Reducing Discrimination with Reviews in the Sharing Economy: Evidence from Field Experiments on Airbnb

Ruomeng Cui*
Kelley School of Business, Indiana University, cuir@indiana.edu

Jun Li*
University of Michigan, Stephen M. Ross School of Business, junwli@umich.edu

Dennis J. Zhang*
Olin Business School, Washington University in St. Louis, denniszhang@wustl.edu

Recent research has found widespread discrimination by hosts against guests of certain races in online marketplaces. In this paper, we explore ways to reduce such discrimination using online reputation systems. We conduct four randomized field experiments among 1,801 hosts on Airbnb by creating fictitious guest accounts and sending accommodation requests to them. We find that requests from guests with African American–sounding names are 19.2 percentage points less likely to be accepted than those with white–sounding names. However, a positive review posted on a guest's page significantly reduces discrimination: When guest accounts receive a positive review, the acceptance rates of guest accounts with white–sounding and African American–sounding names are statistically indistinguishable. We further show that a non-positive review and a blank review without any content can also help attenuate discrimination, but self-claimed information on tidiness and friendliness cannot reduce discrimination, which indicates the importance of encouraging credible peer-generated reviews. Our results offer direct and clear guidance for sharing-economy platforms to reduce discrimination.

Key words: Discrimination, Field Experiment, Information Sharing, Service Operations, Sharing Economy.

1. Introduction

Discrimination has become an important issue in the recent development of sharing-economy marketplaces. Edelman et al. (2017) raises a serious concern over racial discrimination on Airbnb. Using a field experiment, they show that guests with African American–sounding names are 16 percent less likely to be accommodated relative to identical guests with white–sounding names.

* The authors contribute equally to the paper and the authorship is determined alphabetically. The authors thank Ryan Buell, Benjamin G. Edelman, Chiara Farronato, Christopher Tang, Michael W. Toffel, Joseph Xu, Gerard Cachon, Michael Luca and Jan A. Van Mieghem for their helpful comments on the paper.

Many African American users have expressed frustration on social media about how they were allegedly denied a booking request by Airbnb hosts because of their race.[1]

As discrimination in the sharing economy has become a well-known issue to the public (also see Ge et al. (2016) for evidence of discrimination among Uber and Lyft drivers), reducing discrimination has become an important issue in marketplace design and operations. For example, Airbnb states that "we welcome the opportunity to work with anyone that can help us reduce potential discrimination in the Airbnb community."[2] While there is a burgeoning literature in operations management that focuses on the design of marketplaces to improve market efficiency (e.g., Cohen and Harsha 2013, Cachon et al. 2015, Zhang et al. 2017, Feldman et al. 2018) and social welfare (e.g., Ashlagi and Shi 2016, Feldman et al. 2017), discriminatory behavior is often an overlooked factor that hinders effective market mechanisms.

Recognizing this opportunity, we investigate whether peer-generated reviews can reduce discrimination. In particular, we want to understand whether hosts' reviews of guests affect other hosts' discriminatory behavior against these guests based on their race. Reviews may help attenuate discrimination for two main reasons. First, online reviews have been shown to be a credible source of information to signal product or user quality (Bolton et al. 2004, Chevalier and Mayzlin 2006). Reviews on Airbnb provide valuable information, such as safety, tidiness, and friendliness, about a prospective guest. Therefore, additional information from reviews would help hosts make more informed decisions, rather than basing their decisions on guests' race. Second, reviews could also help establish an inclusive normative behavior among community members. The fact that other hosts have accepted a guest encourages a host to accept the guest, regardless of her race.

Moreover, we are also interested in understanding what characteristics of review are most critical in reducing discrimination, including sentiment (i.e., positive or non-positive descriptions of a prior experience with a guest), credibility (i.e., peer-generated or self-claimed quality information),[3] and the existence vs. content of a review (i.e., the fact that a review exists on a guest profile vs. what the host said in the review).

We conduct four randomized field experiments on Airbnb to address these questions. In each experiment, we manipulate both guests' race and the review (or self-claimed) information by employing a 2×2 design. In the first field experiment, we create eight fictitious guest accounts; four accounts do not have reviews and the other four accounts have one positive review written by the same host at the same time. Within each treatment arm, four guest accounts differ only by

[1] Twitter hashtag #Airbnbwhileblack. https://twitter.com/hashtag/airbnbwhileblack

[2] "Discrimination by Airbnb Hosts Is Widespread, Report Says," *New York Times*. December 11, 2015.

[3] On Airbnb, the content of reviews is mostly positive. Some reviews are less positive than others but rarely very negative. Therefore, we test the impact of such "non-positive" reviews rather than negative ones.

name but otherwise are identical. Two guests have white–sounding names and two have African American–sounding names.[4] We then randomly assign these guest accounts to Airbnb hosts in three major U.S. cities—Boston, Chicago, and Seattle—and send out accommodation requests from our guest accounts to these hosts. We record hosts' reply messages and compare acceptance rates across guest accounts. Because we randomly assign Airbnb hosts to guest accounts that differ only by names and reviews, the observed difference in acceptance rate is causally driven by race and review information. We refer to this experiment as the *positive review experiment*.

In the second field experiment, we create another eight fictitious guest accounts and repeat the previous experimental design with one change: the latter four accounts in the review condition receive a non-positive review instead of a positive review. We designate this experiment the *non-positive review experiment*. In the third field experiment, we create another 16 guest accounts and repeat the experimental design with another change: all guest accounts lack reviews, and the latter eight guests self-claim to be neat and friendly in their accommodation request messages. This enables us to test whether self-claimed and unverified information can reduce discrimination; we refer to this experiment as the *self-claimed information experiment*. In the last field experiment, we again create 16 new accounts and repeat the experimental design with one modification: each of the latter 8 guest accounts in the review condition has one blank review without content. This setup allows us to separate the effect of a review's existence from its content. We denote this experiment as the *blank review experiment*. We conduct the four experiments in September 2016, October and November 2016, July and August 2017, and March and April 2018 respectively.

Our positive review experiment suggests that discrimination exists when guest accounts have no review: the average acceptance rates of guests with white–sounding names and guests with African American–sounding names are 47.9 percent and 28.8 percent, respectively, and the difference is 19.2 percentage points (p-value = 0.0002). This result is consistent with Edelman et al. (2017), which finds that guests with white–sounding names are more likely to be accepted compared to guests with African American–sounding names.[5] However, when there is one positive review, discrimination is significantly reduced: the acceptance rate is 56.2 percent for guests with white–sounding names and 58.1 percent for guests with African American–sounding names, and the difference between acceptance rates for white and African American guests is statistically indistinguishable (p-value = 0.8774). Moreover, irrespective of a guest's race, the acceptance rate is higher when the guest has a positive review.

[4] Following the past literature (Bertrand and Mullainathan 2004), the white–sounding names and African American–sounding names are chosen based on name frequency data published by the U.S. Census Bureau's Population division. Please refer to Demographic Aspects of Surnames from Census 2000.

[5] Our acceptance rate of white guests is similar to that of Edelman et al. (2017), while our acceptance rate of African guests is lower.

The remaining three experiments demonstrate whether different types of information can help attenuate discrimination. In particular, our non-positive review experiment suggests that a non-positive review can significantly reduce discrimination: in the absence of reviews, guests with white–sounding names are 21.4 percentage points more likely to be accepted than guests with African American–sounding names (p-value $= 0.0287$); when there is a non-positive review, the acceptance difference between white guests and African American guests becomes statistically indistinguishable (p-value $= 0.6813$). Moreover, the blank review experiment also suggests that the existence of a blank review statistically significantly reduces discrimination. While both non-positive and blank reviews can reduce discrimination, our self-claimed information experiment shows that the self-claimed information by guests themselves on tidiness and friendliness fails to reduce discrimination; guests with white–sounding names are 12.8 percentage points more likely to be accepted compared to guests with African American–sounding names (p-value $= 0.019$) even with self-claimed information, which is statistically indifferent from the gap without self-claimed information.

Our paper contributes to the emerging literature on marketplace innovation by providing evidence that peer-generated reviews can reduce discrimination in the sharing economy. Although several recent studies have documented evidence of discriminatory practices on sharing economy platforms such as Airbnb, Uber and Lyft, none provide concrete methods to mitigate these actions and our paper is the first to do so. We show that different types of reviews—positive, non-positive and even a blank review—can reduce discrimination. Moreover, we show that in contrast to peer-generated reviews, self-claimed information cannot reduce discrimination. This result highlights that verifiability and credibility of a review (i.e., a review is linked to a valid transaction on the platform) is crucial for reducing discrimination. Our findings have several implications for sharing-economy platform owners. To attenuate discriminatory behavior and to improve operational efficiency, platform owners should better leverage online reputation systems to encourage and facilitate information sharing among participants. For example, this may be achieved by sending reminders or offering incentives to users to write reviews of one another, especially when one of them is a first-time user. Platform owners should also carefully validate reviews on their platforms by linking reviews to transactions in order to successfully leverage online reviews to reduce discrimination in the sharing-economy platforms.

2. Literature Review

Our research contributes to the literature on discrimination, online reputation systems, information sharing and operational transparency in service operations, and marketplace operations.

2.1. Discrimination

Our paper examines the existence of racial discrimination in the sharing-economy context and investigates how to correct it with field experiments. While discrimination has been widely documented in the literature, most studies do not provide implementable mechanisms to reduce discrimination with very few exceptions that generate mixed and opposing results (see Bertrand and Duflo 2016 and references therein). Bertrand and Mullainathan (2004) and Nunley et al. (2014) investigate the impact of information revealed in resumes on discrimination in labor markets and conclude that additional information revealed in resumes cannot reduce discrimination in labor markets. In contrast, Kaas and Manger (2012) found that discrimination is significantly reduced when a reference letter with soft information regarding productivity is included in a job application. In our paper, we find that peer-generated reviews can reduce discrimination but self-claimed information cannot. This result helps reconcile previous findings in the literature that letters of reference with information on productivity reduce discrimination in labor markets (Kaas and Manger 2012) but self-revealed information in resumes does not (Bertrand and Mullainathan 2004, Nunley et al. 2014).

2.2. Online Reputation Systems

Our paper relates closely to the literature that studies the value of online reputation systems. Past research has found that a seller's good reputation increases the transaction price and success rate in peer-to-peer online auction markets such as eBay (Bajari and Hortacsu 2003). Moreno and Terwiesch (2014) show similar findings in an online service auction market, where service providers bid for projects posted by buyers. The authors find that buyers are willing to accept higher bids from service providers with comparatively better reputations. Bolton et al. (2004) find that online feedback substantially improves efficiency in distributing surplus among buyers and sellers using lab experiments. We complement this line of literature by showing that online reputation systems can reduce social bias (i.e., racial discrimination), which in turn leads to more successful transactions in online peer-to-peer marketplaces.

Moreover, we investigate what aspects of online reviews are critical in reducing discrimination. Previous literature shows that the credibility of reviews affects whether they are effective in shaping consumers' attitudes and influencing their purchasing decisions (Cheung et al. 2012, Luca 2016). Consistent with this literature, we find reviews on Airbnb, which provide credible information about guests, reduce hosts' biases against minorities, but information claimed by guests themselves does not. While a review's credibility is critical in reducing discrimination, the actual content and sentiment of the review is not. Specifically, we find that a positive review, a non-positive review, and a blank review all reduce discrimination. In other words, the existence of a review alone provides credible information regarding the legitimacy of a guest and establishes an inclusive norm that a guest should be accepted irrespective of her race. This can be partially due to the fact that reviews on Airbnb are overwhelmingly positive, as illustrated by recent studies of Airbnb's review system (Zervas et al. 2015, Fradkin 2016). Consequently, hosts may not pay as much attention to the review's content as its presence.

2.3. Operational Transparency and Information Sharing in Service Operations

A stream of literature in service operations studies the effect of operational transparency and information sharing in service industries. Buell and Norton (2011) and Buell et al. (2016) show how transparency of service processes helps improve customer satisfaction and appreciation of the service provided. Several studies have shown how sharing availability information helps customers make more informed purchasing decisions (Allon et al. 2011, Allon and Bassamboo 2011, Gallino and Moreno 2014, Cui et al. 2017). A recurring theme in this literature is the verifiability of the information shared. When the information provided is non-verifiable, a cheap talk equilibrium may emerge where customers ignore the information (Allon and Bassamboo 2011). Allon et al. (2011) investigate conditions under which firms can credibly communicate non-verifiable information in a queueing service setting. In this case, provision of even non-verifiable information can improve company profits and customer utility. Gallino and Moreno (2014) investigate the impact of sharing availability information empirically in a retail setting and find that verifiable availability information provided online increases offline sales and promotes channel integration.

Our papre shows that increased information transparency provided through online reputation systems can help reduce behavioral biases (i.e., racial discrimination) in online service marketplaces. We also show that an important reason why an online reputation system can effectively reduce discrimination is that it is verified and credible. Our results highlight the need for transparent marketplaces: Encouraging sharing rather than concealing information among market participants will help reduce discriminatory behavior and improve market efficiency.

2.4. Marketplace Operations

We study how to leverage peer-generated reviews to improve market efficiency in a sharing-economy marketplace. In this respect our work is related to research in operations management that studies the design and efficiency of sharing-economy business models in various contexts, including bike sharing (Kabra et al. 2015), ride sharing (Cachon et al. 2015, Taylor 2016, Bimpikis et al. 2016), and home sharing (Zervas et al. 2016, Fradkin 2016, Li et al. 2016, Li and Netessine 2018). A common theme of these papers is how various operations levers, such as pricing and matching, can be used to improve the market efficiency and social welfare. For example, Cachon et al. (2015) study how dynamic pricing can be used to improve social welfare in a ride-sharing context, and

Bimpikis et al. (2016) demonstrate how spatial balance of a ride-sharing system can improve its market efficiency and how this balance can be achieved through compensation schemes. Our paper investigates an important but overlooked factor in the sharing economy that could substantially affect market efficiency and social welfare: discriminatory behavior. Our paper sheds light on how a marketplace can be designed to encourage information sharing so that market inefficiency caused by discrimination is reduced.

3. Hypotheses Development

In this section, we develop hypotheses on discrimination in the sharing economy. In particular, we focus on online reviews as a means of reducing discrimination in sharing-economy marketplaces. We also investigate several moderating factors to further understand what aspects of online reviews are critical in reducing discrimination.

3.1. Discrimination in the Sharing Economy

Discrimination against minority groups has been documented in past decades using field evidence in various markets (see Bertrand and Duflo (2016) for a comprehensive review), including labor (e.g., Bertrand and Mullainathan 2004), rentals (e.g., Carpusor and Loges 2006, Hanson and Hawley 2011, Ewens et al. 2014), retail (e.g., Ayres and Siegelman 1995), healthcare (Williams and Mohammed 2009), and education markets (e.g., Milkman et al. 2012). Consistent with recent studies (e.g., Edelman et al. 2017, Ge et al. 2016), we hypothesize discrimination exists in the sharing economy; that is, members of a minority group are treated less favorably than members of a majority group with identical characteristics. Following seminal work by Bertrand and Mullainathan (2004), we use the correspondence method to test and measure discrimination, and we use names to signal minority status (e.g, African American– or white–sounding names). Specifically, we offer the following hypothesis:

Hypothesis 1. *Discrimination exists in the sharing economy: Guests with African American–sounding names experience lower acceptance rates on Airbnb than guests with white–sounding names.*

3.2. Using Online Reviews to Reduce Discrimination in the Sharing Economy

While past studies have documented the existence of discrimination in various contexts, as pointed out by Bertrand and Duflo (2016), most of these studies do not provide implementable mechanisms to reduce discrimination, which is the focus of this paper. In particular, we hypothesize that online reviews will help reduce discriminatory behavior in the context of a sharing economy. Below, we discuss below two primary reasons why reviews would help reduce discrimination.

First, there is a stream of literature that believes discrimination is driven, at least partially, by a lack of information (e.g., Arrow 1973, Kaas and Manger 2012). In the home rental sharing marketplaces, hosts may decide to accept or reject a request based on perceived guest quality, such as safety, reliability, tidiness, and friendliness. Some risks of home-sharing rentals can lead to severe consequences, even though they may be rare, such as personal and property security risks (theft, personal harassment, property damages, etc.).[6] Because the quality of a guest is not fully observable, a host may infer quality based on limited information about the guest, such as name and profile photo, and the host's prior belief about the average quality of a guest of a certain type. For example, if a host believes it is less safe to host African American guests on average, he may decide to reject a request from an African American guest lacking specific information about her. However, if the host sees a positive review of the guest from other hosts, he can then update his belief about the guest's quality and may be more likely to accept her request. In other words, online reviews of guests serve as a quality signal that allows hosts to update their belief of a prospective guest's quality and make more informed decisions, thus reducing discriminatory biases. We formalize this in a theoretical model in Appendix A.

Another reason why online reviews may reduce discrimination is that they help establish a more inclusive normative behavior—that one should not base their acceptance decisions on a guest's race. The fact that other hosts have accepted the guest regardless of her race and have written a positive review can help establish a nondiscriminatory norm in the sharing-economy community. The literature on social norms (Aarts and Dijksterhuis 2003, Cialdini and Goldstein 2004, Schultz et al. 2007) suggests that people tend to follow the normative behavior. In this case, the existence of a social norm will allow hosts to accept a guest more easily when others have done so too. In summary, we hypothesize that reviews, particularly reviews with positive information, will attenuate discriminatory behavior on sharing-economy platforms.

Hypothesis 2. *Positive reviews reduce discrimination: Hosts' positive reviews of guests on Airbnb can reduce the acceptance rate gap between African American guests and white guests.*

3.3. Characteristics of Reviews

In this subsection, we examine what aspects of online reviews are most critical in reducing discrimination. In particular, we examine how a review's credibility, sentiment, and existence affect its ability to reduce discrimination.

3.3.1. Review Sentiment. Even though 96 percent of reviews posted on Airbnb contain positive information that compliments guests (Zervas et al. 2015, Fradkin 2016), occasionally there are non-positive reviews that criticize an experience with a guest; for example, the guest did not

[6] "Airbnb Sued by Guest Who Says a Host Sexually Assaulted Her." *New York Times*. August 2, 2017. "Out-of-Control Partygoers Attack Neighbours and Trash Rented $3 Million Home." Yahoo News. July 1, 2018.

keep the property clean, was not respectful of the neighborhood community, etc. Therefore, we want to test how the sentiment of a review affects its ability to reduce discrimination.

When a review provides positive information about a guest, hosts put more emphasis on the review than on race when making their acceptance decisions. As a result, guests of either race may experience a higher acceptance rate, and the discrimination gap will shrink. However, when a review provides non-positive information about a guest, it is less clear how it may affect the host's acceptance decision. On one hand, the non-positive review may significantly lower expectations of the quality of either type of guests (African American and white guests). As a result, discrimination gap may also shrink. Moreover, the non-positive review can still demonstrate that other hosts have accepted these guests and therefore establish a nondiscriminatory social norm. On the other hand, the non-positive review may exacerbate a host's prejudice against minority guests due to, for example, confirmation bias, thereby leading to wider discrimination gaps (Myrdal 1944, Wason and Johnson-Laird 1972). We hypothesize that:

Hypothesis 3. *Non-positive reviews reduce discrimination: Non-positive reviews of guests on Airbnb reduce the acceptance rate gap between African American and white guests.*

3.3.2. Review Credibility. The literature shows that the credibility of reviews affects whether they are effective in conveying quality (Cheung et al. 2012, Luca 2016). Airbnb has taken several steps to ensure the credibility of reviews posted on its site. First, hosts and guests can only leave a review after the arranged stay is completed. Second, Airbnb actively monitors reviews' validity and quality, and may remove or alter a review when it violates the platform's guidelines aimed at promoting honesty and transparency.[7] As a result, reviews on Airbnb provide relevant, useful and credible information.

On Airbnb, users can also self-identify to be trustworthy by providing information in their accommodation requests. For example, a guest could claim in the message sent to hosts that she is tidy, friendly and reliable. However, this information is usually non-verifiable from the host's perspective. On one hand, such information may be helpful in reducing discrimination since this self-claimed information can serve as a signal that allows hosts to update their belief of a prospective guest's quality and make more informed decisions, thus reducing discriminatory biases. On the other hand, the validity of the self-claimed quality cannot be verified and hosts might not trust this information (Pavlou and Gefen 2004). We hypothesize that:

Hypothesis 4. *Self-claimed information reduces discrimination: Self-claimed quality information in correspondence messages reduces the acceptance rate gap between African American and white guests.*

[7] Airbnb claims, "Our community relies on honest, transparent reviews. We will remove or alter a review if we find that it violates our review guidelines." Please refer to `https://community.withairbnb.com/t5/Hosting/All-About-Reviews-A-Community-Help-Guide/td-p/38099`

3.3.3. Existence of A Review Typical sharing-economy platforms, such as Airbnb and Uber, stipulate that participants can only leave a review after a transaction is completed. Therefore, a review on these platforms conveys two types of information: First, on Airbnb the existence of a review conveys that the guest has stayed with a host before and the host has written a review. Second, the content of the review shows how satisfied the host is about the guest and the stay. We would like to separate the effect of existence of a review from that of the content. In other words, we would like to test whether the existence of a review alone (a blank review without any content) would reduce discrimination or not. The fact that a review exists on a guest's profile provides several important quality signals about the guest. First, the guest has prior experience with the platform and with other hosts; therefore, he might be more familiar with the process than a first-time guest. Second, the guest's identity is truthful; otherwise, the host would have reported the incident to the platform operator and the guest account would have been suspended. Third, there was no serious breach of the contract caused personal or property security damages; otherwise, the host would have reported the incident or sought legal settlement and the guest account would have been suspended. Thus, even without explicit discussion of a guest's manner or personality, reviews provide valuable information about the guest's legitimacy. Besides the valuable information a review's existence provides, it also helps to establish a norm of inclusive behavior because the guest has been accepted by other community members. Therefore, we hypothesize that:

Hypothesis 5. *Review existence alone reduces discrimination: Blank guest reviews reduce the acceptance rate gap between African American guests and white guests.*

4. About Airbnb

We conduct the field experiments on Airbnb.com. Airbnb is a sharing-economy marketplace that connects hosts who have empty rooms to potential renters. Airbnb defines itself as "a trusted community marketplace for people to list, discover, and book unique accommodations around the world." Hosts on Airbnb can list and rent their properties for a processing fee. Founded in 2008, Airbnb's marketplace has expanded exponentially in the last few years. As of 2016, Airbnb had more than 2 million listings, which have collectively served 60 million guests in 191 countries.[8]

On Airbnb, hosts create profiles for themselves and their properties. Prospective hosts can list their entire apartments or just spare bedrooms on the Airbnb platform, choose their own prices and set guidelines for guests. Information on each host page includes a profile photo, listing information, rental requirements, reviews written by previous guests who have stayed at the host's properties, and Airbnb-certified contact information. Figure 5 in the Appendix shows the information displayed on a typical Airbnb listing. Guests also need to create personal profiles in order to request a booking.

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

Each guest page contains a profile photo, the guest's first name, reviews written by hosts about past stays, and Airbnb-certified contact information. Figure 6 in the Appendix shows a typical Airbnb guest profile page. Airbnb encourages users to verify their identification, email address, phone number and social network profiles.

When they need an accommodation, guests can search for and choose from available listings at the destination for the selected travel dates. If guests have questions regarding a lisiting's rental requirements, they can contact the host by clicking the "Contact Host" button. Upon receiving contact information, hosts can preapprove a guest. Guests can also click the "Request to Book" button to directly send a booking request. Hosts can approve or reject this booking request. Hosts may ask guests to provide further information regarding their travel purpose or contact verification. Once guests and hosts have confirmed travel dates and expectations, guests submit their payment to Airbnb, which holds the payment until 24 hours after the reservation begins.

4.1. Reviews on Airbnb

Airbnb has built a reputation system that enables both guests and hosts to post reviews of each stay on the platform. Once a transaction is complete, the guest and host have 14 days to leave a review for each other. If both of them leave a review within 14 days, they can see each other's review immediately after the review is created. If only one of them leaves a review, this review can only be seen 14 days after the check-out date. In creating a review, users are given an opportunity to rate the experience and then write about it in their own words. Reviews on hosts, which are displayed on host profiles, contain both the rating as well as the textual content of the review. However, unlike other platforms, *reviews on guests displayed on guest profiles only contain the textual content.* Figure 6 (see Appendix) shows an example of a guest profile, and it is evident that reviews on guests are displayed on an easily visible spot on their profiles. Since Airbnb monitors the transaction and only allows users to create reviews after a successful transaction, past reviews serve as a credible signal of user quality. When it includes a review, a guest's profile page proves that the guest has had at least one successful transaction with another host and is familiar with the accommodation process.

4.2. Platform Data

We collect data from all listings offered in several metropolitan areas with high Airbnb usage. This data comprises two types: listing characteristics and host characteristics. First, the listing characteristics include the type of room offered, number of bedrooms, number of reviews, and listing location. Second, we gather the demographics of the host associated with each listing in our experiment. We employ research assistants to classify hosts' race (white, African American, Asian, Hispanic, Arab, unidentifiable), gender (female, male, unidentifiable), and age (0–30, 30–45, 45–60,

Table 1 Experiment Design Summary

	Positive Review Experiment	Non-positive Review Experiment	Self-Claimed Information Experiment	Blank Reivew Experiment
Design	White or African American Names × No or Positive Review	White or African American Names × No or Non-positive Review	White or African American Names × No or Self-Claimed Information	White or African American Names × No or Blank Review
Guest Accounts per Condition	2 guest accounts per condition	2 guest accounts per condition	4 guest accounts per condition	4 guest accounts per condition
Guest Account Verification	Email Address and Phone Number	Email Address and Phone Number	Email Address and Phone Number	Email Address, Phone Number and Government ID
White Names	Scott Baker, Colin Murphy	Scott Baker, Colin Murphy	Scott Mueller, Scott Baker, Colin Murphy, Colin Moore	Scott Mueller, Scott Baker, Colin Murphy, Colin Moore
African American Names	DeAndre McCray, DeShawn Washington	DeAndre McCray, DeShawn Washington	DeAndre McCray, DeAndre Jackson, DeShawn Washington, Tyrone Washington	DeAndre McCray, DeAndre Jackson, DeShawn Washington, Tyrone Washington
Experiment Date	Sept. 4, 2016 to Sept. 27, 2016	Oct. 23, 2016 to Nov. 21, 2016	Jul. 27, 2017 to Aug. 12, 2017	Mar. 05, 2018 to Apr. 10, 2018
Experiment City	Chicago, Boston, Seattle	Boston, Seattle	Boston, Seattle, Austin	Los Angeles
Planned Sample Size	1200	400	1200	500
Actual Sample Size	598	250	660	293

Note: For the first experiment, the difference between the planned sample size and the actual sample size is due to listings' unavailability and Airbnb banning our accounts. In particular, the planned sample size is $1,200$. Excluding the unsent listings due to suspended accounts leaves us with 856 listings. In all other three experiments, because we had backup accounts in case the original accounts were suspended, the difference between planned and actual sample size is only due to listings' unavailability. To avoid being suspended and ensure an adequate sample size, we use more names in the other three experiments. Based on our observations, account suspension happens randomly across different names (i.e., accounts with African American names are not more likely to be suspended) and searches on two famous Airbnb host forums (i.e., airhostsforum.com and community.withairbnb.com) demonstrate that account suspension does not trigger hosts' discussion of accounts with the names used by us.

beyond 60, unidentifiable). We code characteristics as "unidentifiable" when the photo has multiple people in it or the picture does not show a person. Specifically, we hire two graduate students as research assistants to evaluate each image; if there is a disagreement, we code it manually.

5. Experimental Design

We conduct four 2×2 field experiments using listings from Boston, Chicago, Seattle, Austin and Los Angeles. The experiment design details are presented in Table 1. The first field experiment tests the effect of positive reviews on discrimination. The remaining three experiments test the effect of review sentiment, review credibility and review existence without content in reducing discrimination. All experiments are conducted with the approval of the Institutional Review Board, and we discuss how we protect experiment subjects in Appendix B.

5.1. Positive Review Experiment

We conduct the positive review experiment in September 2016. In this experiment, we create eight fictitious guest accounts with and without positive reviews under names that signal different races to test the existence of discrimination as well as the impact of one positive review on discrimination. In order to represent an average Airbnb account, all eight accounts have verified their email addresses and phone numbers (see Figure 7a in the Appendix for an example).[9] The first four fictitious guest accounts do not have reviews and are identical except for names: two accounts have African American–sounding names and two have white–sounding names.

All names are drawn from a pool of most frequently used white and African American–sounding names (Bertrand and Mullainathan 2004) to signal African American or white race. To avoid potential confounds, we use multiple names in the experiment (Wells and Windschitl 1999). To validate the selected names, we perform an image search on Google and verify that most of the people who appear in the results are of a race consistent with the name we've chosen. We also conduct a survey on Amazon Mechanical Turk to assure that average Airbnb users can recognize the race from the names. Table 9 in the Appendix shows that more than 80 percent of the survey participants can correctly identify the race for all the eight names we use.

The second four fictitious guest accounts are identical to the prior four accounts except that they have one positive review written by the same host at the same time. Figure 2a shows the content of this positive review. *Note that the reviews on guests do not contain any ratings except the textual content on Airbnb*, and therefore identical positive reviews represent reviews with the identical positive content. To generate these identical positive reviews, we create one fictitious host account and use this host account to leave a review on the fictitious guest accounts. The fictitious host is named Scott and located in a midwestern city.[10] Because Airbnb allows hosts and guests to post a review of each other after the guest has checked out, we create a transaction between our host and each of our guest accounts on Airbnb and we write an identical positive review of all guest accounts after the checkout date. Because Airbnb only displays the month and year when a review is written, the time of review is the same across our guest accounts if the review dates are within a month of each other. Since Airbnb requires a guest to upload a picture prior to contacting hosts,

[9] Figure 7b in the Appendix shows verification status of 1,000 randomly chosen guest accounts within the experimental cities. It is clear that the most popular verification among these accounts is phone numbers and email addresses. In the blank review experiment, we further verify the government ID for each fictitious account.

[10] The ideal case would be to randomize the name of the host who leaves the review. But because one can only leave a review after a transaction, it is difficult to randomize host names in the experiment. We choose a white–sounding name since, because more than 80 percent of hosts within our sample are white.

Figure 1 Messages in the Experiment

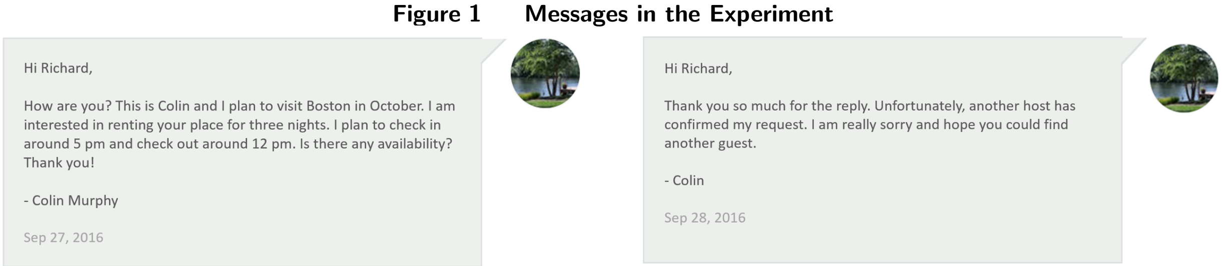

(a) Sent Message in the Experiment

(b) Reply Message in the Experiment

all guest accounts also include a scenery picture, as shown in Figure 1a. We use a neutral-looking scenery photo to prevent the content of a photo from confounding our results.

We then randomly assign Airbnb hosts in our sample to the fictitious guest accounts with different names (i.e., white or African American–sounding names) and different reviews (i.e., no review or one review) and send accommodation requests from our guest accounts to prospective hosts. Each prospective host will be contacted once at most to avoid a host seeing two identical messages from different guest accounts. If a host owns multiple properties, we inquire about only one of the properties. The accommodation requests are sent from guest accounts with and without reviews simultaneously to ensure comparability. We send messages to request a stay of two consecutive nights with check-in dates about three weeks from the date of initial contact.[11] Figure 1a presents the content of a sent message. The name, city and date information vary by request. Since our guest accounts are identical except for having different names and number of reviews, by comparing acceptance rates across guests with different names and review conditions, we can test whether there exists significant discrimination and whether one positive review reduces discrimination.

After sending messages to hosts, we check for the host's reply at least five times: about 5 hours, 10 hours, 24 hours, 48 hours and 5 days after the request was sent. When a host replies to our guests, we follow up with an immediate response stating that we found another place to stay so that the host does not hold the inventory for our fictitious guests. Figure 1b presents the content of our reply. We record host responses within five days after sending requests. We code each response in three categories: "Decline" if the host declines the request; "Accept" if the host accepts the request; "Further information" if the host asks for further information. Following Edelman et al. (2017), we focus on the "Accept" response; our results do not change qualitatively if we switch to a broader definition of acceptance by considering "Further information" as acceptance.

[11] We adhere to the following procedure to find available nights from the host's calendar: We first find two available consecutive nights available given the range of check-in dates. If there are fewer than two nights available, we find one available night. If the host has a "3 nights minimum stay" rule, we find three consecutive nights. If all of the above fail, we do not send a request for the listing.

Figure 2 Review Information Displayed on the Fictitious Guest's Profile Page

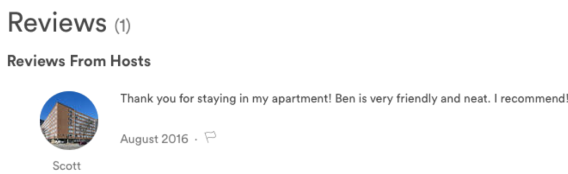

(a) Positive Review Information

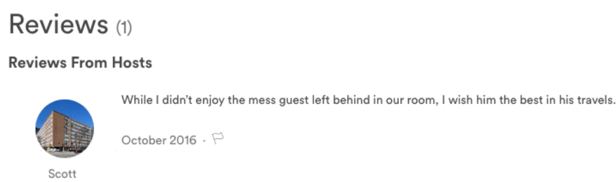

(b) Non-positive Review Information

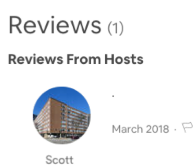

(c) Blank Review

Note: Because the location of the fictitious guest may indicate the location of the authors' institutions, we remove it for the review process.

5.2. Other Experiments

5.2.1. Non-positive Review Experiment. We conduct the non-positive review experiment in October and November of 2016 to test how a non-positive review affects discrimination. In this experiment, we create another eight fictitious guest accounts with identical profile pictures, the same profile setup as the previous experiment: white or African American–sounding names and with a non-positive review or none at all. Similar to the positive review experiment, all eight accounts have verified email addresses and phone numbers. The major difference from the previous experiment is that four accounts in this experiment have a non-positive review instead of a positive review. Figure 2b shows the content of the non-positive review. Similarly, we generate this review after creating a transaction between our fictitious host account (the same one used in the previous experiment) and guest accounts; these non-positive reviews on guest profiles do not contain ratings.

We then randomly assign listings to the fictitious guest accounts and send out accommodation requests to hosts in Boston and Seattle. The listings and hosts are mutually exclusive from those used in the previous experiment. We follow the same rules to select available hosts and dates from their calendars, reply to hosts and record host responses. This experiment allows us to detect whether discrimination exists (at a different time with a different set of hosts) and whether a non-positive review helps attenuate discrimination.

5.2.2. Self-Claimed Information Experiment. In July and August 2017, we conduct another field experiment to test whether self-claimed information from guests helps attenuate discrimination. In this experiment, we employ a 2×2 design and randomly assign 16 fictitious guest accounts with verified phone numbers and email addresses to (a) white– or African American–sounding names and (b) no information or self-claimed information in the accommodation requests. The accounts in the no information condition are similar to those in the no review condition of previous experiments: eight fictitious guest accounts—four with white–sounding names and four with African American–sounding names—are created without reviews and their accommodation requests have the basic content as in Figure 1a. In the self-claimed information condition, the eight fictitious guest accounts do not have a review, and their accommodation requests have the basic request plus the sentence, "I am a tidy and friendly person. I like to keep places clean and organized. Let me know if you have any questions." This additional sentence represents the self-claimed information from the guest. This sentence signals the guest's tidiness and friendliness, which is consistent with the key information contained in the positive review shown in Figure 2a. This experiment uses a total of 8 names (instead of 4 names) and 16 accounts (instead of 8 accounts).

We then randomly match a set of listings and the corresponding hosts from Boston, Seattle and Austin to our fictitious guest accounts. Again, the listings and hosts are mutually exclusive from the first two experiments. We follow the same rules to select available hosts and dates from their calendars, send requests, record hosts' responses and measure acceptance rates. This experiment allows us to test if discrimination exists and whether information provided by guests themselves about their tidiness and friendliness helps attenuate discrimination.

5.2.3. Blank Review Experiment. In March and April 2018, we conduct the fourth field experiment to test whether the existence of a review without content (i.e., a blank review) can help reduce discrimination. In this experiment, similar to the self-claimed information experiment, we employ a 2×2 design and create another 16 fictitious guest accounts and randomly assign 4 accounts to each condition. Each condition is a combination of the account names (i.e., white or African American–sounding names) and account review status (i.e., no review or a blank review). The major difference between this experiment and the previous positive review and non-positive review experiments are that (a) these 16 accounts have verified government IDs in addition to verified email addresses and telephone numbers as in previous experiments, and (b) we use our fictitious host account to leave an identical blank review without content instead of a positive or non-positive review. Figure 2c provides an example of this blank review. Similarly, we generate

Table 2 Response Summary Statistics by Guest Race in the Positive Review Experiment

Guest Race	No. of Listings	No. of Requests Sent	No. of Requests Accepted	Prob. of Acceptance
Panel A: Summary Statistics for Guests Without Reviews				
White	272	188	90	47.9%
African American	272	188	54	28.7%
Average Differences				19.2%
Difference 95% Confidence Interval				(9%, 29%)
P-value of Proportion Test				0.0002
Panel B: Summary Statistics for Guests with Positive Reviews				
White	153	105	59	56.2%
African American	159	117	68	58.1%
Average Differences				-1.9%
Difference 95% Confidence Interval				(-16%, 12%)
P-value of Proportion Test				0.8774

this review after creating a transaction between our fictitious host account (the same one used in the previous experiment) and guest accounts; these blank reviews on guests do not contain ratings. We then randomly assign 400 listings in Los Angeles to the fictitious guest accounts and send out accommodation requests to the corresponding hosts.[12] We follow the same rules to select available hosts and dates from their calendars, reply to hosts and record host responses. This experiment allows us to detect whether discrimination exists in March and April 2018 in Los Angeles and whether a blank review helps reduce discrimination.

6. Main Results from the Positive Review Experiment

In this section, we show the results from the positive review experiment. Our findings confirm the existence of racial discrimination and reveal that positive reviews can help attenuate discrimination.

6.1. Existence of Racial Discrimination

Panel A of Table 2 displays the summary statistics of listing availability and request acceptance rates for guest accounts with no review in the first experiment, stratified by guest race. Panel A shows that without reviews, a guest with a white–sounding name is accepted with a 47.9 percent probability. By comparison, a guest with an African American–sounding name only has a 28.7 percent probability of being accepted. In other words, a guest with an African American–sounding name is 19.2 percent less likely to be accepted.

Using a nonparametric proportion test, we demonstrate that the 19.2 percent difference in the acceptance rate is statistically significant (p-value $= 0.0002$). Panel A of Table 2 shows that the p-value of the proportion test is 0.0002.[13] Given that hosts are randomly assigned to guest accounts

[12] We choose hosts from Los Angeles in this fourth experiment since we have exhausted available and eligible hosts in cities used in previous experiments.

[13] Note that we start with the nonparametric proportion test because it requires fewer assumptions on the data generating process compared to regression-based methods, e.g., Logit model.

Table 3 Main Results of the Positive Review Experiment

	Dependent Variable: Acceptance					
	Linear Probability			*Logit*		
	(Review=0)	(Review=1)	(All)	(Review=0)	(Review=1)	(All)
	I	II	III	IV	V	VI
White	0.200^{***}	-0.045	0.186^{***}	0.947^{***}	-0.235	0.900^{***}
	(0.048)	(0.063)	(0.048)	(0.236)	(0.312)	(0.233)
Review			0.212^{***}			1.013^{***}
			(0.055)			(0.265)
White × Review			-0.236^{***}			-1.140^{***}
			(0.079)			(0.376)
Host Controls	Yes	Yes	Yes	Yes	Yes	Yes
Listing Controls	Yes	Yes	Yes	Yes	Yes	Yes
Request Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	376	222	598	376	222	598

Note: *p<0.1; **p<0.05; ***p<0.01

and the guest accounts are identical except for names, the empirical evidence suggests that guest race causally affects the acceptance rates for their accommodation requests. In other words, racial discrimination exists.

To formally test the existence of racial discrimination, we use the following specification to analyze how guest race affects the acceptance rate of a request while controlling for characteristics of the host, the listing and the specific request:

$$\text{Accept}_i = f(\alpha + \beta \text{Race}_i + X_i \gamma + \epsilon_i). \tag{1}$$

We choose two specifications for the choice function $f(\cdot)$: the linear probability model and the ogit model. X_i includes host characteristics (i.e., host gender and host race), listing characteristics (i.e., room type, number of bedrooms, number of reviews and listing location) and request characteristics (i.e., the date and length of the request).

Columns I and IV of Table 3 present the results based on linear probability and Logit models respectively. The findings are robust: a guest with a white–sounding name has a higher acceptance rate. In particular, after controlling for host, listing and request characteristics, being associated with a white–sounding name can causally increase the acceptance rate by 20.0 percentage points. This result is consistent with and further confirms the finding of Edelman et al. (2017) that racial discrimination exists on Airbnb.

To summarize, the results from no-review guest accounts demonstrate that racial discrimination exists on Airbnb. Guests with white–sounding names are 19 percentage points more likely to be accepted than guests with African American–sounding names. This empirical evidence supports Hypothesis 1.

6.2. The Effect of One Positive Review

Panel B of Table 2 presents the summary statistics of listing availability and request acceptance rates for guest accounts with a positive review, stratified by guest race. Panel B shows that, with a positive review, racial discrimination is significantly attenuated: white and African American guests receive non-statistically distinguishable acceptance rates (56.2 percent and 58.1 percent respectively with p-value $= 0.8774$). This suggests that *one* positive review significantly attenuates discrimination on Airbnb.

A potential concern is whether this null result is driven by the lack of power in our experiment. Prior to the experiment, we conduct a power analysis based on the estimation results of Edelman et al. (2017) and select sample sizes accordingly.[14] We also conduct a post-experiment power analysis to show the power of detecting a significant discrimination effect given our sample size if a positive review does not reduce discrimination. In particular, given the discrimination effect size of 19.2 percent (based on the no review condition) and the sample size of 222 (for the positive review condition), we have 90 percent power to detect the discrimination at a 0.05 significance level and 95 percent power at a 0.1 significance level. This indicates that the null result with p-value $= 0.8774$ observed in our experiment is unlikely to be due to the lack of power.[15]

As expected, a positive review provides credible information about guest quality and improves the acceptance rate for guests of both races. Guests with white–sounding names experience an increase from 47.9% to 56.2% (p-value $= 0.0968$) in the acceptance rate, while guests with African American–sounding names experience a more substantial increase, from 28.7% to 58.1% (p-value < 0.00001). Figure 3 presents this change visually.

We then follow the same specification as in Equation (1) to formally analyze the effect of a review on discrimination while controlling for host, listing and request characteristics. Columns II and V of Table 3 summarize the results. Both the linear and the Logit models show that in the presence of a positive review, the acceptance rate is the same across different races, so there is no sign of discrimination. In order to formally identify whether having one positive review significantly reduces discrimination (Gelman and Stern 2006), we estimate this effect using the following specification:

$$\mathrm{Accept}_i = f(\alpha + \beta_0 \mathrm{Race}_i + \beta_1 \mathrm{Review}_i + \beta_2 \mathrm{Review}_i \times \mathrm{Race}_i + X_i \gamma + \epsilon_i), \tag{2}$$

[14] The sample size in this experiment was selected based on the findings of Edelman et al. (2017) of an effect from 42 percent (African American accounts) to 50 percent (white accounts) and a power level of 0.8. Based on the one-sided proportion test and these estimates, we need 300 observations per treatment condition at a 0.05 significance level and 222 observations per treatment condition at a 0.1 significance level. Note that we use the one-sided power calculation because Edelman et al. (2017) show that accounts with African American–sounding names have a lower acceptance rate than accounts with white–sounding names.

[15] The post-experiment power drops to 60 percent at a 0.1 significance level if the discrimination effect size is 10 percent instead of 19.2 percent. Therefore, we claim throughout the paper that the empirical evidence shows that discrimination is significantly attenuated (instead of being eliminated) with one positive review, even though the observed discrimination effect size is not significantly different from 0 when there is one positive review.

Figure 3 Results for the Positive Review Experiment

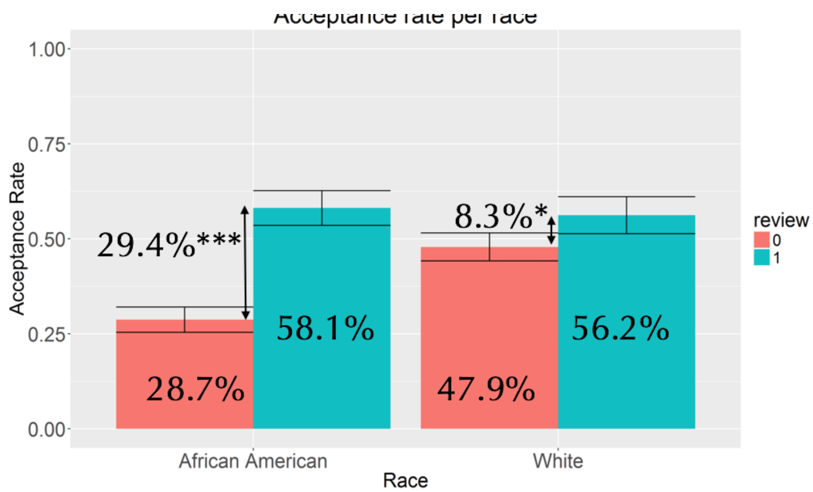

where $\text{Review}_i = 1$ indicates that the guest account i has one positive review and $\text{Review}_i = 0$ otherwise.

Columns III and VI of Table 3 display the results based on linear probability and Logit models respectively. The coefficient of the interaction of having a positive review and having a white–sounding name is significantly negative, which suggests that a positive review causally reduces discrimination.

To summarize, we find that one positive review can significantly reduce racial discrimination. This empirical evidence supports Hypothesis 2.

6.3. Robustness of Results

We next show the robustness of our results: the impact of a review on discrimination is not driven by a particular subset of hosts or listings. We split our samples by listing and host characteristics and follow the Logit regression specification in Equation (2) to test the existence of discrimination in each subsample. Table 4 presents the estimation results. Columns I and II of Table 4 show that discrimination exists for both low-review (LR) hosts and high-review (HR) hosts, and a positive review significantly reduces discrimination for both types of hosts. Low review is defined as the number of reviews being lower than the median (i.e., 13). We also split the sample by city and demonstrate that our results are robust across the three cities (i.e., Chicago, Boston and Seattle), shown in Table 8 in the Appendix.

Airbnb classifies listings as "Entire Apartment," "Private Room" and "Shared Room." Private and shared rooms require guests to stay with hosts, whereas the entire-apartment type does not. We classify "Private Room" and "Shared Room" as shared listings and "Entire Apartment" as non-shared listings. Columns III and IV of Table 4 show that discrimination exists for both non-shared and shared rooms and a positive review significantly mitigates discrimination for both types

Table 4 Positive Review Experiment Results Moderated by Listing Characteristics

	Dependent Variable: Acceptance							
	(LR)	(HR)	(Non-Shared)	(Shared)	(Female)	(Non-Female)	(White)	(Non-White)
	I	II	III	IV	V	VI	VII	VIII
White	0.156^{**}	0.248^{***}	0.167^{***}	0.250^{***}	0.201^{***}	0.191^{***}	0.222^{***}	0.164
	(0.070)	(0.069)	(0.062)	(0.077)	(0.070)	(0.068)	(0.053)	(0.123)
Positive Review	0.169^{**}	0.311^{***}	0.198^{***}	0.313^{***}	0.266^{***}	0.174^{**}	0.258^{***}	0.161
	(0.077)	(0.083)	(0.069)	(0.097)	(0.082)	(0.078)	(0.062)	(0.142)
White $\times$	-0.237^{**}	-0.309^{***}	-0.213^{**}	-0.378^{***}	-0.277^{**}	-0.234^{**}	-0.293^{***}	-0.220
Positive Review	(0.110)	(0.113)	(0.100)	(0.130)	(0.117)	(0.109)	(0.087)	(0.190)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	305	293	357	241	305	293	479	119

Note: $*p < 0.10; **p < 0.05; ***p < 0.01$. Standard errors are robust. The characteristics in Panel C include all the host, listing and request characteristics in Table 3.

of rooms. Note that the magnitude of discrimination without reviews is larger for shared rooms. This provides evidence that when hosts need to stay with guests, they become more careful in inferring guests' quality and tend to be more discriminative.

We also split our samples based on hosts' gender. We classify the host gender as female or non-female (including male and unidentified hosts). Columns V and VI of Table 4 show that both female and non-female hosts are likely to discriminate against guests based on race, as signaled by their names in the absence of reviews. One positive review on the guest profile significantly reduces discrimination for both female and non-female hosts.

Last, we split our samples based on hosts' race. We classify the host of each listing as white and non-white (including non-white hosts and unidentified hosts). Column VII of Table 4 shows that white hosts are likely to discriminate against guests based on race signaled by names; one positive review significantly reduces such discrimination. Column VIII of Table 4 shows that the discrimination effect for non-white hosts is positive but insignificant; because the discrimination effect is insignificant, the value of a positive review in reducing discrimination is also insignificant. Note that this result does not necessarily mean that non-white hosts do not exhibit discriminatory behavior because the sample size is small ($N = 119$) and the null result may be due to low statistical power. The point estimates of discrimination effects in Columns VII and VIII of Table 4 may suggest that non-white hosts are less likely to discriminate; therefore, one positive review helps attenuate discrimination more for white hosts.

Next, we provide evidence on whether hosts require additional information upon receiving requests from guest accounts of different races, with and without reviews, from accounts in different conditions. On Airbnb, hosts can reply to guests' sent messages and ask for more information about

Figure 4 Information Requested in the Positive Review Experiment

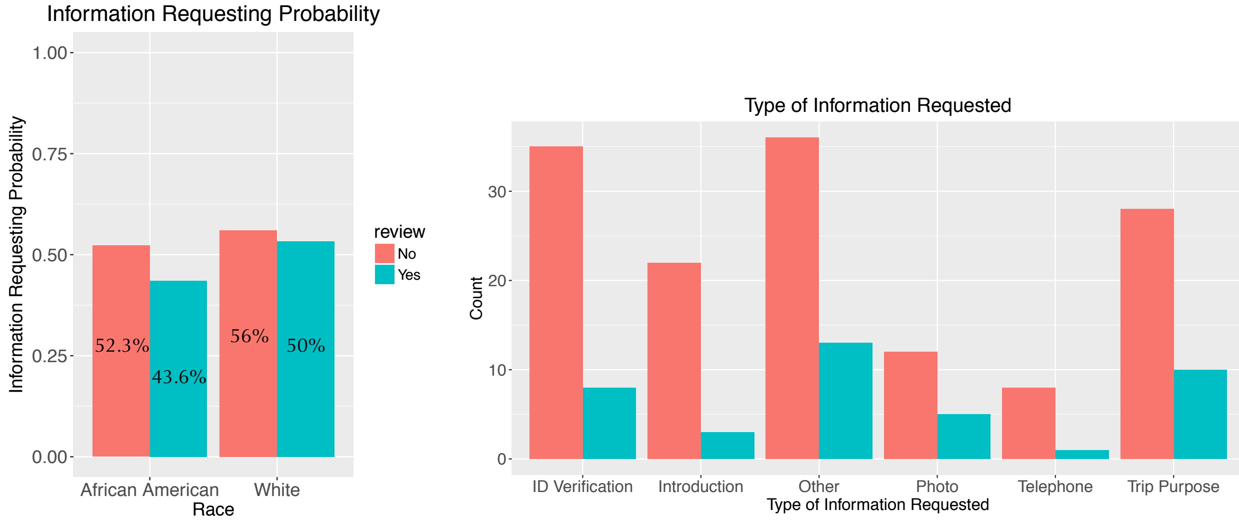

(a) Information Request Probability

(b) Type of Information Requested

the traveling party or the purpose of the intended trip. Some hosts who do not immediately accept the accommodation use this opportunity. In this case, if a review provides additional valuable information to hosts, a host would be more likely to request additional information from a guest with no review. Moreover, if we believe that hosts discriminate against African American guests because of the lack of information, we would observe that the probability of requesting information conditional on rejecting the offer drops more severely for African American accounts than for white accounts. Figure 4a shows that this is indeed the case: the probability of information request has dropped from 56 percent to 50 percent for white accounts (i.e., the drop is 6 percent, with p-value < 0.1), and from 52.3 percent to 43.6 percent for African American accounts (i.e., the drop is 8.7 percent, with p-value < 0.05). This indicates that reviews substitute for guest quality information, which is consistent with statistical discrimination. Moreover, we classify information requested by the hosts into the following categories: ID Verification, Photo, Self-Introduction, Telephone, Trip Purpose and Other. Figure 4b shows the number of times each type of information is requested for both no review and positive review conditions.

In summary, we confirm that our main result—i.e., discrimination exists and a positive review significantly reduces discrimination—is robust across different types of listings and hosts. We also show that hosts are less likely to request guest quality information when guests have a review because reviews signal the quality information that hosts find useful in making their accommodation decisions.

7. Results from Other Experiments

In this section, we demonstrate the results from the non-positive review experiment, the self-claimed information experiment, and the blank review experiment. Our findings demonstrate that both a non-positive review and a blank review can help reduce discrimination, while self-claimed information cannot.

7.1. Non-positive Review Experiment

Panels A and B of Table 5 show the number of requests sent and acceptance rates for each condition in the non-positive review experiment. Panel A demonstrates that when there are no reviews, discrimination exists. Guests with white–sounding names are accepted 62.7 percent of the time, which is 21.4 percent higher than guests with African American–sounding names (p-value = 0.02874). The baseline acceptance rates in the non-positive review experiment are different from those in the positive review experiment. This may be due to the fact that these experiments were conducted with hosts from different cities and at different times. This result is consistent with the 19.2 percent difference we found in the previous round of experimentation. However, with one non-positive review, discrimination becomes significantly attenuated. The acceptance rates for white and African American guests are 58.2 percent and 57.4 percent, respectively, and the difference is statistically insignificant (i.e., p-value = 0.6813).[16]

Comparing Panel A and B of Table 5 shows that a non-positive review increases the acceptance rate of guest accounts with African American–sounding names by 16.1% (p-value = 0.0965). This could be because hosts' prior beliefs about an average African American guest's quality are even lower than the quality revealed by the non-positive review. This could also be because the existence of a review shows that the guest has completed a transaction with other hosts. Such completion signals several qualities of the guest. It shows that the guest is likely more familiar with the process than a first-time guest. It also shows that the guest's identity is truthful because otherwise the host would have reported the incident to the platform in which case the guest account would have been suspended. In other words, even though the guest may be messy, at least he is safe to host and is familiar with the process. The completion of a transaction may also demonstrate that the guest has been accepted by other hosts in the community, which establishes a social norm that the guest should be accepted. In Section 7.3, we formally tease apart the impact of a review's existence (i.e., the signal that a guest has completed a transaction on the platform) and its content in reducing discrimination.

[16] We conduct a post-experiment power analysis for this review condition. The identified effect of racial discrimination in the non-review condition is a 21.4 percentage reduction from a 62.7 percent acceptance rate for white accounts to a 41.3 percent acceptance rate for African American accounts, and the total sample size of the review condition is 128. Assuming the non-positive review does not change the discrimination effect, we have a power of 79 percent and 69 percent to detect the discrimination effect at the 0.1 and 0.05 significance levels. This shows that the null results we observe are unlikely to be driven by the lack of power.

Table 5 Response Summary Statistics and Regression Results from Non-positive Review Experiment

Guest Race			No. of Listings	No. of Requests Sent	No. of Requests Accepted	Prob. of Acceptance
			Panel A: Summary Statistics for Guests without Reviews			
White			100	59	37	62.7%
African American			100	63	26	41.3%
Average Differences					21.4%	
Difference 95% Confidence Interval					(2.4%, 40%)	
P-value of Proportion Test					0.0287	
			Panel B: Summary Statistics for Guests with Non-positive Reviews			
White			100	67	39	58.2%
African American			100	61	35	57.4%
Average Differences					-1.9%	
Difference 95% Confidence Interval					(-17.1%, 18.8%)	
P-value of Proportion Test					0.6813	

Panel C: Non-positive Review Experiment Regression Results

	Linear Probability			*Logit*		
	(Review=0)	(Review=1)	(All)	(Review=0)	(Review=1)	(All)
	I	II	III	IV	V	VI
White	0.219^{**}	0.036	0.212^{**}	1.079^{**}	0.145	0.979^{**}
	(0.086)	(0.090)	(0.087)	(0.426)	(0.385)	(0.403)
Non-positive Review			0.150^{*}			0.699^{*}
			(0.085)			(0.392)
White ×			−0.198			-0.928^{*}
Non-positive Review			(0.123)			(0.563)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	122	128	250	122	128	250

Note: $*p < 0.10; **p < 0.05; ***p < 0.01$. Standard errors are robust. The characteristics in Panel C include all the host, listing and request characteristics in Table 3.

Last, we follow the same specifications as in Table 3 and formally test the existence of discrimination and the impact of one non-positive review on discrimination; the results are shown in Panel C of Table 5. Columns I and IV demonstrate that in absence of review information, having a white–sounding name can causally increase the acceptance rate by 21.9 percentage point; Columns II and V show that with a non-positive review, having a white–sounding name no longer significantly affects the acceptance rate. Columns III and VI suggest that having a non-positive review significantly reduces the racial discrimination caused by names.

To summarize, we show that a non-positive review can also effectively reduce racial discrimination, which supports Hypothesis 3.

7.2. Self-Claimed Information Experiment

We are also interested in whether the credibility of reviews on Airbnb is a key reason that reviews can reduce discrimination. In particular, we test whether information that guests provide can reduce discrimination. Panel A of Table 6 shows that when there is no self-claimed information, discrimination exists. Guests with white–sounding names are accepted 53.3 percent of the time,

Table 6 Response Summary Statistics and Regression Results from Self-Claimed Information Experiment

Guest Race	No. of Listings	No. of Requests Sent	No. of Requests Accepted	Prob. of Acceptance
Panel A: Summary Statistics for Guests without Reviews				
White	300	165	88	53.3%
African American	300	163	69	42.3%
Average Differences			11.0%	
Difference 95% Confidence Interval			(1%, 21%)	
P-value of Proportion Test			0.046	
Panel B: Summary Statistics for Guests with Self-Claimed Information				
White	300	181	89	49.2%
African American	300	151	55	36.4%
Average Differences			12.8%	
Difference 95% Confidence Interval			(2%, 23%)	
P-value of Proportion Test			0.019	

Panel C: Self-Claimed Information Experiment Regression Results

	Linear Probability			*Logit*		
	(Review=0)	(Review=1)	(All)	(Review=0)	(Review=1)	(All)
	I	II	III	IV	V	VI
White	0.106**	0.156***	0.106**	0.630**	0.690***	0.457**
	(0.086)	(0.090)	(0.087)	(0.426)	(0.385)	(0.403)
Self-Claimed			−0.069			−0.307
Information			(0.055)			(0.240)
White × Self-			0.042			0.196
Claimed Information			(0.076)			(0.329)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	328	332	660	328	332	660

Note: $*p < 0.10; **p < 0.05; ***p < 0.01$. Standard errors are robust. The characteristics in Panel C include all the host, listing and request characteristics in Table 3.

while guests with African American–sounding names are accepted 42.3 percent of the time. The difference is 11.0 percent (p-value $= 0.046$). The magnitude of the discrimination effect is different in this experiment compared to the previous two, which may be driven by different cities or experiment times. Nevertheless, discrimination still exists in this experiment. Moreover, Panel B shows that with self-claimed information, discrimination still exists; the acceptance rates for white and African American guests are 49.2 and 36.4 percent respectively, and the difference is 12.8 percentage points (p-value $= 0.019$). This provides initial evidence that self-claimed information about the guest's friendliness and tidiness does not reduce discrimination.

Similarly, following the specification in Table 3, we formally test the existence of discrimination and the impact of self-claimed information on discrimination; the results are displayed in Panel C of Table 6. Columns I and IV show that in the absence of self-claimed information and reviews, having a white–sounding name can causally increase the acceptance rate by 13.9 percentage points, controlling for host, listing and request characteristics. Columns II and V show that the acceptance rate of guests with white–sounding names is 14.8 percentage points higher than that of guests

Table 7 Response Summary Statistics and Regression Results from Blank Review Experiment

Guest Race			No. of Listings	No. of Requests Sent	No. of Requests Accepted	Prob. of Acceptance
			Panel A: Summary Statistics for Guests without Reviews			
White			150	87	33	37.9%
African American			150	89	18	20.2%
Average Differences					17.7%	
Difference 95% Confidence Interval					(4.3%, 31.1%)	
P-value of Proportion Test					0.009	
			Panel B: Summary Statistics for Guests with Blank Review			
White			100	53	18	34.0%
African American			100	64	22	34.4%
Average Differences					-0.4%	
Difference 95% Confidence Interval					(-18.0%, 17.2%)	
P-value of Proportion Test					0.963	
			Panel C: Blank Review Experiment Regression Results			
	Linear Probability				*Logit*	
	(Review=0)	(Review=1)	(All)	(Review=0)	(Review=1)	(All)
	I	II	III	IV	V	VI
White	0.185^{***}	-0.020	0.192^{***}	0.843^{**}	-0.126	0.877^{**}
	(0.070)	(0.092)	(0.071)	(0.349)	(0.403)	(0.344)
Blank Review			0.127^{*}			0.720^{*}
			(0.077)			(0.373)
White × Blank Review			-0.196^{*}			-0.885^{*}
			(0.113)			(0.524)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	176	117	293	176	117	293

Note: $*p<0.10; **p<0.05; ***p<0.01$. Standard errors are robust. The characteristics in Panel C include all the host, listing and request characteristics in Table 3.

with African American–sounding names when guests state that they are friendly and tidy in the accommodation request. Last, Columns III and VI show that the self-claimed information about friendliness and tidiness does not reduce racial discrimination.

In summary, the self-claimed information about guests' tidiness and friendliness does not reduce discrimination. This empirical evidence rejects Hypothesis 4.

7.3. Blank Review Experiment

Last, we separate the effect of a review's existence from its content in reducing discrimination. In order to do this, we test whether a blank review can help reduce discrimination. Panel A of Table 7 shows that when there is no review, discrimination still exists in this fourth experiment (i.e., March and April 2018 in Los Angeles). Guests with white–sounding names are accepted 37.9 percent of the time, while guests with African American–sounding names are accepted 20.2 percent of the time. The difference is 17.7% percentage points (p-value $= 0.009$), which shows that discrimination still exists even with verified government ID (see Table 1 for all verification). Furthermore, Panel B shows that a blank review significantly reduces discrimination; the acceptance rates for white and

African American guests are 34.0 percent and 34.4 percent, and the difference is −0.4 percentage points (p-value = 0.963). This provides initial evidence that a blank review left by hosts can reduce discrimination. Following the same specification as in Table 3, we provide the regression results of this blank review experiment in Panel C of Table 7. Columns I and IV show that in the absence of self-claimed information and reviews, having a white–sounding name can causally increase the acceptance rate by 18.5 percentage points, controlling for host, listing and request characteristics. Columns II and V show that the acceptance rate of guests with white–sounding names is not significantly different from that of guests with African American–sounding names when there is a blank review on all guests' profiles. Last, Columns III and VI show that the blank review can help reduce discrimination, which is an interesting result given that the past literature (John et al. (2016)) has shown the failure to disclose information can cause observers to assume the worst. Moreover, the fact that estimated effect of blank review in Column III is smaller than the effect of the positive review in Table 3 is also consistent with this literature.

In summary, the existence of a review without content (i.e., a blank review) can help reduce discrimination; this provides empirical evidence to support Hypothesis 5.

8. General Discussion

Since discrimination has been documented in the sharing economy, there have been heated discussions on how to reduce it. One way to reduce discrimination is through legislation. Anti-discrimination laws have successfully reduced discrimination in housing and rental markets in the past few decades (US Department of Housing and Urban Development 2013). Several states such as California and Massachusetts forbid prospective landlords from asking applicants about their race, religion, gender or possible disabilities. However, the prevailing view among legal scholars is that online marketplaces like Airbnb fall into a gray area of the law (Todisco 2015); therefore, existing legal frameworks may not provide an adequate solution to discrimination issues on Airbnb (Belzer and Leong 2016). There also have been discussions about enforcing complete anonymization in online home-rental marketplaces. Such attempts, however, might jeopardize the core idea of a sharing economy—that is, building trust. Airbnb's chief executive, Brian Chesky, said that "access is built on trust, and trust is built on transparency."[17]

Airbnb has also been active in fighting discrimination. The platform has banned hosts due to the usage of discriminatory language.[18] On November 1, 2016, Airbnb issued a nondiscrimination policy and required its members to adhere to it. Under this policy, Airbnb may "take steps up to and including suspending the host from the Airbnb platform," if "the host improperly rejects

[17] "Discrimination by Airbnb Hosts Is Widespread, Report Says," *New York Times*, December 11, 2015.

[18] "Airbnb Bans N. Carolina Host as Accounts of Racism Rise," *USA Today*, June 1st, 2016.

guests on the basis of protected class."[19] However, it is not clear how the platform could monitor hosts' discriminatory behaviors and how such policies could affect discrimination on the platform.[20]

Our paper suggests reviews as an alternative approach to effectively reduce discrimination in the sharing economy. We find that in the absence of a review, an accommodation request made by a guest with an African American–sounding name is 19 percentage points less likely to be accepted by Airbnb hosts. However, a positive review can significantly reduce the observed racial discrimination based on a name's perceived racial origin. The findings are robust across various listings and host characteristics, including cities, listing types, hosts' past reviews, and hosts' gender and race.

We conduct further field experiments to explore how different a review's characteristics could affect its ability to reduce discrimination. Our non-positive review and blank review experiments demonstrate that a non-positive review or a blank review could also effectively reduce discrimination. This suggests that the existence of a review, rather than its sentiment or content, could help attenuate discrimination. Our self-claimed information experiment shows that the credibility of a review is essential to reduce discrimination; when the guests self-claim their friendliness and tidiness in accommodation requests, such unverified information cannot reduce discrimination.

8.1. Practical Implications

Our results shed light on several important managerial implications on how to leverage online reputation systems to prevent discrimination in the sharing economy. First, our positive review, non-positive review and blank review field experiments show that if a guest has one review (regardless of the sentiment), discrimination is significantly reduced. Therefore, we encourage platform owners to incentivize users to review one another after a transaction. Given that the review system is a part of Airbnb's platform design, such a recommendation is easy to implement and does not require major changes to the way that Airbnb is organized. There are several actions that a platform owner may take to encourage users to write reviews: for example, sending email reminders or offering monetary incentives such as discounts or credits for future transactions. These actions will be particularly effective in reducing discrimination if they target relatively new users who do not have a review yet.[21]

[19] "Airbnb's Nondiscrimination Policy: Our Commitment to Inclusion and Respect," Airbnb.

[20] For example, Airbnb can easily detect offensive language and ban the users. But it is difficult for Airbnb to react to weak signals like a rejection. Therefore, our suggestion to use reputation to prevent discrimination serves as a proactive approach rather than a punitive reaction.

[21] If the minority group receives poorer reviews on the platform, then having more reviews would not necessarily alleviate discrimination. In order to provide some insight that this is not the case on Airbnb, we randomly sample 5,000 guests who have transacted with hosts in our sample, collect the reviews these guests receive, and classify the review content as positive or negative. Table 10 presents summary statistics: African American guests in general have fewer reviews but the percentage of positive reviews is almost identical across African American and white guests. This provides some side evidence that African American guests are not more likely to receive poorer reviews.

Second, we also highlight that for first-time users, discrimination exists and may be substantial. The self-claimed information experiment shows that a mere description of someone's friendliness or tidiness is unlikely to reduce discrimination. First-time users should be cautious that discrimination may exist before the first review and they should also try hard to get the first review as soon as possible.

Third, we show that one blank review helps reduce discrimination while self-claimed information about tidiness and friendliness does not. This suggests that the credibility of a review—the fact that a review is verified by the platform and is linked to a completed transaction—is crucial for the review to reduce discrimination. The underlying mechanism may be (a) the completion of a transaction sends additional information about the guests' quality to reduce statistical discrimination or (b) the completion of a transaction signals that the guests, regardless of their race, have been accepted by other hosts and establishes a social norm for other potential hosts to accept these guests. Under either mechanism, our experiments show that in order to fight discrimination, platforms should ensure the credibility of reviews by monitoring the review system and only allowing either side of the platform to leave reviews after completing a transaction.

8.2. Limitations and Future Research

Our paper has several limitations that future research could help resolve. The first limitation, similar to Edelman et al. (2017), is that race signaled by names might be associated with socioeconomic status. Past research shows that African American–sounding names can be correlated with lower socioeconomic status (Fryer Jr and Levitt 2004). Thus, the level of identified discrimination could be driven by both perceived socioeconomic status and the race, signaled by names. Such limitation is embedded in the method of measuring discrimination in this type of correspondence study—using names to signal races. Future research could use other information to signal race, which may potentially disentangle the discrimination caused by race and socioeconomic status.

Moreover, the observation that online reviews reduce discrimination in our setting can potentially be explained by two mechanisms. The first is that reviews provide valuable information regarding guest quality, and as a result, hosts do not focus as much on race. The other is that reviews establish an inclusive normative behavior that make hosts more likely to accept guests who have been accepted by others, and either mechanism could explain all four of our experiments. Teasing out them is a subtle and tricky task because in order to establish a "credible" norm for inclusive behavior one has to show that a guest has been hosted by other members in the community regardless of her race. However, such endorsement by other hosts also signals guest quality such as safety, legitimacy, and familiarity with the process. As a result, we are not able to distinguish the two mechanisms completely in this specific context even with the non-positive review and blank review experiments. Moreover, it is possible there is no single mechanism that could fully explain the findings, and both mechanisms may be present. We therefore leave this interesting question for future research.

References

Aarts, Henk, Ap Dijksterhuis. 2003. The silence of the library: environment, situational norm, and social behavior. *Journal of personality and social psychology* **84**(1) 18.

Allon, Gad, Achal Bassamboo. 2011. Buying from the babbling retailer? the impact of availability information on customer behavior. *Management Science* **57**(4) 713–726.

Allon, Gad, Achal Bassamboo, Itai Gurvich. 2011. We will be right with you: Managing customer expectations with vague promises and cheap talk. *Operations research* **59**(6) 1382–1394.

Arrow, Kenneth J. 1973. The theory of discrimination. *Discrimination in Labor Markets* **3**(10) 3–33.

Ashlagi, Itai, Peng Shi. 2016. Optimal allocation without money: An engineering approach. *Management Science* **62**(4) 1078–1097.

Ayres, Ian, Peter Siegelman. 1995. Race and gender discrimination in bargaining for a new car. *American Economic Review* **85**(3) 304–321.

Bajari, Patrick, Ali Hortacsu. 2003. The winner's curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *The RAND Journal of Economics* **34**(2) 329–355.

Belzer, Aaron, Nancy Leong. 2016. The new public accommodations. *Georgetown Law Journal* Forthcoming.

Bertrand, Marianne, Esther Duflo. 2016. Field experiments on discrimination. Prepared for the Handbook of Field Experiments.

Bertrand, Marianne, Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *The American Economic Review* **94**(4) 991–1013.

Bimpikis, Kostas, Ozan Candogan, Daniela Saban. 2016. Spatial pricing in ride-sharing networks .

Bolton, Gary E, Elena Katok, Axel Ockenfels. 2004. How effective are electronic reputation mechanisms? an experimental investigation. *Management science* **50**(11) 1587–1602.

Buell, Ryan W, Tami Kim, Chia-Jung Tsay. 2016. Creating reciprocal value through operational transparency. *Management Science* Forthcoming.

Buell, Ryan W, Michael I Norton. 2011. The labor illusion: How operational transparency increases perceived value. *Management Science* **57**(9) 1564–1579.

Cachon, Gerard P, Kaitlin M Daniels, Ruben Lobel. 2015. The role of surge pricing on a service platform with self-scheduling capacity. *Available at SSRN* .

Carpusor, Adrian G., William E. Loges. 2006. Rental Discrimination and Ethnicity in Names. *Journal of Applied Social Psychology* **36**(4) 934–956.

Cheung, Cindy Man-Yee, Choon-Ling Sia, Kevin KY Kuan. 2012. Is this review believable? a study of factors affecting the credibility of online consumer reviews from an elm perspective. *Journal of the Association for Information Systems* **13**(8) 618.

Chevalier, Judith A., Dina Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* **43**(3) 345–354.

Cialdini, Robert B, Noah J Goldstein. 2004. Social influence: Compliance and conformity. *Annu. Rev. Psychol.* **55** 591–621.

Cohen, Maxime, Pavithra Harsha. 2013. Designing price incentives in a network with social interactions .

Cui, Ruomeng, Dennis J Zhang, Achal Bassamboo. 2017. Learning from inventory availability information: Field evidence from amazon. *Management Science* Forthcoming.

Edelman, Benjamin, Michael Luca, Dan Svirsky. 2017. Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics* **9**(2) 1–22.

Ewens, Michael, Bryan Tomlin, Liang Choon Wang. 2014. Statistical discrimination or prejudice? a large sample field experiment. *Review of Economics and Statistics* **96**(1) 119–134.

Feldman, Jacob, Dennis Zhang, Xiaofei Liu, Nannan Zhang. 2018. Taking assortment optimization from theory to practice: Evidence from large field experiments on alibaba .

Feldman, Pnina, Jun Li, Tiffany Tsai. 2017. Welfare implications of congestion pricing: Evidence from sfpark. Working paper, Boston University.

Foster, Andrew D, Mark R Rosenzweig. 1995. Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy* 1176–1209.

Fradkin, Andrey. 2016. Search frictions and the design of online marketplaces. Working Paper.

Fryer Jr, Roland G, Steven D Levitt. 2004. Understanding the black-white test score gap in the first two years of school. *The Review of Economics and Statistics* **86**(2) 447–464.

Gallino, Santiago, Antonio Moreno. 2014. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science* **60**(6) 1434–1451.

Ge, Yanbo, Christopher R. Knittel, Don MacKenzie, Stephen Zoepf. 2016. Racial and gender discrimination in transportation network companies. Working paper, University of Washington.

Gelman, Andrew, Hal Stern. 2006. The difference between “significant” and “not significant” is not itself statistically significant. *The American Statistician* **60**(4) 328–331.

Hanson, Andrew, Zackary Hawley. 2011. Do landlords discriminate in the rental housing market? evidence from an internet field experiment in us cities. *Journal of Urban Economics* **70**(2-3) 99–114.

John, Leslie K, Kate Barasz, Michael I Norton. 2016. Hiding personal information reveals the worst. *Proceedings of the National Academy of Sciences* **113**(4) 954–959.

Kaas, Leo, Christian Manger. 2012. Ethnic discrimination in Germany's labour market: A field experiment. *German Economic Review* **13**(1) 1–20.

Kabra, Ashish, Elena Belavina, Karan Girotra. 2015. Bike-share systems: Accessibility and availability. *Chicago Booth Research Paper* (15-04).

Li, Jun, Antonio Moreno, Dennis J Zhang. 2016. Pros vs Joes: Agent pricing behavior in the sharing economy. *Available at SSRN* .

Li, Jun, Serguei Netessine. 2018. Market thickness and matching (in)efficiency: Evidence from a quasi-experiment. Working paper, University of Michigan.

Luca, Michael. 2016. Reviews, Reputation, and Revenue: The Case of Yelp.com. Havard Business School Working Paper.

Milkman, Katherine L., Modupe Akinola, Dolly Chugh. 2012. Temporal distance and discrimination: An audit study in academia. *Psychological Science* **23**(7) 710–717.

Moreno, Antonio, Christian Terwiesch. 2014. Doing business with strangers: Reputation in online service marketplaces. *Information Systems Research* **25**(4) 865–886.

Myrdal, Gunnar. 1944. *An American Dilemma: The Negro Problem and Modern Democracy*. Harper, New York.

Nunley, John M., Adam Pugh, Nicholas Romero, R. Alan Seals. 2014. An examination of racial discrimination in the labor market for recent college graduates: Estimates from the field. Working Paper.

Pavlou, Paul A, David Gefen. 2004. Building effective online marketplaces with institution-based trust. *Information Systems Research* **15**(1) 37–59.

Schultz, P Wesley, Jessica M Nolan, Robert B Cialdini, Noah J Goldstein, Vladas Griskevicius. 2007. The constructive, destructive, and reconstructive power of social norms. *Psychological science* **18**(5) 429–434.

Taylor, Terry. 2016. On-demand service platforms. *Available at SSRN 2722308* .

Todisco, Michael. 2015. Share and share alike? considering racial discrimination in the nascent room-sharing economy. *Standford Law Review* **67** 121–129.

US Department of Housing and Urban Development. 2013. Housing discrimination against racial and ethnic minorities 2012. `https://www.huduser.gov/portal//Publications/pdf/HUD-514_HDS2012.pdf`.

Wason, Peter Cathcart, Philip N Johnson-Laird. 1972. *Psychology of reasoning: Structure and content*, vol. 86. Harvard University Press.

Wells, Gary L, Paul D Windschitl. 1999. Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin* **25**(9) 1115–1125.

Williams, David R, Selina A Mohammed. 2009. Discrimination and racial disparities in health: evidence and needed research. *Journal of behavioral medicine* **32**(1) 20–47.

Zervas, Georgios, Davide Proserpio, John Byers. 2015. A first look at online reputation on Airbnb, where every stay is above average. *Where Every Stay is Above Average (January 28, 2015)* .

Zervas, Georgios, Davide Proserpio, John Byers. 2016. The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry. *Boston U. School of Management Research Paper* (2013-16).

Zhang, Dennis J, Hengchen Dai, Lingxiu Dong, Fangfang Qi, Nannan Zhang, Xiaofei Liu, Zhongyi Liu, Jiang Yang. 2017. How do price promotions affect customer behavior on retailing platforms? evidence from a large randomized experiment on alibaba. Tech. rep., working paper.

Appendix

A. Discrimination Model

We study a model where a host receives an accommodation request and decides whether to accept it. Following the literature on economics of discrimination, a host makes a renting decision by considering aspects of guests' quality, such as tidiness. We assume that hosts are homogeneous in the market. A host receives utility u_i from an accommodation request i from a guest, which consists of guest quality and guest race,

$$u_i = \eta_i + \alpha r_i,$$

where η_i is the guest quality and r_i is the race of the guest in request i. We focus on two races in this paper (i.e., white and African American). Let $r_i = 1$ if the guest in request i is white, and $r_i = 0$ if the guest is African American.

Consistent with the majority of learning models found in the literature, we assume that the host, who possesses a prior belief about the quality of a potential guest based on race, uses the guest's personal information to update that belief and arrives at a final rental decision. Since the guest's quality is not observable to hosts, the host relies on a prior belief of the guest's quality η_i. We assume that the prior belief about guests with race r_i is drawn from a commonly known normal distribution, denoted as $\eta_i \in N(\bar{\eta}_{r_i}, \sigma_\eta^2)$, where $\bar{\eta}_0$ and $\bar{\eta}_1$ represent the prior belief of the average quality of white guests and African American guests respectively. The host then uses the guest's public profile information—e.g., a review by left by past hosts—to ultimately infer the guest's quality and make a rental decision. Specially, hosts receive a signal, $\tilde{\eta}_i$, of guest quality, η_i,

$$\tilde{\eta}_i = \eta_i + \epsilon_i,$$

where $\epsilon_i \in N(0, \sigma_s^2)$ represents the noise level of the signal.

After receiving a request i, the host updates his utility based on signal $\tilde{\eta}_i$ and the guest's race r_i. Let $\mathbb{E}[u_i|\tilde{\eta}, r_i]$ denote the expected utility that a risk-neutral host derives from request i conditional on the observed information. The host accepts the request i if and only if the expected utility derived from accepting the request is higher than the cost, c_i, i.e., $\mathbb{E}[u_i|\tilde{\eta}, r_i] > c_i$. The cost c_i includes both the physical cost and the opportunity cost of renting out the property, and it is drawn from a commonly known distribution with cumulative density function $F(\cdot)$.

We denote the renting decision by A_i, where $A_i = 1$ represents the condition in which the host has accepted the request i and $A_i = 0$ otherwise. The following proposition demonstrates the acceptance probability for request i.

Proposition 1. *For any request i, the probability that the request is accepted is $P(A_i = 1|\tilde{\eta}_i, r_i) = F((1 - \beta)\bar{\eta}_{r_i} + \beta\tilde{\eta}_i + \alpha r_i)$, where $\beta = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_s^2}$.*

Proposition 1 demonstrates that accepting a request is a weighted decision based on hosts' prior beliefs, received signal and guests' respective races. When hosts place a certain weight on race, i.e., $\alpha \neq 0$, race directly dictates the probability of accepting a request. When hosts place a non-zero weight on their prior beliefs $\bar{\eta}_{r_i}$, i.e., $\beta \neq 1$, race also affects hosts' decisions if the prior belief about group average varies by race, i.e., $\bar{\eta}_0 \neq \bar{\eta}_1$.

Proof of Proposition 1: The prior belief of guest quality in request i follows a normal distribution, i.e., $\eta_i \in N(\bar{\eta}_{r_i}, \sigma_\eta^2)$. The signal in request i is the linear combination of the true quality and normal noise: $\tilde{\eta}_i = \eta_i + \epsilon_i$ where $\epsilon_i \sim N(0, \sigma_s^2)$. In other words, $\tilde{\eta}_i \sim N(\eta_i, \sigma_s^2)$. Therefore, after receiving request i, the host observes a Gaussian signal $\tilde{\eta}_i$ as well as the group average $\bar{\eta}_{r_i}$ and uses this information to update the Gaussian prior. Based on the normal learning model (Foster and Rosenzweig 1995), the updated posterior also follows a Gaussian distribution,

$$\eta_i|\tilde{\eta}_i \sim N(\frac{\sigma_s^2\bar{\eta}_{r_i} + \sigma_\eta^2\tilde{\eta}_i}{\sigma_s^2 + \sigma_\eta^2},\ \frac{\sigma_s^2\sigma_\eta^2}{\sigma_s^2 + \sigma_\eta^2}).$$

The posterior distribution can be rewritten as

$$\eta_i|\tilde{\eta}_i \sim N((1-\beta)\bar{\eta}_{r_i} + \beta\tilde{\eta}_i,\ \frac{\sigma_s^2\sigma_\eta^2}{\sigma_s^2 + \sigma_\eta^2}),$$

where $\beta = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_s^2}$ represents the weight that a host puts on the observed signal to infer the guest's quality and $(1-\beta)$ represents the weight that the host puts on the prior group average.

Hence, the expected utility of accepting request i is

$$\begin{aligned} E[u_i|\tilde{\eta}_i, r_i] &= E[\eta_i + \alpha_i r_i|\tilde{\eta}_i, r_i] \\ &= E[\eta_i|\tilde{\eta}_i, r_i] + \alpha r_i \\ &= (1-\beta)\bar{\eta}_{r_i} + \beta\tilde{\eta}_i + \alpha r_i. \end{aligned}$$

Since the host only accepts the request if the expected utility is larger than the cost (i.e., $E[u_i|\tilde{\eta}_i, r_i] > c_i$) and the cost is distributed according to $F(\cdot)$, the probability of accepting request i becomes,

$$\begin{aligned} \mathbb{P}(E[u_i|\tilde{\eta}_i, r_i] > c_i) &= \mathbb{P}((1-\beta)\bar{\eta}_{r_i} + \beta\tilde{\eta}_i + \alpha r_i > c_i) \\ &= F((1-\beta)\bar{\eta}_{r_i} + \beta\tilde{\eta}_i + \alpha r_i). \end{aligned}$$

Hence, the acceptance rate of request i depends on the observed signal quality $\tilde{\eta}_i$, the average guest quality of race group $\bar{\eta}_{r_i}$ and the race of the guest r_i. ■

When request i and request j contain identical information, $\tilde{\eta}_i = \tilde{\eta}_j = \tilde{\eta}$, and the guests' races are different, $r_i \neq r_j$, discrimination exists if and only if

$$\mathbb{P}(A_i = 1|\tilde{\eta}, r_i) \neq \mathbb{P}(A_j = 1|\tilde{\eta}, r_j).$$

Based on Proposition 1, we define the discriminatory acceptance gap for signal $\tilde{\eta}$ as the difference between acceptance rates of two requests with identical signal $\tilde{\eta}$ but different races,

$$G(\tilde{\eta}) = F((1-\beta)\bar{\eta}_1 + \beta\tilde{\eta} + \alpha) - F((1-\beta)\bar{\eta}_0 + \beta\tilde{\eta}). \tag{3}$$

The discriminatory acceptance gap between white ($r_i = 1$) and African American guests ($r_i = 0$) stems from two mechanisms: taste-based discrimination and statistical discrimination. According to the definition, taste-based discrimination occurs when there is an inherent disutility related to certain race, i.e., $\alpha > 0$; statistical discrimination occurs when decision makers lack information and base their decisions on prior beliefs, i.e., $\alpha = 0$, $\sigma_s^2 > 0$ and $\bar{\eta}_0 \neq \bar{\eta}_1$. As the information signal becomes more informative (i.e., σ_s^2 decreases), hosts will place less weight on the prior belief (i.e., $1-\beta$ decreases). If discrimination is statistical, the discriminatory gap (i.e., $(1-\beta)\bar{\eta}_{r_i}$ will diminish. If discrimination is taste-based, the discriminatory gap will still exist.

B. Institutional Review Board

Since the study involves deceiving subjects, we establish to the board that deception is the only feasible means of conducting the research and it will pose no or minimal risks to subjects. In order to test whether hosts discriminate against guests based on race, the research randomly assigns guest races while keeping other characteristics the same. By concealing that these are in fact fictitious guest accounts, hosts will respond the same way as they would to a request from a real guest. Not concealing such fact will result in no response from hosts and hence a failure of this study. We take the following steps to minimize the potential risks to subjects: First, subjects involved in this study only have to read a short message with no more than 50 words and decide whether to respond; a typical response usually contains 5–10 words indicating whether the request is approved. Second, each selected host will only be contacted once through Airbnb's reservation request. Third, the request will be cancelled immediately upon receiving the reply such that the host will not have to block his calendar on the requested dates. We are granted a waiver of informed consent upon establishing that the study posts no or minimal risks to subjects and that requesting informed consent will violate the design of the study; hosts should not be aware of the study to ensure honest and accurate replies.

C. Auxiliary Figures and Tables

Figure 5 Airbnb Host Inquiry and Booking Page Example

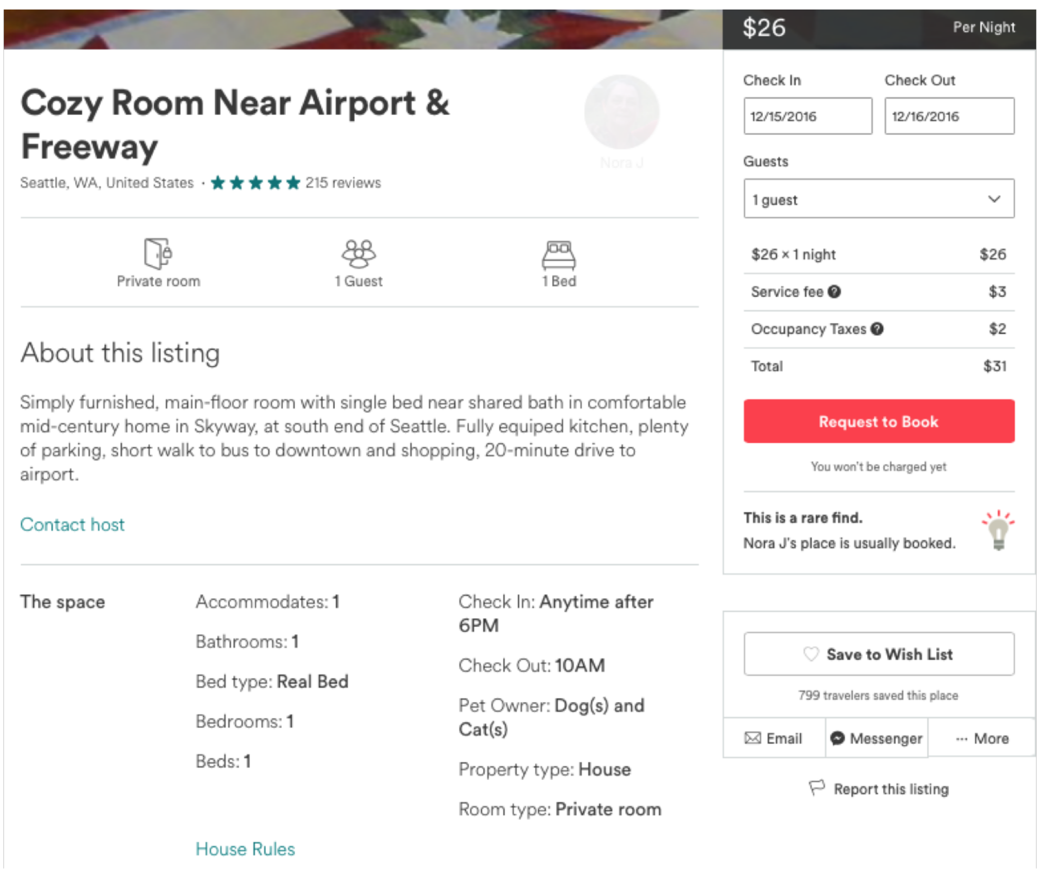

Figure 6 Airbnb Guest Page Example

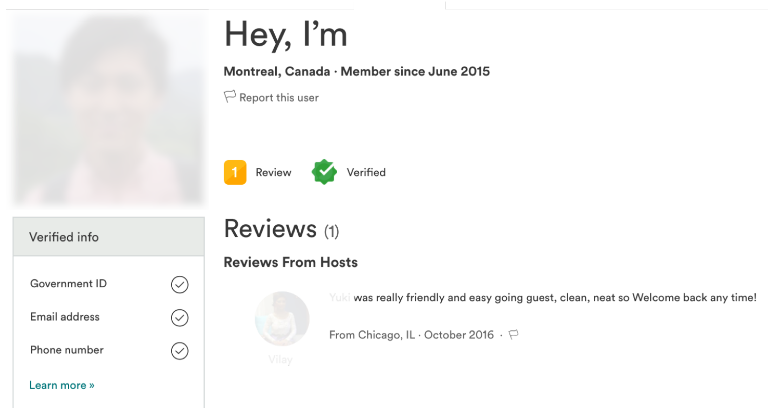

Figure 7 Guest Account Verification

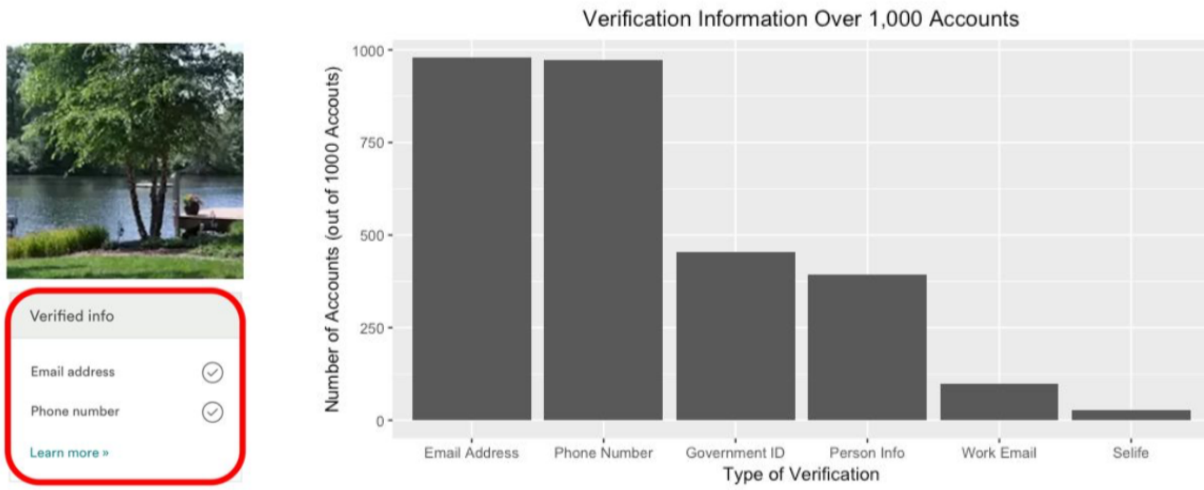

(a) Guest Account Verification Example

(b) Frequency of Different Types of Verification among 1,000 Randomly Chosen Guest Accounts

Table 8 Main Results from the Positive Review Experiment per City

	Dependent Variable: Acceptance		
	Chicago	Boston	Seattle
White	0.187^{***}	0.250^{**}	0.202^{*}
	(0.062)	(0.123)	(0.109)
Positive Review	0.185^{**}	0.239^{**}	0.301^{***}
	(0.081)	(0.118)	(0.110)
White × Positive Review	-0.244^{**}	-0.341^{**}	-0.265^{*}
	(0.124)	(0.171)	(0.150)
Host Characteristics	Yes	Yes	Yes
Listing Characteristics	Yes	Yes	Yes
Request Characteristics	Yes	Yes	Yes
Observations	318	135	145
Adjusted R^2	0.028	0.047	0.093

Note: *p<0.1; **p<0.05; ***p<0.01

Table 9 Amazon Mechanical Turk Survey Results on Names

Name	White	African American	Other	I cannot tell	No. of Participants	Intended Race
Colin Moore	78.85%	3.85%	3.85%	13.46%	52	White
Colin Murphy	84.62%	9.62%	1.92%	3.85%	52	White
Scott Mueller	88.46%	1.92%	0.00%	9.62%	52	White
Scott Baker	88.46%	1.92%	3.85%	5.77%	52	White
DeShawn Washington	1.92%	90.38%	5.77%	1.92%	52	African American
DeAndre McCray	7.69%	86.54%	0.00%	5.77%	52	African American
DeAndre Jackson	3.85%	92.31%	0.00%	3.85%	52	African American
Tyrone Washington	5.77%	88.46%	1.92%	3.85%	52	African American

Table 10 Reviews of Guests per Race

Race	Number of Guests	Average Number of Reviews	Average Number of Positive Reviews	Percentage of Positive Reviews
White	3752	6.20	6.16	99.64%
African American	109	5.17	5.16	98.97%
Difference		1.02	1.01	0.68%
P-value		0.1135	0.1184	0.8665

Note: Among 5,000 randomly selected guest accounts, 3,752 are White guests and 109 are African American guests.

Table 11 Summary Statistics of Host and Listing Characteristics per Experiment Conditions

Variable	Condition	No. of Obs	Mean	Std	P-value
Panel A: Positive Review Experiment					
No. of Bedrooms	White	293	1.32	0.81	0.8286
	African American	305	1.31	0.77	
No. of Reviews	White	293	28.03	36.66	0.2457
	African American	305	24.67	33.97	
Female Host	White	293	0.52	0.50	0.6759
	African American	305	0.50	0.50	
White Host	White	293	0.81	0.39	0.4986
	African American	305	0.79	0.41	
Panel B: Non-positive Review Experiment					
No. of Bedrooms	White	126	1.37	0.85	0.4603
	African American	124	1.28	0.86	
No. of Reviews	White	126	27.55	37.09	0.5487
	African American	124	30.49	40.34	
Female Host	White	126	0.56	0.50	0.3571
	African American	124	0.62	0.49	
White Host	White	126	0.77	0.42	0.5851
	African American	124	0.80	0.40	
Panel C: Self-Claimed Information Experiment					
No. of Bedrooms	White	346	1.14	0.79	0.021
	African American	314	1.01	0.59	
No. of Reviews	White	346	20.30	38.46	0.7153
	African American	314	19.26	34.92	
Female Host	White	346	0.56	0.49	0.8563
	African American	314	0.58	0.49	
White Host	White	346	0.76	0.43	0.7816
	African American	314	0.75	0.43	

Note: The number of bedrooms is significantly different between the white–name and African American–name conditions in the self-claimed information experiment due to large variations in house size in Austin. The difference is statistically insignificant if we restrict to listings with fewer than three bedrooms or listings in Seattle and Boston. Our main analyses are robust with and without controlling for the number of bedrooms.