58%	40.24%	36.18%	4.06%
				[320]	[546]	[491]	$p = 0.0876$ **
Words, dropping audits with equal word length replies	31.5	28.36	3.14***				
	(50.18)	(50.40)	$p = 0.000$				

Notes: Columns (1)–(3) use word counts from all e-mails, while columns (4)–(7) use landlord level replies. Row 1 examines the word count for all audits and includes counting non-response as zero words. p -values in columns (1)–(3) results are from standard difference in means z tests. p -values in columns (4)–(6) are from McNemar paired difference in means tests. Row 2 examines only landlords that replied to both e-mails. Row 3 excludes replies that were of equal length and does not count non-responses. Standard Deviations are reported in (), number of landlords reported in []. ** 0.05 significance. *** 0.01 significance.