Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment[†]

Benjamin Edelman,[*] Michael Luca,[**] and Dan Svirsky[***]

January 6, 2016

Abstract

Online marketplaces increasingly choose to reduce the anonymity of buyers and sellers in order to facilitate trust. We demonstrate that this common market design choice results in an important unintended consequence: racial discrimination. In a field experiment on Airbnb, we find that requests from guests with distinctively African-American names are roughly 16% less likely to be accepted than identical guests with distinctively White names. The difference persists whether the host is African-American or White, male or female. The difference also persists whether the host shares the property with the guest or not, and whether the property is cheap or expensive. We validate our findings through observational data on hosts' recent experiences with African-American guests, finding host behavior consistent with some, though not all, hosts discriminating. Finally, we find that discrimination is costly for hosts who indulge in it: hosts who reject African-American guests are able to find a replacement guest only 35% of the time. On the whole, our analysis suggests a need for caution: while information can facilitate transactions, it also facilitates discrimination.

JEL Classification: D40, D47, J71

[†] We thank Ian Ayres, Larry Katz, Kevin Lang, Sendhil Mullainathan, Devah Pager, and seminar participants at eBay, Harvard Law School, Indiana University, Northwestern University, and Stanford University for valuable feedback. We thank Haruka Uchida for tireless research assistance.
[*] Harvard Business School, bedelman@hbs.edu
[**] Harvard Business School, mluca@hbs.edu
[***] Harvard Business School and Harvard University Department of Economics, dsvirsky@hbs.edu

1. Introduction

Online marketplaces such as Airbnb, Uber, and Upwork have moved an increasing share of the economy online and created a series of new markets. These platforms facilitate arms-length transactions that would previously have been infeasible both due to the difficulty of matching buyers and sellers as well as uncertain quality on both sides. To address these challenges, platform designers provide pricing and reputation mechanisms to reduce search frictions and build trust.

On first-generation online platforms, a seller agrees to do business with a given buyer before learning the name or identity of the buyer. For example, a customer booking a hotel at Expedia or Priceline need not reveal any personal information until finalizing the reservation. Moreover, a customer who makes a reservation with a valid payment card can be assured of a room being available; hotels do not reject customers based on customer name or identity. Similarly, neither Amazon nor a third-party seller in Amazon Marketplace can reject buyers on the basis of personal information. Some eBay sellers previously restricted purchases to users with certain types of verification or prior feedback, but even these restrictions have largely fallen by the wayside.

Online platforms have the potential to create fairer, more inclusive transactions. Because they facilitate commerce at a distance, online platforms can conceal information that might otherwise enable discrimination. For example, Ayres and Siegelman (1995) find that African-American car buyers pay a higher price than white car buyers at dealerships, whereas Scott Morton et al. (2003) find no such racial difference in online purchases. Similarly, platforms such as Amazon, eBay, and Expedia offer little scope for discrimination, as sellers effectively pre-commit to accept all buyers regardless of race or ethnicity.

However, these advantages are by no means guaranteed, and in fact they depend on design choices made by each online platform. Over time, platforms have moved toward systems that favor more revealing profiles that reduce anonymity for users. New platforms also often grant sellers the ability to handpick the people they transact with. If a hotel lists a room on Expedia, platform rules effectively prevent the hotel from rejecting a guest based on perceived race, ethnicity, or almost any other factor. But if the same hotel lists a room on Airbnb (which some hotels have begun to do), it could reject a guest based on these factors or others. Much of this shift has been framed in terms of the creation of a "sharing" economy, suggesting a shift towards social transactions involving interpersonal interactions and judgments.

In this paper, we investigate the existence and extent of racial discrimination on Airbnb, the canonical example of the sharing economy. Airbnb allows hosts to rent out houses, apartments, or rooms within an apartment. To facilitate these transactions, Airbnb promotes properties to prospective guests, facilitates communication, and handles payment and some aspects of customer service. Airbnb also requires hosts and guests to present their first names. After receiving an inquiry from a guest seeking to stay at a given property, the host can review the guest's profile and accept or reject the request.

To investigate discrimination, we conduct a field experiment in which we inquire about the availability of roughly 6,400 listings on Airbnb across five cities. Specifically, we create guest accounts that differ by name but are otherwise identical. Drawing on the methodology of a labor market experiment by Bertrand and Mullainathan (2004), we select two sets of names—one distinctively African-American and the other distinctively White.

Overall, we find widespread discrimination against African-American guests. Specifically, African-American guests received a positive response roughly 42% of the time, compared to roughly 50% for White guests. This 8 percentage point (roughly 16%) penalty for African-American guests is particularly noteworthy when compared to the discrimination-free setting of competing short-term accommodation platforms such as Expedia. The penalty is consistent with the racial gap found in contexts ranging from labor markets to online lending to classified ads to taxicabs.[1]

Airbnb is an appealing context to study online discrimination both because of its novelty, size, and importance, and because of the depth and breadth of data we are able to collect about which hosts discriminate and what happens before and afterwards. Combining our experimental results with observational data from Airbnb's site, we investigate whether different types of hosts discriminate more, and whether discrimination is more common at certain types of properties based on price or local demographics. In traditional audit studies, researchers have generally been unable to collect detailed data on who discriminates. For example, in mailing out resumes, a researcher does not know the identity of the specific person evaluating candidates, absent some follow-up communication (Carlsson & Rooth 2007). In contrast, Airbnb's site gives us significant information about the host including the host's name and photo, property location and details, and the mix of guests that recently stayed with a given host.

[1] Doleac & Stein (2013) find a 62% to 56% gap in offer rates for online classified postings. Bertrand and Mullainathan (2004) find a 10% to 6% gap in callback rates for jobs. Pope & Sydnor (2011) find a 9% to 6% gap in lending rates in an online lending market. Ayres et al. (2005) find a 20% to 13% gap in how often taxi drivers receive a tip.

We also expand on prior audit studies by observing what happens after discrimination occurs. In other audit studies, researchers document racial gaps in (for example) interview requests, but they cannot determine whether this gap persists when the job is ultimately filled. In contrast, our data allows us to explore these issues, and we can check whether a host ultimately finds a replacement guest. These additional analyses yield a more nuanced sense of the problem and possible policy interventions.

On the whole, we find that results are remarkably persistent. Both African-American and White hosts discriminate against African-American guests; both male and female hosts discriminate; both male and female African-American guests are discriminated against. Effects persist both for hosts that offer an entire property and for hosts who share the property with guests. Discrimination persists among experienced hosts, including those with multiple properties and those with many reviews. Discrimination persists and is of similar magnitude in high and low priced units, in diverse and homogeneous neighborhoods.

Because hosts' profile pages contain reviews (and pictures) from recent guests, we can cross-validate our experimental findings using observational data on whether the host has recently had an African-American guest. One might worry that our results are driven by some anomaly in our treatment, perhaps hosts distrusting some aspect of our guest profiles. To that end, if we continue to find a race gap even among hosts with a documented willingness to host African-Americans in the recent past, then this would cast doubt on the effectiveness of our treatment. But when we limit our analysis to such hosts, the effect disappears. This indicates that our results measure an aspect of host behavior, not anomalies from our specific treatment.

To explore the cost to a host of discriminating, we check whether each property is ultimately rented for the weekend we inquired about. Combining that information with the price of each listing, we estimate that a host incurs a cost of roughly $65-$100 in foregone revenue by rejecting an African-American guest.

Our results suggest an important tradeoff for market designers, who set the rules of online platforms, including the pricing mechanisms (Einav et al 2013) and the information that is available and actionable at the time of transaction (Luca forthcoming). Market design principles have generally focused on increasing the information flow and quality within a platform (Bolton et al 2013, Che and Horner 2014, Dai et al 2014, Fradkin et al 2014), but we highlight a situation in which platforms may be providing too much information. In particular, revealing too much information may sometimes have adverse effects, including facilitating discrimination. In the discussion section, we explore market design and policy implications of our results.

2. About Airbnb

Airbnb is a popular online marketplace for short-term rentals. Founded in 2008, the site gained traction quickly and, as of November 2015, it offers 2,000,000 listings worldwide.[2] This is more than three times as many as Marriott's 535,000 rooms worldwide. Airbnb reports serving over 40 million guests in more than 190 countries.

While the traditional hotel industry is dominated by hotels and inns that each offer many rooms, Airbnb enables anyone to post even a single room that is vacant only occasionally. Hosts provide a wealth of information about each property, including the

[2] https://www.airbnb.com/about/about-us

type of property (house, apartment, boat, or even castle, of which there are over 1400 listed), the number of bedrooms and bathrooms, the price, and location. Each host also posts information about herself. An interested guest can see a host's profile picture as well as reviews from past guests. Airbnb encourages prospective guests to confirm availability by clicking a property's "Contact" button to write to the host.[3] In our field experiments (described in the next section), we use that method to evaluate a host's receptiveness to a booking from a given guest.

3. Experimental Design

3.1 Sample and data collection

We collected data on all properties offered on Airbnb in Baltimore, Dallas, Los Angeles, St. Louis, and Washington, D.C. as of July 2015. We chose these cities in order to focus on major United States metropolitan areas with varying levels of Airbnb usage. Baltimore, Dallas, and St. Louis offer several hundred listings each, while Los Angeles and Washington, D.C. have several thousand.

Because some hosts offer multiple properties, we selected only one property per host using a random number generator. This helped to reduce the burden on any given host, and it also prevented a single host from receiving multiple identical emails. Each host was contacted for no more than one transaction in our experiment.

We also collected data from each host's profile page. This allowed us to analyze host characteristics in exceptional detail. First, we saved the host's profile image. We then employed Mechanical Turk workers to assess each host image for race (White,

---

[3] See "How do I know if a listing is available", https://www.airbnb.com/help/question/137.

African-American, Asian, Hispanic, multiracial, unknown), gender (male, female, two people of the same gender, two people of different genders, unknown), and age (young, middle-aged, old). We hired two Mechanical Turk workers to assess each image, and if the workers disagreed on race or gender, we hired a third to settle the dispute. If all three workers disagreed (as happened, for example, for a host whose profile picture was an image of a sea turtle), we manually coded the picture. (We coded race as "unknown" when the picture did not show a person.) Through this procedure, we roughly categorized hosts by race, gender, and age.

Profile pages also revealed other variables of interest. We noted the number of properties each host offers on Airbnb, anticipating that professional hosts with multiple properties might discriminate less often than others. We retrieved the number of reviews the host has received, a rough measure of whether the host is an avid Airbnb user or a casual one. We further checked the guests who had previously reviewed each host. Airbnb posts the photo of each such guest, so we used Face++, a face-detection API, to categorize past guests by race, gender, and age.[4] This allows us to examine relationships between a host's prior experience with African-American guests and the host's rejection of new African-American requests.

We also collected information about each property. We recorded the price of the property, the number of bedrooms and bathrooms, the cancelation policy, any cleaning fee, and the property's ratings from past guests. We also measured whether the property was for an entire unit or just a room in a larger unit, yielding a measure of how much the

[4] In addition to detecting race, gender, and age, Face++ estimates its confidence for each trait. When Face++ was unable to make a match or its confidence was below 95 out of 100, we used Mechanical Turk, to categorize the past guest via the method described above.

host interacts with the guest. Though renting out a room in one's house does not guarantee that the host and guest will share the property, that is the typical arrangement in such properties.

Each property listing included a longitude and latitude, which allowed us to link to census demographic data to assess the relationship between neighborhood demographics and racial discrimination. After linking the latitude and longitude to a census tract, we used census data on the number of African-American, Hispanic, Asian, and White individuals. Table 1 presents summary statistics about the hosts and listings.

We later checked each listing to see whether hosts were ultimately able to fill openings. Our guests inquired about reservations eight weeks in advance. Thus, if a guest sent a message on August 1 about the weekend of September 25, we checked on Friday, September 24 to see whether the specified property was still listed as available.

3.2 Treatment groups

Our analysis used four main treatment groups based on the perceived race and gender of the test guest accounts. Hosts were contacted by guests with names that signaled African-American males, African-American females, White males, and White females.

To experimentally vary the perceived race of the guest, we began with a list of names that are distinctively White and names that are distinctively African-American, drawn from Bertrand and Mullainathan (2004). The list was based on the frequency of names from birth certificates of babies born between 1974 and 1979 in Massachusetts. Distinctively White names are those that are most likely to be White, conditional on the name, and similarly for distinctively African-American names. To validate the list, we conducted a survey in which we asked participants to quickly categorize each name as White or African-American. With just three seconds permitted for a response, survey takers had little time to think beyond a gut response. The survey results, presented in the Appendix, confirm that the names continue to signal race.[5]

We then created twenty Airbnb accounts, identical in all respects except for guest names. Our names included ten that are distinctively African-American names and ten distinctively White names, divided into five male and five female names within each group. To avoid the confounds that would result from pictures, we use only names; our Airbnb profiles include no picture of the putative guest.[6] From these twenty guest accounts, we sent messages to prospective hosts. Figure 1 presents a representative email from one of our guests to an Airbnb host. The name and dates changed depending on the message sender and when the message was sent. In choosing the dates, we asked hosts about a weekend that was approximately eight weeks distant from when the message was sent. We limited our search to those properties that were listed as available during the weekend in question.

3.3 Experimental procedure

We sent roughly 6,400 messages to hosts between July 7, 2015 and July 30, 2015. Each message inquired about availability during a specific weekend in September. When a host replied to a guest, we replied to the host with a personal message clarifying that we

[5] On a scale of 0 to 1, where 0 is African-American, the White female names each had an average survey response of 0.90 or above, and the African-American female names all had an average score of 0.10 or below. The male names showed slightly more variation but tell the same story: all the White male names scored 0.88 or above, and all the African-American male names except for Jermaine Jones scored 0.10 or below. The Appendix presents the full results of the survey.

[6] Accounts with no profile pictures are common among casual Airbnb users. We collected a random set of 449 profile pictures; 44% of these profiles lacked a picture.

(as the guest) were still not sure if we would visit the city or if we would need a place to stay. We sent this reply in order to reduce the likelihood of a host holding inventory for one of our hypothetical guests.

We tracked host responses over the 30 days that followed each request. A research assistant then coded each response into categories. The majority of responses were in one of six groups: “No response” (if the host did not respond within 30 days); “No or listing is unavailable”; “Yes”; “Request for more information” (if the host responded with questions for the guest); “Yes, with questions” (if the host approved the stay but also asked questions); “Check back later for definitive answer”; and “I will get back to you.”

As these categories show, our initial categorizations used subtle distinctions between possible responses. In our analyses below, however, we restrict our attention to the simplest response: “Yes.” All of the main results are robust to using “No” instead, as well as a variety of interpretations of the intermediate responses. The results are also robust to using the time until a host responds as a dependent variable.

We collected all data using scrapers we built for this purpose. We sent inquiries to Airbnb hosts using web browser automation tools we built for this purpose. Our Institutional Review Board approved our methods before we began collecting data.

4. Results

Table 2 presents the main effect. We find that inquiries from guests with White-sounding names are accepted roughly 50% of the time. In contrast, guests with African-American-sounding names are accepted roughly 42% of the time. Relative to the 50% base response rate, the eight percentage point difference represents a 16% reduction in the acceptance rate for African-American guests. Columns 2 and 3 introduce additional control variables related to the host or the property. The race effect stays constant at a roughly eight percentage point gap across these specifications, controlling for the host's gender, race, an indicator for whether the host has multiple listings, an indicator for whether the property is shared, host experience (whether the host has more than ten reviews), and the log of the property price.

To put this effect into perspective, suppose we were to run the same experiment on a standard hotel-booking platform such as Expedia or Priceline. Due to their design, these platforms would necessarily have zero race gap under our experiment because they do not allow hotels to decide whether to accept a prospective guest based on the guest's name. Similarly, Scott Morton et al. (2003) find no difference by race in price paid for cars in online purchases—a sharp contrast to traditional channels. They conclude that the Internet is reducing discrimination. However, our result contributes to a small but growing body of literature suggesting that discrimination persists—and we argue may even be exacerbated—in online platforms. Edelman and Luca (2014) show that African-American hosts on Airbnb seek and receive lower prices than White hosts, controlling for the observable attributes of each listing. Pope and Sydnor (2011) find that loan listings with pictures of African-Americans on Prosper.com are less likely to be funded than similar listings with pictures of White borrowers. Doleac and Stein (2013) show that buyers are less likely to respond to Craigslist listings showing an iPod held by a Black hand compared to an identical ad with a White hand.

The additional information provided in the newest online platforms highlights the crucial role of market designers in choosing the scope for discrimination. While it is unlikely that Airbnb will achieve this discrimination-free benchmark, design changes might lead to lower levels of discrimination.

4.1. Effects by host characteristics

In this section, we explore the relationship between the extent of discrimination and the characteristics of hosts. Overall, our results suggest that discrimination is not limited to one type of host or one particular situation. The effect is stable, and it persists across hosts of diverse races, ages, and levels of experience. It exists for hosts with a single property, as well as for those with multiple listings.

One possible explanation of discrimination is homophily (in-group bias). According to this theory, hosts might simply prefer guests of the same race. If homophily was the primary factor driving differential guest acceptance rates, then African-American guests would face higher acceptance rates from African-American hosts. Table 3 presents regressions that include guest race, host race, and an interaction term. Across the entire sample of hosts, the interaction between the race and guest of the host is not significantly different from zero, but the point estimate is noisy. This result masks heterogeneity across genders. Columns 2 and 3 of Table 3 report the same regression limited to male hosts and female hosts, respectively. Among male hosts, the interaction between the host's race and guest's race shows a widening of the race gap by 11 percentage points, whereas among females, the race gap narrows by 11 percentage points. Both estimates are noisy; we cannot reject coefficients of zero.[7]

[7] Table 4 explores the effect of the host's race with more nuance. It shows the proportion of Yes responses from each gender/race cell among hosts in response to each gender/race cell among guests. African-American male hosts discriminate against African-American male and female guests. White hosts of both genders are more likely to accept white guests of either gender. African-American female hosts are the only

Discrimination may also be influenced by a host's proximity to the guest. For example, Becker (1957) formalizes racial discrimination as distaste for interactions with individuals of a certain race. On Airbnb, a host must classify each property as offering an entire unit, a room within a unit, or a shared room. We classify anything other than an entire unit as a "shared property." Column 1 shows that the race gap is roughly the same whether or not a property is shared.

One might expect a distinction between casual Airbnb hosts who occasionally rent out their homes, versus professional hosts who offer multiple properties. Roughly a sixth of Airbnb hosts manage multiple properties, and roughly 40% of hosts have at least 10 reviews from past guests. Columns 2 and 3 explore the extent of discrimination among hosts with multiple locations, and those with more than 10 reviews. Across these specifications, the race gap persists with roughly the same magnitude.

To the extent that discrimination rates are changing over time, one might expect racial discrimination to be less common among younger hosts. To assess this possibility, we employed Mechanical Turk workers to categorize hosts as young, middle-aged, or old. Column 4 shows that discrimination also persists across the age categories with roughly the same magnitude.

4.2. Effects by location characteristics

Just as racial discrimination was robust across host characteristics, we find that discrimination does not vary based on the cost or location of the property.

---

exception: they accept African-American female guests more than any other group. Thus, with the exception of African-American females, the data is inconsistent with homophily.

Column 1 of Table 6 shows that, overall, listings above the median price are more likely to reject inquiries. However, the extent of discrimination remains the same: hosts with expensive listings are just as likely to discriminate as those with less expensive listings.

We also hypothesized that the extent of discrimination might vary with the diversity of a neighborhood. More generally, one might expect that geography matters and that discrimination is worse in some areas than others, due to market structure or underlying rates of discrimination among a population. Merging data on neighborhoods by census tract, Column 2 shows that the extent of discrimination does not vary with the proportion of nearby residents who are African-American. Column 3 shows that discrimination is ubiquitous: it does not vary with the number of Airbnb listings within the census tract. The point estimates are also similar across cities in our sample population.

4.4 Robustness – effects by name

Table 7 shows the proportion of positive responses broken down by name. The effect is robust across choice of names. For example, the African-American female name with the most positive responses (Tamika) received fewer positive responses than the White female name with the fewest positive responses (Kristen). Similarly, the African-American males with the most positive responses (Darnell and Rasheed) received fewer acceptances than the White male with the fewest positive responses (Brad).

4.5 Comparing experimental results with observational patterns

Audit studies have yielded important insights in labor economics, and promise further benefits in studying online platforms. That said, audit studies suffer an important potential limitation: the experimenter designs the profiles, chooses the names, and chooses the sample population, so even a carefully designed experiment might not fully reflect broader conditions. In this section, we exploit the richness of our data to assess the external validity of our results.

This further analysis is grounded in the detailed data Airbnb posts about each host, notably including prior reviews of the host by past guests. These reviews are in turn accompanied by guests' photos, which allow us to see which hosts previously accepted African-American guests.

For this analysis, we collected profile pictures from the ten most recent reviews evaluating each of the hosts we had contacted. We then categorized these past guests by race and gender, and we regressed the likelihood of a host responding positively to our inquiry on the race of the guest, whether the host has at least one recent review from an African-American guest, and an interaction between these variables. Column 5 of Table 5 reports the results. We find that the race gap drops sharply among hosts with at least one recent review from an African-American guest, and we cannot reject zero difference for requests from our African-American test accounts versus requests from our White test accounts.[8]

This finding reinforces our interpretation of our main effects, including the role of race and the interpretation that observed differences reflect racial discrimination by Airbnb hosts. Put another way, if our findings are driven by a quirk of our experimental design, rather than race, then it is difficult to explain why the race gap disappears

[8] These findings are robust to alternative specifications of a host's past guests. The same substantive results hold if we instead look at the raw number of reviews from African-Americans, rather than whether there is at least one such review. The same is true if we use the proportion of reviews from African-American guests.

precisely among hosts with a history of accepting African-American guests. While we hesitate to generalize to other audit studies in other contexts and with other methodologies, the apparent validity of our approach confirms the general attractiveness of the audit study approach.

4.6 How much does discrimination cost hosts?

A host incurs a cost for discriminating when rejecting a guest causes a unit to remain empty. The expected cost depends on the likelihood of the property remaining vacant, which in turn depends on the thickness of the market. If a host can easily find a replacement guest, then discrimination is nearly costless for the host. But if a property remains vacant after the host rejects a guest, then discrimination imposes a more significant cost. In other words, the impact on net revenue from discriminating depends on the likelihood of filling a unit with someone of the host's preferred race after rejecting a guest of a disfavored race.

Because we collect data about each property's availability after a host declines a guest, we can estimate the cost in net revenue from discrimination. Suppose a host charges price p for a listing and pays listing fees f to Airbnb. Let $\pi_{replace}$ be the probability of filling the property after rejecting a guest in our study. Then the cost in net revenue of discrimination is as follows:

$$\Delta Net\ Revenue = (p - f) - \pi_{replace} \cdot (p - f)$$

$$= (1 - \pi_{replace}) \cdot (p - f)$$

That is, the cost of discrimination, in terms of net revenue, is the revenue that the host forgoes if the listing remains empty multiplied by the probability that the listing remains empty.

In our data, hosts who rejected or never responded to our inquiries had properties with a median price of \$163 and a mean price of \$295. The numbers are similar and slightly higher if we restrict the sample further to those hosts who rejected African-American guests.[9] Airbnb charges each host a fee equal to 3% of the listing price.

After our inquiries, roughly 25.9% of the listings in our study remained vacant on the dates we requested after rejecting or not responding to one of our guests. Another 37.9% remained listed but were no longer available on those dates, suggesting that the host either found another guest or decided to no longer make the property available on the specified dates. The remaining 36.1% of properties were no longer listed on Airbnb. Because it is unclear whether the hosts who exit should be excluded from the sample or treated as not having found a replacement, we develop two estimates.

If we exclude these disappearing hosts from our calculation, 59.4% of hosts found a replacement guest. Setting p equal to the median price (\$163) and fees at 3% of the median price:

$$\Delta Net\ Revenue = (1 - .594) \cdot (\$163 - .03 \cdot \$163) \approx \$64.19$$

If we treat disappearing listings as vacancies, in effect assuming that the host of a dropped listing was not able to find a replacement guest, then only 37.9% of hosts found a replacement guest. The cost of discrimination rises as a result:

$$\Delta Net\ Revenue = (1 - .379) \cdot (\$163 - .03 \cdot \$163) \approx \$98.19$$

In this analysis, we focus on the net revenue, which does not incorporate the marginal cost of each night the listing is rented, since we do not directly observe costs. The cost of hosting includes various types of host effort or wear-and-tear to the property.

[9] In calculating price, we sum the listing price and any cleaning fee.

In principle, hosting also entails a risk of damage by a guest, though throughout the relevant period Airbnb automatically provided all hosts with property insurance, which reduces the risk. Our calculation also excludes unobserved benefits of hosting, such as the possibility that a positive review draws more guests in the future and improves the listing position on Airbnb. A full estimate of profit would also need to consider the time cost of looking for new guests after rejecting someone on the basis of race.

While these estimates are clearly noisy, they suggest that hosts incur a real cost by discriminating. The median host who rejects a guest because of race is turning down between \$65 and \$100 of revenue.

5. Discussion

Online platforms such as Airbnb create new markets by eliminating search frictions, building trust, and facilitating transactions (Lewis 2011, Luca forthcoming). With the rise of the sharing economy, however, comes a level of racial discrimination that is impossible in the hotel reservations process. Clearly, the manager of a Holiday Inn cannot examine names of potential guests and reject them based on race. Yet, this is commonplace on Airbnb, which now accounts for a growing share of the hotel market. In this section, we discuss implications for market designers and policy-makers.

5.1 Statistical versus taste-based discrimination

Economic models of discrimination often distinguish between statistical and taste-based discrimination. Our findings suggest a more nuanced story than either of the classic models.

Our findings are inconsistent with the simplest versions of taste-based discrimination. Under this view, hosts may dislike interactions with people of a certain race. If this is true, then one would expect increased proximity to the guest to worsen the race gap. But we find no such evidence. Similarly, our data is inconsistent with a simple theory of homophily, in which people have a distaste for people of other racial or gender groups. We find homophily among African-American females, but not among other race/gender combinations.

We also find some evidence against pure statistical discrimination. Under statistical discrimination, hosts use race as a proxy for some undesirable trait, like the likelihood that a guest does damage to the property. As noted above, we find that hosts who have had an African-American guest in the past exhibit less discrimination than other hosts. This suggests that, at the very least, hosts are using a variety of statistical models as they evaluate potential guests. Similarly, there is some evidence that hosts are using incorrect statistical models. Across all hosts, African-American guests fared worse than White guests, and male guests fared worse than female guests. Both findings are consistent with demographics on crime rates by race and gender, which suggests that statistical discrimination is occurring. But within race/gender units, African-American females fared worse than White males, which is inconsistent with the same statistics.[10] Indeed, in survey data linking the names in our study to the likelihood of property damage, respondents offered no significant difference in their assessment of African-American male and female names, and both were rated as more likely to cause damage

[10] This mirrors findings by Pager (2003), which uses an audit study to measure the effects of race and a criminal record on job callbacks. Pager (2003) finds that African-American applicants *without* a criminal record fared worse than white job applicants *with* one. Taken together, both Pager (2003) and our results suggest that hosts and employers find race to be a more salient proxy for risk than other traits, like gender or a criminal history.

than White names of either gender.[11]

In sum, the richness of our data paints a more nuanced picture than any single economic model of discrimination. If hosts exhibit homophily, the effect is heterogeneous across race and gender groups. If hosts are using statistical discrimination, their statistical models are inconsistent and by all indications inaccurate.

5.2 Can discrimination persist in a market?

As discussed above, when firms discriminate, they forego revenue that they would otherwise receive, making them less competitive. Becker (1957) suggests that this loss will help drive discriminating firms out of competitive markets. But in the Airbnb context, discrimination can persist in the long run because competition is not perfect: listings are differentiated, and some hosts may be willing to trade off revenue with their preference for White guests. Consistent with this assessment, columns 2 and 3 of Table 5 reports that experienced hosts (hosts with a long history of reviews as well as hosts with multiple properties) were just as likely to discriminate as others. These findings dim hopes of competition preventing discrimination. In that light, the next two sections discuss actions that Airbnb might take to reduce discrimination, as well as policy implications.

5.3 Designing a discrimination-free marketplace

Because online platforms choose which information is available to parties during a transaction, they can prevent the transmission of information that is irrelevant or potentially pernicious. Our results highlight a platform's role in preventing discrimination

[11] These findings come from a survey on Mechanical Turk. We presented each participant with one name from our study and asked whether an Airbnb guest with this name is more or less likely to commit property damage than the population of Airbnb guests as a whole.

or facilitating discrimination, as the case may be. If a platform aspires to provide a discrimination-free environment, its rules must be designed accordingly.

Airbnb has several options to reduce discrimination. For example, it could conceal guest names, just as it already prevents transmission of email addresses and phone numbers so that guests and hosts cannot circumvent Airbnb's platform and its fees. Communications on eBay's platform have long used pseudonyms and automatic salutations, so Airbnb could easily implement that approach.

Alternatively, Airbnb might further expand its "Instant Book" option, in which hosts accept guests without screening them first. Closer to traditional hotels and bed and breakfasts, this system would eliminate the opportunity for discrimination. This change also offers convenience benefits for guests, who can count on their booking being confirmed more quickly and with fewer steps. However, in our sample, only a small fraction of hosts currently allow instant booking. Airbnb could push to expand this fraction, which would also serve the company's broader goal of reducing search frictions.

5.4 Policy Implications

Because the legal system grants considerable protection to online marketplaces, Airbnb is unlikely to be held liable for allowing discrimination on its platform. As discussed in Edelman and Luca (2014), any changes by Airbnb would likely be driven by ethical considerations or public pressure rather than law.

In contrast, some hosts on Airbnb may be liable for discriminating against guests. Within the United States, the Civil Rights Act of 1964 prohibits discrimination in hotels (and other public accommodations) based on race, color, religion, or national origin. While some hosts may be exempt from these requirements due to their small size, hosts and managers with sufficiently many properties (or rooms) might be held liable (Todisco 2015). One clear policy implication is that regulators may want to audit Airbnb hosts using an approach based on our paper—much like longstanding efforts to reduce discrimination in offline rental markets. One might have hoped that online markets would cure discrimination, and it seems a different design might indeed do so. Regrettably, our analysis indicates that at Airbnb, this is not yet the case.

References

Ayres, I., & Siegelman, P. (1995). Race and Gender Discrimination in Bargaining for a New Car. *American Economic Review*, *85*(3), 304–321.

Ayres, I., F. Vars, & N. Zakariya. (2005). To Insure Prejudice: Racial Disparities in Taxicab Tipping. *Yale Law Journal*, *114*(7), 1613-1674.

Becker, G. (1957). The Economics of Discrimination. *The University of Chicago Press.*

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A field experiment on albor market discrimination. *American Economic Review*, *94*(4), 991–1013.

Bolton, G., B. Greiner, and A. Ockenfels. (2013). Engineering Trust: Reciprocity in the Production of Reputation Information. *Management Science*, *59*(2), 265-285.

Carlsson, M., & Rooth, D. (2007). Evidence of Ethnic Discrimination in the Swedish Labor Market Using Experimental Data. *Labour Economics*, *14*(4), 716-729.

Che, Y. and J. Hörner. (2014). Optimal Design for Social Learning. *Working Paper*.

Dai, W., G. Jin, J. Lee, & M. Luca. (2014). Optimal Aggregation of Consumer Ratings: an Application to Yelp.com. *NBER Working Paper*.

Doleac, J., & L. Stein. (2013). The Visible Hand: Race and Online Market Outcomes. *Economic Journal*, *123*(572), 469-492.

Edelman, B, and M. Luca. (2014). Digital Discrimination: The Case of Airbnb.com. *Harvard Business School Working Paper*.

Einav, L., C. Farronato, J. Levin, & N. Sundaresan. (2013). Sales Mechanisms in Online Markets: What Happened to Internet Auctions? *Working paper.*

Fradkin, A., E. Grewal, D. Holtz, and M. Pearson, (2014). Bias and Reciprocity in Online Reviews: Evidence from Field Experiments on Airbnb. *Working Paper*

Lewis, G. (2011). Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors. *American Economic Review*, 101(4): 1535-46.

Luca, M. forthcoming. User-generated Content and Social Media. *Handbook of Media Economics.*

Pager, D. (2003). The Mark of a Criminal Record. *American Journal of Sociology, 108*(5), 937–975.

Pope, D. G., & Sydnor, J. R. (2011). What's in a Picture?: Evidence of Discrimination from Prosper.com. *Journal of Human Resources*, *46*(1), 53–92.

Scott Morton, F., F. Zettelmeyer, and J. Silva-Risso, 2003. Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities? *Quantitative Marketing and Economics*, *1*(1), 65-92.

Todisco, M. (2015). Share and Share Alike? Considering Racial Discrimination in the Nascent Room-Sharing Economy. *Stanford Law Review Online*, *67*, 121–129.

Figures

Figure 1: Sample Treatment

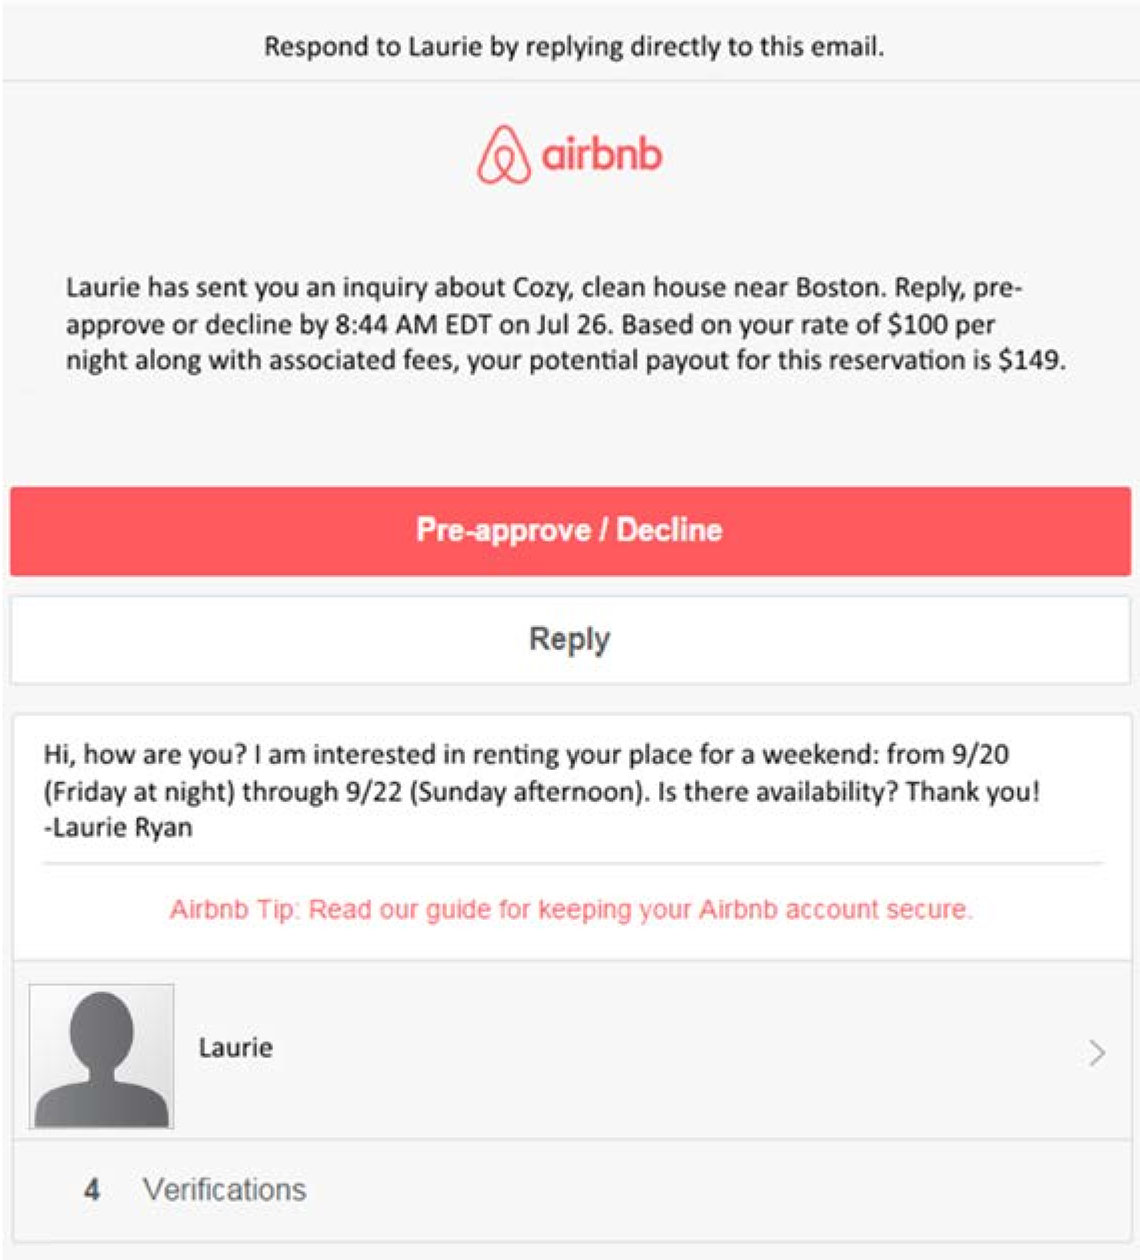

Tables

Table 1. Summary Statistics

Variables	Mean	Std. Dev.	25th %tile	75th %tile	Obs
Host is White	0.63	0.48	0	1	6,392
Host is African-American	0.08	0.27	0	0	6,392
Host is female	0.38	0.48	0	1	6,392
Host is male	0.30	0.46	0	1	6,392
Price ($)	181.11	1280.23	75	175	6,302
Number of bedrooms	3.18	2.26	2	4	6,242
Number of bathrooms	3.17	2.26	2	4	6,285
Number of reviews	30.87	72.51	2	29	6,390
Host has multiple listings	0.16	0.36	0	0	6,392
Airbnb listings per Census tract	0.14	0.2	0.03	0.14	6,378
% population black (Census tract)	9.51	9.28	2	14	6,392

Table 2. The Impact of Race on Likelihood of Acceptance

	Dependent Variable: 1(Host Accepts)		
Guest is African-American	-0.08*** (0.02)	-0.08*** (0.02)	-0.09*** (0.02)
Host is African-American		0.07** (0.02)	0.09*** (0.02)
Host is Male		-0.05** (0.02)	-0.05** (0.02)
Host has Multiple Listings			0.09*** (0.01)
Shared Property			-0.07*** (0.02)
Host has 10+ Reviews			0.12*** (0.01)
ln(Price)			-0.06*** (0.01)
Constant	0.49*** (0.01)	0.50*** (0.01)	0.76*** (0.07)
Observations	6,235	6,235	6,168
Adjusted R^2	0.006	0.009	0.040

Notes: A host's response is coded as a "Yes" if, in her reply to the guest, she invites the guest to stay at the property, if she offers a special deal ("book within 24 hours and get a discount"), or approves the guest while also asking some clarifying question ("You can stay, but how many people will you have with you?"). Standard errors are clustered by guest name and are reported in parentheses.

* p < .10. ** p < .05. *** p < .01.

Table 3: Race gap by race of the host

	Dependent Variable: 1(Host Accepts)		
	All Hosts	Male Hosts	Female Hosts
Guest is African-American	-0.08***	-0.09**	-0.09***
	(0.02)	(0.02)	(0.02)
Host is African-American	0.06	0.19**	-0.00
	(0.04)	(0.06)	(0.04)
Host is African-American * Guest is African-American	0.01	-0.11	0.11
	(0.05)	(0.07)	(0.06)
Constant	0.48***	0.44***	0.50***
	(0.01)	(0.02)	(0.02)
Observations	6235	1854	2336
Adjusted R^2	0.007	0.015	0.007

Notes: Standard errors are clustered by guest name and are reported in parentheses.
* p < .10. ** p < .05. *** p < .01.

Table 4. Proportion of Positive Responses by Race and Gender

		Guest Race / Gender			
		White Male	Black Male	White Female	Black Female
Host Race / Gender	White Male	0.42	0.35	0.49	0.32***
	Black Male	0.64**	0.40	0.59	0.43
	White Female	0.46	0.35	0.49	0.44
	Black Female	0.43	0.38	0.53	0.59***

Notes: This table shows the proportion of Yes responses by hosts of a certain race/gender to guests of a certain race/gender.

* p < .10. ** p <.05. *** p<.01. P-values from testing that proportion of Yes responses in a specific cell is equal to the proportion of Yes responses from the other cells in that colum

Table 5. **Are Effects Driven by Host Characteristics?**

	Dependent Variable: 1(Host Accepts)				
Guest is African-American	-0.07*** (0.02)	-0.08*** (0.02)	-0.09*** (0.02)	-0.11*** (0.02)	-0.09*** (0.02)
Shared Property	0.00 (0.02)				
Shared Property * Guest is African-American	-0.02 (0.03)				
Host has Multiple Listings		0.14*** (0.02)			
Host has Multiple Listings * Guest is African-American		-0.01 (0.03)			
Host has 10+ Reviews			0.14*** (0.02)		
Host has Ten+ Reviews * Guest is African-American			0.01 (0.03)		
Host Looks Young				-0.03 (0.02)	
Host Looks Young * Guest is African-American				-0.01 (0.03)	
Host has 1+ reviews from an African-American guest					0.10*** (0.02)
Host has 1+ reviews from an African-American guest * Guest is African-American					0.06* (0.03)
Constant	0.49*** (0.01)	0.46*** (0.01)	0.42*** (0.01)	0.50*** (0.01)	0.46*** (0.01)
Observations	6,235	6,235	6,235	6,235	6,235
Adjusted R^2	0.006	0.014	0.027	0.011	0.019

Notes: Standard errors are clustered by guest name and are reported in parentheses.
* p < .10. ** p < .05. *** p < .01.

Table 6. Are Effects Driven by Location Characteristics?

	Dependent Variable=1(Host Accepts)		
Guest is African-American	-0.08**	-0.08***	-0.09**
	(0.02)	(0.02)	(0.02)
Price > Median	-0.06***		
	(0.02)		
Guest is African-American * (Price > Median)	-0.01		
	(0.03)		
Share of Black Population in Census Tract		0.05	
		(0.05)	
Guest is African-American * (Share of Black Population in Census Tract)		0.02	
		(0.08)	
Airbnb Listings per Census Tract			0.00
			(0.00)
Guest is African-American * (Airbnb Listings per Census Tract)			0.00
			(0.00)
Constant	0.52***	0.48***	0.49***
	(-0.02)	(-0.01)	(-0.02)
Observations	6235	6223	6235
Adjusted R^2	0.01	0.006	0.006

Notes: Standard errors are clustered by guest name and are reported in parentheses.
* p < .10. ** p < .05. *** p < .01.

Table 7. Proportion of Positive Responses, by Name

Entire Sample		0.43 (6,390)	
White Female		Black Female	
Allison Sullivan	0.49 (306)	Lakisha Jones	0.42 (324)
Anne Murphy	0.56** (344)	Latonya Robinson	0.35** (331)
Kristen Sullivan	0.48 (325)	Latoya Williams	0.43 (327)
Laurie Ryan	0.50 (327)	Tamika Williams	0.47** (339)
Meredith O'Brien	0.49 (303)	Tanisha Jackson	0.40 (309)
White Male		Black Male	
Brad Walsh	0.41* (317)	Darnell Jackson	0.38 (285)
Brent Baker	0.48 (332)	Jamal Jones	0.33 (328)
Brett Walsh	0.44 (279)	Jermaine Jones	0.36 (300)
Greg O'Brien	0.45 (312)	Rasheed Jackson	0.38 (313)
Todd McCarthy	0.43 (314)	Tyrone Robinson	0.36 (254)

Notes: The table reports the proportion of Yes responses by name. The number of messages sent by each guest name is shown in parentheses.

* $p < .10$. ** $p < .05$. *** $p < .01$. P-values from test of proportion. Null hypothesis is that the proportion of Yes responses for a specific name are equal to the proportion of Yes responses for all other names of the same race*gender cell.

Appendix

Results of survey testing races associated with names

White Female		White Male	
Meredith O'Brien	0.93	Greg O'Brien	0.88
Anne Murphy	0.95	Brent Baker	0.90
Laurie Ryan	0.97	Brad Walsh	0.91
Allison Sullivan	0.98	Brett Walsh	0.93
Kristen Sullivan	1.00	Todd McCarthy	0.98
Black Female		Black Male	
Tanisha Jackson	0.03	Tyrone Robinson	0.00
Lakisha Jones	0.05	Rasheed Jackson	0.06
Latoya Williams	0.05	Jamal Jones	0.07
Latonya Robinson	0.07	Darnell Jackson	0.10
Tamika Williams	0.07	Jermaine Jones	0.26

Notes: "White" is coded as 1. "African-American" is coded as 0. Sample size = 62.