Preliminary Draft: Do Not Cite or Distribute.

Market Design Choices, Racial Discrimination, and Micro-Entrepreneurship in Digital Marketplaces

Mohammed Alyakoob
Marshall School of Business, University of Southern California, Los Angeles, CA, 90089
alyakoob@marshall.usc.edu

Mohammad Rahman
Krannert School of Management, Purdue University, West Lafayette, IN, 47907
mrahman@purdue.edu

Digital marketplaces remove many traditional entry barriers, especially for minorities, and facilitate micro-entrepreneurship irrespective of the personal characteristics of entrepreneurs – race, origin, ethnicity, etc. The trust mechanism in such marketplaces, however, often require identifying information of the entrepreneurs (e.g., hosts on AirBnb, Taskers on TaskRabbit, etc.), accentuating concerns about discrimination against minority entrepreneurs. In this paper, we utilize a unique quasi-random design to examine discrimination faced by minority Airbnb hosts across 10 large U.S. cities. More importantly, we evaluate the potential for a platform-controlled market design choice to alleviate such discrimination. Specifically, exploiting the discontinuity in the criteria to be a Superhost – an Airbnb certification program, we study the effect of random information shocks, enjoyed by the borderline Superhosts, on the reservations, revenue, and booking windows of Airbnb micro-entrepreneurs (hosts). We find that an average Superhost enjoys a 10% growth in their reservations and 23% increase in corresponding revenue from the platform certification. Consistent with the prediction of statistical discrimination theories, borderline Black Superhosts benefit disproportionately from the information shock – a 20% increase in reservations and a 66% increase in corresponding revenue. We find consistent results when using double machine learning methods to investigate the heterogeneous impact of Superhost assignment for hosts not on the boderline of Superhost status. While the documented inequity faced by the black hosts on AirBnb is troubling, we are encouraged by the fact that platform certification programs can be effective in mitigating the discrimination faced by minority entrepreneurs on digital marketplaces.

1. Introduction

One of the most ubiquitous examples demonstrating the Internet's capacity to easily connect people is the success of digital marketplaces. Providing an efficient platform to connect demand and supply, marketplaces such as eBay, Amazon, Etsy, Airbnb, and TaskRabbit, among others, provide companies and individuals an avenue to easily transact with consumers. This facilitates the entry of small scale sellers and service providers who previously may have found market participation infeasible. For example, an individual seeking to sell a piece of furniture now has access to a vast set of consumers without having to open a store or engage in any market-building activity. The increased access attributable to digital marketplaces is even more pronounced when we consider the recent rise of sharing economy platforms. These platforms enable individuals to rent out (or share) their personal assets, such as their homes and cars. This provisions micro-entrepreneurship that is based on extracting rent from products that an individual already owns. Whether it is selling or sharing a service or product, digital marketplaces generally provide entrepreneurs with increased market accessibility.

Importantly, the increased accessibility enabled by digital marketplaces has the potential to equally benefit all market participants, even those from groups that have historically experienced disparities. For example, a persistent issue when examining entrepreneurship has been the under-representation of minority entrepreneurs. The disparity in entrepreneurial rates has been linked to factors such as a lack of credit access (Blanchflower et al. 2003, Freeland and Keister 2016), a lack of a self employed family member to provide insights (Fairlie and Robb 2007), and a persistent inter-generational wealth gap (Chetty et al. 2020). By providing easier entry for entrepreneurs, digital marketplaces can facilitate an environment that results in lower levels of disparity. This is most evident when considering sharing economy platforms, where the only requirement for entry is to make an already owned product (i.e., a home or a car) available for rent when the owner is not using the product.[1]

[1] We recognize that there are disparities in the initial access to these products. For example, whether an individual has access to an extra room, which they can monetize by renting through the home-sharing site Airbnb, is not independent of the traditional factors that lead to disparate outcomes. Nonetheless, it is clear that these platforms lower the bar considerably for the requirements to enter in many markets as an entrepreneur.

While digital marketplaces have decreased barriers to entry, their success hinges on fostering trust between both sides of the market (supply and demand). Initial skeptics of digital marketplaces suggested that the information uncertainty inherent in peer-to-peer interactions would limit the willingness of market participants to engage with other side participants. The success of early marketplace platforms, such as eBay, can be directly attributed to their implementation of feedback and reputation mechanisms, which helped establish trust between market participants (Tadelis 2016). Most design choices are based on user-generated assessment (i.e., reviews, ratings, and question and answer forums), while others, such as user certification programs, are determined by the platform itself. One persistent feature of most marketplace design options related to feedback and reputation systems is that they do not historically utilize the participant's identity (i.e., gender, race, etc.). This is useful when promoting the narrative that digital marketplaces reduce the potential for identity driven disparate market outcomes.

In recent years, digital marketplaces have incorporated design choices that make participants' identities more visible. This has mostly occurred on platforms where trust building requires personal identification. A prominent example is sharing economy platforms, which have been the fastest growing digital marketplaces over the past decade and have become increasingly mainstream.[2] Due to the intimately personal nature of transactions on these platforms—for example, sharing a car/home with an individual requires a substantially higher level of intimacy then purchasing a product on Amazon—enabling the participants to observe the name and image of other participants is usually deemed necessary by the platform. For example, Airbnb, Uber, and TaskRabbit allow participants to observe the names and images of other users. Unfortunately, revealing participants' identities creates an environment that can foster discrimination and, thereby, cause disparity in outcomes.[3] While revealing images might cause issues, anti-discrimination policies are difficult to

[2] For example, Uber and Airbnb, two of the most visible sharing economy digital platforms, have a combined market capitalization of over $200 billion as of March 2021.

[3] The demand-side of these platforms—requesters on TaskRabbit and visitors on Airbnb—are not currently required to show their images prior to the transaction. Previously, when identifiable information on the demand-side consumers were available, several studies have documented discrimination faced by minority market participants (Cui et al. 2020, Edelman et al. 2017, Pope et al. 2017). This policy is starkly different for supply-side participants. For example, workers on the labor matching platform TaskRabbit and hosts on the home-sharing platform Airbnb both provide their images to be viewed by other-side market participants.

enforce, especially with regard to potential discrimination facing supply-side participants. Platforms cannot observe the decision making processes of demand-side participants, unlike the more observable supply-side participants.[4] Therefore, platforms are generally left with market design solutions that disincentivize discriminatory behavior on the platforms.

Past research documenting racial disparity in outcomes of supply-side participants (i.e., entrepreneurs) is at best suggestive of discrimination (Laouénan and Rathelot 2020, Mohammed 2020, Edelman and Luca 2014).[5] This is because documenting discrimination faced by supply-side participants is notoriously challenging in digital marketplaces.[6] However, with the recent trend toward incorporating racially identifiable information of the entrepreneurs in digital marketplaces, it is evermore important to investigate the role of discrimination in these marketplaces. More importantly, as highlighted previously, platforms cannot actively police consumer prejudices and discriminatory purchasing behavior. Consequently, it is important to identify market design solutions that can minimize discrimination in such marketplaces.

This dilemma in documenting discrimination is further complicated because it can take various forms. Economics generally delineates discrimination into two types: taste-based and statistical-based discrimination. Taste-based discrimination (Arrow 1998) is based on a general dislike towards interactions with members of a specific group. Statistical discrimination (Arrow 1998, Aigner and Cain 1977) is based on perceived statistical group averages.[7]

To illustrate, assume we observe that digital marketplace participants of a specific race, in a marketplace where the participants' identifying features (image, name, etc.) can be observed, derive

[4] This is similar to the broader regulatory challenges in implementing anti-discrimination policies and laws. For example, based on the Civil Rights Act of 1964, while the U.S. federal government can punish sellers (i.e., supply-side participants) for discriminating against certain buyers, they are unable to punish buyers who choose not to purchase from certain sellers because of race, or other discriminatory reasons.

[5] Disparity in outcome can happen as a result of discrimination. However, it can also be a result of non-racially motivated differences in consumer preferences or simply differences in the inherent characteristics of the goods (Lang and Spitzer 2020).

[6] Two exceptions are Doleac and Stein (2013) and Jin and Kato (2006) where the authors document discrimination by experimentally varying the race of sellers in an e-commerce setting by changing the color of the hand holding the product.

[7] Taste-based and statistical discrimination can be difficult to disentangle, especially in cases where the discrimination is caused by invalid statistical inferences, meaning the perceived statistics are not accurate (Bohren et al. 2019).

lower demand levels. This disparity may be caused by various factors, including: (i) their products are inferior for an unobserved reason that is correlated with race; (ii) demand-side participants have an aversion to interactions with member of the race and engage in discrimination (taste-based discrimination); (iii) demand-side participants, based on their statistical interpretations, have lower expectations for the products offered by members of the race and, therefore, discriminate (statistical discrimination). To identify discrimination related to (ii), we would need to examine whether randomly removing the identity indicators would reduce the levels of disparity. Alternatively, since the discrimination described in (iii) is statistical in nature, we would need to randomly provide more information about participants and examine whether all participants experience homogeneous demand changes, regardless of race.

Pragmatically, platform design choices are more likely to solve statistical discrimination. Resolving taste-based discrimination would require removing identifiable information all together. Additionally, in labor market studies, research has shown that most discrimination is statistical rather than taste-based. Thus, if consumers are provided with sufficient information about individual entrepreneurs that dispels quality related concerns, racial disparity should subside. Therefore, we focus on certification as a design choice, which provides a reputation signal about entrepreneurs, as a potential anti-discriminatory force in digital marketplaces. Certification provides a clearer and more powerful signal than user-generated contents such as ratings and reviews.[8]

One method that researchers have frequently used to investigate discrimination is through audit (or correspondence) studies. Researchers create fictitious profiles with identical features except for the identity features (i.e., race, gender, etc.) and examine the differences in outcomes.[9] Similar approaches have been used to study potential discrimination faced by demand-side consumers on

[8] User generated contents are organic and, therefore, cannot be systematically relied on as a design intervention. Moreover, reviews often have biases and can lack credibility in the eyes of consumers (Tadelis 2016, Zervas et al. 2020).

[9] Previously, researchers had used multivariate regressions with a dummy variable indicating the race of the individual. Critics highlighted the difficulty in causally prescribing discrimination because of the difficulty in adequately controlling for all relevant observable variables. Refer to Lang and Spitzer (2020) for a recent survey of the discrimination literature in Economics.

sharing economy platforms. Audit studies on Airbnb (Cui et al. 2020, Edelman et al. 2017) and Uber (Ge et al. 2020) find that supply-side providers—hosts on Airbnb and drivers on Uber—discriminate against demand-side participants. To examine whether the discrimination faced by Airbnb guests is driven by statistical discrimination, Cui et al. (2020) investigate whether the disparity persists after randomly allocating ratings and reviews to the Airbnb demand-side fictitious accounts (the visitors). They find that providing hosts more information about the visitors does alleviate the disparity, which suggests that parts of the discrimination facing Airbnb hosts are associated with statistical discrimination.

Using audit studies to disentangle disparity and discrimination directed towards the supply-side participants is practically infeasible. Audit studies creating fictitious supply-side participants are extremely costly to implement, notwithstanding the legal and ethical ramifications of such efforts since they may adversely impact a platform's market equilibrium.[10] One method to overcome this challenge is to randomly enhance the reputation of a subset of supply-side participants. If the disparity on the digital marketplace is not driven by discrimination, all market participants whose reputation is randomly elevated should benefit equally from their newly enhanced reputations. Since the reputation boost is random, it cannot be correlated with other confounding factors. On the other hand, if market participants do *not* benefit equally, and the differences are based on the identities of the participants, this indicates that the initial disparity is associated with discrimination. Specifically, parts of the disparity are driven by discrimination that can potentially be alleviated by providing reputation enhancing signals.

Fortunately, platform designated certification programs can be used to construct an empirical design that provides a quasi-random reputation boost. Digital marketplaces often provide certification to users—usually supply-side users—who meet a specific quality standard. Platform certification is a design choice that can alleviate the aforementioned information asymmetry challenges by

[10] For example, Cui et al. (2020) create fictitious demand-side profiles on Airbnb, or potential visitors. They simply request and then immediately cancel the booking if accepted. The booking are for a few months away so there is not too much impact on the host. On the other hand, the analogous supply-side study would require the creation of multiple host accounts. Beyond the fact that it is unrealistic to purchase/rent properties by researchers, it may also impact the demand of other hosts.

elevating the reputation of higher quality users, facilitating trust between market participants.[11] Since the certification is determined by the platform, it can provide a clearer signal than user generated metrics such as ratings and/or reviews. By providing a clear reputation boost to entrepreneurs participating on sharing economy platforms, certification programs can potentially mitigate statistical discrimination and alleviate disparity in sharing economy supply-side outcomes. Importantly, unlike user-generated reputation mechanisms, platforms have significantly greater influence on the design and implementation of certification programs. Therefore, if certification programs do alleviate disparities, platforms have a more manageable design choice with which to combat discrimination on their platforms.

Specifically, we study the effect of a certification program, Superhost, used by the home-sharing marketplace Airbnb. Airbnb enables hosts (supply-side participants) to share their properties with visitors (demand-side participants) for short-term accommodations. Before booking, visitors can view information about the host (including name and picture) and listing (including pictures of the property, user reviews, bathrooms, bedrooms, and a plethora of other characteristics). Visitors are also shown whether the host has been certified as a Superhost by the platform. Airbnb states the following when describing Superhosts:

"Superhosts are experienced hosts who provide a shining example for other hosts, and extraordinary experiences for their guests."

Our empirical design exploits the clear discontinuity in the criteria for the Airbnb hosts to obtain Superhost status. Specifically, we identify the Airbnb hosts that are on the margins of fulfilling the Superhost criteria. The hosts at the margins of the criteria discontinuity, given that they are unable to precisely manipulate which side of the cutoff they eventually land on, provide a quasi-experimental setup to estimate the causal impact of Superhost assignment for Airbnb hosts. More importantly, this unique empirical design enables us to investigate whether the quasi-random

[11] Example of certification programs include: Airbnb assigns high quality hosts as "Superhosts", TaskRabbit assigns elite taskers as "TaskRabbit Elite", and Yelp assigns premier reviewers to their "Yelp Elite Squad". Certification programs have been documented to improve the performance of supply-side certification holders (Elfenbein et al. 2015).

certification impact is homogeneous for all hosts, regardless of race. This provides a setting where we can identify the extent of racial discrimination directed towards the entrepreneurs on sharing economy platforms and to study whether certification programs are an effective platform design choice to mitigate racial disparity on these platforms.

2. Related Literature

Our work is related to the broad economics literature investigating discrimination (Arrow 1998). While the challenges associated with studying discrimination limit the volume of such studies (Lang and Spitzer 2020), this void is acutely evident when considering discrimination facing supply-side digital platform participants. Much of the work investigating digital platforms has focused on documenting discrimination facing demand-side participants. As stated previously, the importance of such studies has accelerated since many of the rapidly growing platforms require participants to reveal their identities to facilitate trust. This practice has caused concern because it enables participants to engage in discriminatory behavior on the platforms. Edelman et al. (2017) and Cui et al. (2020), employing an audit study approach, find that Airbnb hosts are less likely to accept a visit request from a black individual. A similar approach has also been utilized to document disparity facing demand-side consumers on ride-sharing platforms (Ge et al. 2020). While similar discrimination most likely faces supply-side participants, documenting such discrimination using similar audit study based approaches is generally impractical. The dearth of studies investigating supply-side discrimination is amplified because, as stated previously, platforms are more limited in policing the behavior of demand-side participants. Therefore, an examination of both the potential discrimination targeting supply-side digital platform participants, as well as potential platform solutions is imperative to further this stream of literature.

Due to the mentioned difficulties in conducting supply-side randomized experiments, investigations of discrimination facing supply-side digital marketplace participants generally employ cross-sectional analysis to highlight the disparity in minority outcomes.[12] For example, Pope et al. (2017)

[12] Identifying disparity is relatively more straightforward than identifying discrimination since disparities can be caused by other factors such as differences in preferences or characteristics (Lang and Spitzer 2020).

and Younkin and Kuppuswamy (2018) find that funding requests by black participants are less likely to obtain funding on crowdfunding platforms. On the home-sharing platform Airbnb, Edelman and Luca (2014), Laouénan and Rathelot (2020), and Mohammed (2020) find that minority hosts generally post lower prices.[13] Laouénan and Rathelot (2020), using a conceptual framework, investigate the discrimination associated with the differences in posted prices, finding that it is mainly driven by statistical discrimination. They find no statistical racial difference in posted prices for Airbnb hosts with fifty or more reviews. In contrast, we go beyond factors dictated by supply-side participants (posted price) and investigate discrimination directly tied to consumer behavior (# of reservations, revenue, and booking window). This provides a more reflective representation of the discrimination, especially since Airbnb hosts' pricing strategies are often not optimized (Zhang et al. 2021). Moreover, we use a quasi-experimental approach to investigate the potential for discrimination on the platforms. This is unique given the aforementioned difficulties in conducting experimental investigations of discrimination facing supply-side participants.

Specifically, we investigate the effectiveness of reputation and feedback mechanisms to mitigate statistical discrimination in digital marketplaces.[14] Moreover, due to the various issues surrounding user-generated feedback we heed the call of recent scholars to go beyond user-generated feedback as a reputation mechanism(Tadelis 2016, Nosko and Tadelis 2015). Despite their prevalence, user-generated feedback mechanisms suffer from various biases. For example, user-generated feedback is shown to have an upward bias (Dellarocas and Wood 2008, Nosko and Tadelis 2015, Bolton et al. 2013), which is more pronounced in markets with reciprocal reviews, such as Airbnb (Zervas et al. 2020, Fradkin et al. 2015). In addition, recent studies have also outlined the potential for fraudulent reviews to contaminate the review generation process (Mayzlin et al. 2014, Luca and Zervas 2016). Moreover, user generated reviews may also suffer from the same racially motivated

[13] Mohammed (2020) also finds that the racial difference in posted prices does not significantly change when the platform removes the thumbnail of the host picture in the early stages of a visitor's search process.. Importantly, the visitor could still observe the picture, but only after selecting to view a specific listing.

[14] Recall, statistical discrimination can be reduced by providing more information to reduce the impact of perceived statistical averages.

biases previously outlined. Since consumers generally have a variety of alternatives, the issues hindering the effectiveness of user-generated reviews in facilitating trust are more pronounced for supply-side participants. As such, we contribute to recent efforts examining the potential of platform initiated sponsorship to facilitate trust (Horton 2019, Saeedi 2019) and, more importantly, we explore the potential for the clearer signal to alleviate discrimination on digital platforms.

Given the potential for sharing economy platforms to provide a more equal avenue to entrepreneurship, our work also contributes to the stream of entrepreneurship literature that examines the racial differences in entrepreneurial participation and success. Digital platforms are often touted for their capacity to democratize access, reducing the role of traditional obstacles to entrepreneurship (Blanchflower et al. 2003, Freeland and Keister 2016, Chetty et al. 2020). Notably, prior studies have examined the potential for digital marketplaces to directly mitigate traditional obstacles facing minority entrepreneurs. For example, studies find that, when the identity is revealed, funding rates on crowd-funding marketplaces are subject to racial disparity (Pope et al. 2017, Younkin and Kuppuswamy 2018), reducing their capacity to overcome traditional funding disparities. While we do not directly investigate digital marketplaces' capacity to directly alleviate traditional obstacles, digital marketplaces generally have lower barriers to entry, and they may provide greater equity in entrepreneurial opportunities, despite the continued existence of these traditional barriers. However, unfortunately, digital marketplaces design choices geared at facilitating trust between participants may enable platform level discrimination. Therefore, by providing a novel strategy to document discrimination as well as potential platform level solutions, we provide important insights regarding the possibility for digital platforms to provide greater equality of opportunities for micro-entrepreneurs.

3. Empirical Context

3.1. Airbnb Overview

Our empirical context focuses on Airbnb, a home-sharing digital marketplace. Airbnb enables micro-entrepreneurs, referred to as hosts, to list their properties, for rent, on the platform. It is one of the most visible and utilized digital marketplaces for micro-entrepreneurs engaging in the sharing economy, with about 3 million hosts as of 2021.[15] Given the intimate nature of the service—individuals sharing their homes with strangers—Airbnb hosts and visitors generally provide detailed profiles, which include an image, name, and a brief description. As previously prescribed, this has led to scrutiny about the potential for such platform design choices to facilitate discrimination. We study Airbnb activity in 10 major U.S. cities: Boston, Chicago, Dallas, Houston, Los Angeles, New York City, Oakland, Philadelphia, and Washington. The examined cities have high levels of Airbnb activity as well as diverse populations. This provides an ideal setting to investigate discrimination facing supply-side participants on a digital platform.

To assist potential visitors in selecting their desired listing and host combination, the platform displays listing characteristics such as the number of rooms, bedrooms, beds, bathrooms, as well as various pictures of the property. Listings can be shared, meaning the host and the visitor may share amenities during the visitor's stay, or private, meaning the host and visitor do not share the amenities. Figure 1 shows a sample listings page. As shown in Figure 1, the platform also displays information about the host, including the host name and photo, as well as whether the host has been designated a Superhost by the platform. As stated previously, the Airbnb Superhost certification is assigned to users that have a history of providing exemplary experiences to visitors. Moreover, in Section 4.1, we provide a detailed explanation of the Superhost assignment procedure as well as the clear discontinuities which provide the basis of our identification strategy.

3.2. Airbnb Data

We obtain detailed data on Airbnb listings from a third party vendor, AirDNA. Our sample period is from April 2016 - October 2018. For each listing, we observe all listing characteristics available to a potential visitor, such as the reviews, number of bedrooms, Superhost certification status, among others. We also obtain detailed daily reservation information for each listing. Importantly, for each reservation the platform facilitates, we obtain the date the visitor made the booking on the platform—referred to as the *Booking Day*—and the date the visitor will stay at the property—referred to as the *Reserved Day.* This distinction is important since we are concerned with what

[15] https://www.stratosjets.com/blog/airbnb-statistics/

the consumer was able to view on the website whilst making their booking. To clarify the naming conventions, assume a consumer visited the Airbnb website on July 20, 2017 and booked a stay for August 15, 2017. The *Booking Day* for this reservation is July 20, 2017 and the *Reserved Day* is August 15, 2017. We also observe, as of each *Booking Day*, all the *Available Days* for that listing. *Available Days* are specifically important in the home sharing context because hosts may block specific days for reasons unrelated to the platform. For example, a host may choose to make their property available on weekdays but choose to personally use the property on weekends.

For each host, we also acquire the profile picture that the host uses on their Airbnb host page. These profile pictures allow a viewer to perceive the race of the host. Accordingly, we use these images to determine the host race, as outlined in detail in Section 3.3. Moreover, for each listing, we procure all the property images that are made viewable on the platform. Naturally, these images, particularly, the first image, make the first impression. Consequently, for each property image, we utilize recent advances in computer vision algorithms to classify the aesthetic appeal of the image (Lennan et al. 2018). This process provides a score for the aesthetic appeal of each image and allows us to measure the first impression from the property images in a systematic and scalable fashion.

3.3. Determining Host Race

As mentioned, we obtain the profile pictures of all hosts in our sample. For each profile picture, we utilize Amazon Mechanical Turk (MT), a crowdsourcing marketplace, to determine the race of the host in each profile picture. While MT workers may not be able to identify the actual race of the host, they are ideal for our objective since we are concerned with the perceptions of race from the image. An alternative approach that has been used is to utilize artificial intelligence (AI) based methods for predicting race from images. These methods have been scrutinized due to potential biases in labeling images of minorities. Moreover, these approaches are optimized to predict the race of the individual and not necessarily the *perception* of race. For example, these algorithms may identify certain races by a person's bone structure, which may not align with how individuals perceive a person's race. Therefore, whenever logistically and financially viable, we prefer to use an approach that is wholly dependent on human participants' perceptions of race for each profile picture. Specifically, each image is viewed by multiple individuals and a consensus based approach is used to determine race.

Each Airbnb host profile picture was initially assigned to three MT workers. Each MT worker was asked to label the individual(s) in the picture from the following criteria: no faces, more than one face, black, white, east asian (e.g., China, Japan ...), south asian (e.g., India, Pakistan ...), latino/hispanic, unknown, combination of two or more races. If all three MT workers agreed on the race classification, the host was labeled based upon the unanimous classification indicated by the workers. If the three workers did not agree, we assigned an additional two workers to classify the host's profile picture, increasing the number of workers to five for that specific picture. If four out of the five workers agree on a classification, we assign the host the race with the majority classification. In cases where less than four MT workers agreed on a race, we assign a further five MT workers, taking the total number of workers assigned to classify the host's profile picture to 10. If seven out of ten workers agree, we assign the host's race to the majority classification. Otherwise, we assign the host race as unknown since the workers did not reach a consensus on the classification for the host's picture.

If the MT workers classify a host as "no faces", the host does not have a profile picture that shows a face. These could be pictures of rooms, animals, etc. In cases where the MT workers classify the profile picture as "combination of two or more races", we conduct a secondary procedure and request the workers to identify whether all the races in the picture are the same. If the MT workers agree that this is the case, we label the host accordingly. Otherwise, we assign the host race as unknown.

4. Identification Strategy

Our objective in this study is to investigate statistical discrimination facing hosts on Airbnb. As previously outlined, statistical discrimination, in our context, refers to visitors discriminating against hosts because of group level perceptions of quality. If we observe that host performance, after controlling for all observable characteristics, differs based upon the race of the host, then this could suggest discrimination. However, the differences in host performance may also reflect unobserved factors that are correlated with race. If the discrimination is statistical (Arrow 1998), then disentangling the two potential causes (discrimination or unobserved variables) requires providing the visitors with more information about the quality of the hosts and observing the impact on the outcomes. In cases where more information about quality reduces the differences, this indicates that a portion of the disparity is caused by statistical discrimination. Moreover, the hosts that were adversely impacted by visitors' preconceived notions, prior to the additional information availability, should experience greater gains from the new information provided to the visitors.

Fortunately, the Airbnb Superhost program provides an ideal setup to quasi-randomly provide additional information about a subset of hosts. To use the certification program to identify discrimination, two important elements are required. First, the platform certification must be credible such that the visitors accept the signal as a legitimate measure of quality. That is, the visitors must accept the validity of the signal to an extent that mitigates their preconceived prejudices regarding the host's group characteristics. Second, the rationale for assigning certification to a specific host must be independent of the endogenous unobserved factors that may correlate with the host receiving the certification. Airbnb's Superhost certification program provides a quasi-experimental setting that satisfies these conditions. First, recent market design literature has outlined the need for platform initiated signals, as they can overcome the biases related to more common methods such as user generated content (Tadelis 2016). Moreover, in Section 5 we highlight the clear performance gains for hosts that obtain certification. Second, the Superhost certification process provides clear discontinuities in criteria, enabling us to exploit a quasi-random Superhost assignment process, independent of potential unobserved confounders. We outline this process below.

4.1. Superhost Certification Program

The Superhost certification is assigned by Airbnb to hosts that achieve certain criteria. During our sample period, Airbnb hosts that met the following criteria over the previous year were assigned

Superhost status: (i) hosted at least 10 trips; (ii) maintained a 90% response rate when responding to visitor requests; (iii) received a 5-star review at least 80% of the time; (iv) completed all confirmed reservations without cancellation.[16] Notably, Airbnb uses specific "assessment periods" to evaluate whether hosts have met the Superhost certification criteria. For each year, the evaluation periods begin on January 1st, April 1st, July 1st, and October 1st respectively. Each Host's Superhost status is re-evaluated at the beginning of the subsequent assessment period, usually within 10 days of the start of the assessment period (i.e., for a January 1st evaluation period, each host's status is adjusted, if necessary, by January 10th). Superhost status may adjust every three months and our detailed data enables us to determine whether the host was a Superhost when the reservation was made (*Booking Day*).

We construct our data set around the 8 Superhost evaluation periods in our study period.[17] We examine the *Booking Days* in the two months after the evaluation month for all potential *Reserved Days* in the 6 months after each evaluation month. For example, April 2016 is a Superhost evaluation month, which means that Airbnb hosts will be evaluated on the aforementioned Superhost criteria, and will either obtain, retain, or lose their Superhost certification status. The subsequent certification month is July 2016. Therefore, all hosts that obtain the status in April 2016 are guaranteed to keep the certification until the end of June 2016. As such, Superhost certification status will not change during May and June of 2016. For each listing, we identify the *Booking Days* that occurred in May and June of 2016 and that corresponded to *Reserved Days* on days between May 1, 2016 and October 31, 2016 (a six month period after the evaluation month, April 2016). The fact that hosts must continually meet the criteria every three months is useful to alleviate concerns that hosts may reduce their effort levels once they obtain certification.

For each evaluation period, we calculate the *Reservations*, *Revenue* (total dollar amount earned), and *Booking Window* (the average number of days between the *Booking Days* and *Reserved Days* for a listing). We use these variables as the listing performance metrics in our analysis. We also

[16] Note, the criteria was slightly changed in subsequent periods, we investigate this in the Online Appendix.

[17] April (2016, 2017, 2018), July(2016, 2017), October(2016, 2017), and January(2017).

calculate the average listed price for each listing. This is calculated by identifying all the *Available Days* in the 6 month period and obtaining the average of all the listed prices for those days.

We focus our discontinuity identification strategy on the Superhost requirement that 80% of the host's reviews in the past year are five star reviews. For each host, we calculate the number of additional five star reviews the host would need to reach the 80% threshold. We name this variable *Additional 5-Star Reviews Needed*. For example, if a host has received eight total reviews over the past year, and six are five star reviews, that host has 75% ($\frac{6}{8} = 0.75$) five star reviews. A host with eight total reviews would need at least seven five star reviews ($\frac{7}{8} = 0.875$) to meet the 80% minimum threshold and obtain Superhost status. Therefore, the value of *Additional 5-Star Reviews Needed* would be -1 for this host (recall, the host had six five star reviews).[18] This host would not be assigned Superhost status. Hosts, assuming all the other Superhost criteria are met, would need *Additional 5-Star Reviews Needed* to have a value of 0, or higher, to obtain Superhost status.

We utilize the *Additional 5-Star Reviews Needed* because it provides a framework that meets the requirements of a discontinuity specification. Specifically, our setup would be invalid if hosts could *precisely* manipulate *Additional 5-Star Reviews Needed* around the threshold of 0. While Airbnb hosts are able to impact the level of quality they provide, which can ensure average higher ratings, and higher overall values of *Additional 5-Star Reviews Needed*, they cannot dictate whether they receive exactly enough reviews to pass the 80% threshold, or one less review.[19] This inability to precisely manipulate *Additional 5-Star Reviews Needed* around the 0 discontinuity point is exacerbated by the fact that reviewers may give a high quality host a four out of five rating. Therefore, we focus our analysis on the subset of hosts where *Additional 5-Star Reviews Needed* equals 1 (referred to as *Marginal Superhosts*) or 0 (*Marginal Not Superhosts*) for a specific evaluation period. These hosts have met all the Superhost criteria except that the members of the *Marginal Not Superhosts*

[18] To further illustrate, assume a host had eight total reviews in the previous year. The host would need seven five star reviews to achieve Superhost status. If the host had 4, 5, 6, 7, or 8 five star reviews in the previous year, *Additional 5-Star Reviews Needed* would be -3, -2, -1, 0, or 1 for each respective case.

[19] Lee and Lemieux (2010) provide a detailed discussion on this condition. Specifically, they state: "If individuals—even while have some influence—are unable to *precisely* manipulate the assignment variable, a *consequence* of this is that the variation in treatment near the threshold is randomized as though from a randomized experiment.

group are one five star review away from Superhost status. Importantly, while the *Marginal Superhosts* group do obtain Superhost status, they only differ from the *Marginal Not Superhosts* by one non five star review. Table 1 provides summary statistics for our main sample, which includes hosts in either the *Marginal Superhosts* or *Marginal Not Superhosts* groups.

4.2. Summary Statistics

Table 1 provides summary statistics (means) for the variables used in our analysis. Note that each observation corresponds to a unique combination of listing and evaluation period. *Reservations*, *Revenue*, and *Booking Window* are the three main listing performance metrics (dependent variables) that we assess. *Booking Window* refers to the number of days between the booking day and the reserved day. A larger values of the *Booking Window* represents better performance because it signifies the demand for the property well in advance. It is also an outcome variable that can capture the Airbnb guests' willingness to transact with a host from a particular race or with certain characteristics when they there are not under time pressure or supply constraints.

Instant Bookable (1/0) refers to an optional feature which the host can use to enable potential visitors to book their property without the host's approval. This essentially eliminates the screening possibilities by the host. *Shared Listing (1/0)* refers to whether the visitor will be sharing the lodging with the host during their stay. *Zipcode Revenue Share* is the share of all city level Airbnb revenue during a specific period that is attributed to the host's zipcode. This variable accounts for the relative popularity of a property's location temporally. *Image Aesthetic Score* refers to the aesthetic appeal score of the listing's main image. It ranges from 1-10, where a higher number signifies better aesthetic appeal. Recall that we utilize the deep learning approach outlined in Lennan et al. (2018) to obtain this measure.

Columns 1-3 of Table 1 show the summary statistics for all observations, segregated by Superhost status in Columns 2 and 3. The data indicates that on average, as expected, Superhosts outperform Non-Superhosts across the three performance metrics. Specifically, listings managed by Superhosts have more reservations and revenue, as well as longer booking windows. Superhosts have more reviews and higher ratings, while non-Superhosts are more likely to utilize the Instant Bookable option. Superhosts are also more likely to be located in zipcodes that earn a greater share of the overall Airbnb revenue.

While the full sample statistics are important to understand the clear differences between Superhosts and non-Superhosts, our methodology, as described above, focuses on hosts that are on the margins of attaining Superhost status. Columns 4-6 of Table 1 provide statistics for the subset of hosts that are on the margins of fulfilling Superhost criteria. Reassuringly, the gap between Superhosts and non Superhosts, specifically for the non-performance based variables, is significantly reduced in the marginal sample.

Our objective in this study is to examine whether the Superhost assignment has a different impact for Black hosts. If it does, this is indicative of statistical discrimination. Table 2 shows the statistics for the marginal subsample for Black and White hosts respectively. Comparing columns 1 and 4, which show the statistics for all Black and all White hosts, the data shows clear disparity between the two group of hosts. Black hosts, on average, earn less revenue and have smaller booking windows. Interestingly, Black hosts obtain, on average, slightly more reservations, however the lower prices entail that this doesn't necessarily translate to greater revenue. Notably, there is not a large difference between the average rating of the two group of hosts, indicating that the quality of the offering may not be starkly different. The lower revenue and booking windows are observed despite the fact that Black hosts are also much more likely to provide potential visitors the instant booking option. Notably, Black hosts are generally located in zipcodes that obtain lower revenue in general, as observed by the difference in *Zipcode Revenue Share*. Columns 2, 3, 4 and 5 of Table 2 show the differences between Black and White Superhosts. The statistics indicate that Black Superhosts have considerably higher reservations and revenue than Black non-Superhosts. While White Superhosts earn higher revenue and enjoy longer booking windows than White non-Superhosts, the gains are less pronounced than for the analogous difference between Black Superhosts and non-Superhosts.

To further examine the relationship between Black and White Superhosts, Tables 3 and 4 provide analogous statistics, keeping only those listings that are in zipcodes where more than 50% of the residents are Black (Table 3) or White (Table 4). These subsample allow us to compare the Superhost differences in areas that are potentially more similar. This is evident by the similarities in *Zipcode Revenue Share* observed in Tables 3 and 4 as compared to Table 2. Both subsamples indicate that Superhosts perform significantly better than non-Superhosts and, importantly, these differences are much larger for Black hosts than White hosts. This provides further evidence that the trends observed are not simply a result of Black hosts residing in different areas. Even in the majority Black and majority White zipcodes, the trends are consistent.

4.3. Examining Quasi-Random Assignment

Our approach depends on the assumption that the *Marginal Not Superhosts* and the *Marginal Superhosts* groups are similar. Crucially, these similarities imply that the *Marginal Superhosts* are unable to precisely manipulate their reviews and move from *Additional 5-Star Reviews Needed* equals 0 to 1. Figure 2 shows plots the volume of observations in our sample at each value of *Additional 5-Star Reviews Needed*. The evidence does not indicate that there is a clear jump from -1 to 0. Moreover, if *Marginal Superhosts* and *Marginal Not Superhosts* are statistically equivalent, there should not be a statistical difference between the groups based on any predetermined covariates (i.e., bedrooms, bathrooms, etc.). Table 5 presents statistics which examine the differences in the previous period posted price, previous period number of reviews, number of bedrooms, number of bathrooms, and previous period zipcode revenue share between the hosts where *Additional 5-Star Reviews Needed* equals -1 and 0. The results indicate that, in the periods preceding the Superhost assignment, there were not clear differences between the groups (there is a statistical difference in number of bathrooms but this is due to the lack of variation in number of bathrooms offered). Table 6 shows the analogous results for hosts with *Additional 5-Star Reviews Needed* equals -3 and 2 respectively. In contrast to Table 5, we find statistical differences between the group averages. This is reassuring because it entails that, while hosts can effect their overall quality, they are unable to to do so at the cutoff point, which is the sample we are concerned with.

5. Main Results

As previously outlined, our main identification strategy exploits the clear discontinuities in Airbnb Superhost criteria. Moreover, by focusing on listings associated with either the *Marginal Superhosts* or the *Marginal Not Superhosts* samples, we are examining a subset of hosts who could not have plausibly manipulated whether they obtain Superhost status. In addition, we focus on hosts that were not Superhosts in the previous period to further enforce that the two sets are comparable.[20] This indicates that when examining the results of this subsample, the Superhost assignment is quasi-random. We employ the following equation:

$$Y_{i,t} = \alpha_k + \delta_t + \beta_1 \cdot \mathit{Superhost}_{i,t} + \beta_2 \cdot \mathit{Black\ Host\ (1/0)}_i + \beta_3 \cdot \mathit{Superhost}_{i,t} * \mathit{Black\ Host\ (1/0)}_i + X_{i,t} + \epsilon_{i,t} \quad (1)$$

$Y_{i,t}$ represents the outcome variable (log(reservations), log(revenue), and log(Booking Window) for Airbnb listing i, which is located in city k, during evaluation period t. Our main variable of interest is *Superhost*, which is a binary variable indicating whether the listing's host became a Superhost in the evaluation period t. $X_{i,t}$ represents the collection of listing specific variables. These include the listing number of reviews, rating, instant bookable status, whether the listing is shared, number of bedrooms, number of bathrooms, Zipcode revenue share. We also include the log of Available Days and the average listed price for the Available Days as well as Image Aesthetic Score (described in Section 3.2).

Table 7 shows the results of Equation 1 excluding the interaction of host race and Superhost status. Columns 1, 2, and 3 show the results for log(Reservations), log(Revenue), and log(Booking Window) respectively.[21] The results indicate that, on average, Superhost status increases the number of reservations by 9.96% $((e^{0.095} - 1) * 100)$, revenue by 23.12%, and the booking window by 8%. While unsurprising, these results confirm the importance of platform sponsorship on the decisions that platform demand side participants make. Specifically, our results indicate that Airbnb

[20] This additional restriction leaves us with 15,362 marginal Superhosts.

[21] Note that column 3, log(Booking Window), has less observations since we cannot calculate this variable for any listing/evaluation period combinations where the listing had no bookings.

hosts that marginally obtain Superhost status significantly outperform comparable hosts who only missed obtaining Superhost status by one five-star review. The coefficients associated with the various controls are also reassuringly consistent. Listings with more reviews and higher average ratings perform better. Moreover, listings in zipcodes with a greater share of revenue as well as those with a higher aesthetic score for the listing main image perform better as well.

While the Superhost effect documented in Table 7 is reassuring, our main objective is to examine the heterogeneity of this impact as it relates to the host's race. Tables 8, 9, and 10 investigate this issue by estimating Equation 1 for log(Reservations), log(Revenue), and log(Booking Window), respectively. The results examining the heterogeneity related to reservations, Table 8, indicate that Black hosts who obtain the Superhost status increase their reservations by a further 10.1% as compared to non-Black hosts (Column 1). Notably, the results indicate that Black hosts generally perform significantly worse as compared to other hosts, with the coeficient associated with *Black Host (1/0)* indicating that Black hosts have approximately 16% less reservations than other hosts. The total boost Black hosts gain from the Superhost status, compared to all other hosts, is approximately 20.08% ($(e^{0.087+0.096}-1)*100$). Column 2 shows the analogous results for the subsample of hosts that only includes Black and White hosts, meaning that the base case is White hosts. The results are consistent with those in column 1, where the base case considers all non-Black all hosts. Columns 3 and 4 replicate the analysis in columns 1 and 2, except we replace Black hosts with No Face Hosts. Recall that No Face Hosts are those whose profile picture did not provide a portrait of the host. We have total 1,668 observations without profile pictures in this subsample. The results indicate that the effect size for these hosts is significantly smaller than for Black hosts and, notably, is not statistically significant.

Tables 9 and 10 provide the results with *Revenue* and *Booking Window* as the performance metric. The results provide a consistent message. Black hosts average performance is worse than other hosts and Black hosts have substantially higher increases from the Superhost status attainment[22]

[22] The coefficient for the interaction of Superhost and Black Host (1/0) in column 1 of Table 10 has a p-value of approximately 0.13.

These results are consistent regardless of whether the base case is all other hosts (Column 1) or only White hosts (Column 2). Moreover, the results examining No Face Hosts are also consistent, indicating that their benefit from the Superhost status is not statistically different than other hosts.

The results provide significant evidence that, prior to the platform provided information attributed to the Superhost assignment, Black hosts perform significantly worse than their counterparts. On the flip side, the results indicate that Black hosts gain the most from the Superhost assignment, documenting evidence of statistical discrimination on the platform while demonstrating the potential for platform sponsorship to provide necessary information to demand side participants to mitigate discrimination faced by Black hosts. Moreover, the findings collectively provide evidence that hosts that do not have a personal photo as the profile picture do not benefit differently from the others with the Superhost status. Also, a Black Airbnb host is generally better off not having a profile picture than disclosing their racial identity, especially when they are not a Superhost. From the perspective of platform certification as a design choice, while other hosts may also benefit from the additional information, they are likely not as negatively impacted by the lack of information as Black hosts. Overall, our findings indicates that Black hosts are suffering statistical discrimination and that certification programs can be an effective measure to mitigate discrimination facing minority entrepreneurs in digital marketplaces.

6. Double Machine Learning

Thus far, our analysis has focused on the hosts that are on the margins of obtaining Superhost criteria. By exploiting the discontinuity at the margins, we are able to identify the causal gain associated with Superhost status and, more importantly, the heterogeneity of this impact as it relates to host race for the hosts on the margins. Unfortunately, a limitation of the discontinuity design is that it is only able to provide insights for hosts on the margin of attaining Superhost status. It cannot assess whether these results are generalizable to the full set of hosts in our sample. Fortunately, recent advances have provided methods to identify heterogeneous treatment effects in non-randomized observational settings. To illustrate this approach as it relates to our context, consider the following specification:

$$Y_{i,t} = \theta(R) \cdot Superhost_{i,t} + g(R, X) + \epsilon_{i,t} \tag{2}$$

$$Superhost_{i,t} = f(R, X) + \nu_{i,t} \tag{3}$$

where $Y_{i,t}$ refers to the dependent variables considered in our specifciations (*log(Reservations*, *log(Revenue)*, and *log(Booking Window)*).[23] R refers to the race of the host and X refers to a set of observable characteristics of the sample. $g(R, X)$ is an unknown function of high-dimensional vector of observables R and X. $\theta(R)$ is the Superhost treatment effect, which is a function of the race of the host. Moreover, the variation in Superhost status is generated by $f(R, X)$, another unknown function of high-dimensional vector of observables R and X (Equation 3). In our analysis of the marginal set (Section 5 and equation 1), $g(R, X)$ is assumed to take a linear form. More importantly, the treatment assignment was random by design among the marginal Superhosts. However, the treatment assignment is not random in the general sample. Rather, we need to model the treatment assignment with selection on observables for estimating impact of the Superhost status in the overall sample.

One recent method proposed to alleviate this concern is by utilizing the double machine learning approach (Chernozhukov et al. 2017, 2018), which utilize state-of-the-art machine learning methods to non-parametrically estimate $g(R, X)$ and $f(R, X)$. Specifically, we can combine equations 2 and 3 to obtain the following:

$$Y_{i,t} - E[Y_{i,t}|R, X] = \theta(R) \cdot (Superhost_{i,t} - E[Superhost_{i,t}|R, X]) + \epsilon_{i,t} \tag{4}$$

Machine learning methods are used to obtain the best estimates of both conditional expectation functions from Equation 4. Subsequently, $E[Y_{i,t}|R, X]$ and $E[Superhost_{i,t}|R, X]$ are partialled out, leaving the following:

$$\hat{Y}_{i,t} = \theta(R) \cdot \widehat{Superhost}_{i,t} + \epsilon_{i,t} \tag{5}$$

[23] We further assume that $E[\epsilon_{i,t}|R, X] = 0$, $E[\nu i, t|R, X] = 0$, and $E[\epsilon_{i,t} \cdot \nu i, t|R, X] = 0$.

Finally, Equation 5 is estimated using OLS.[24].

There are numerous benefits to this approach. First, as previously mentioned, it allows us to utilize our full data set and examine whether our results are constrained to those hosts at the margins.[25] Second, it enables us to incorporate the richness of our data by including various features. These include a high dimensional set of interactions as well as features utilizing the details in the reviews as well as the listing descriptions. These are described in more detail in Section 6.1. Finally, while this approach allows us to easily estimate the heterogeneity related to race, it also enables us to conduct more thorough heterogeneity analysis. Specifically, in Section 6.3 we investigate the potentially mitigating role of reviews on the differences in benefit derived from Superhost status by host race.

6.1. Machine Learning Features

As mentioned, the benefit of the machine learning approach is that it enables us to leverage the richness of our data. First, using the listing description, we produce n-grams, setting n using a sliding window of 1 to 3 words over all the words in the listing descriptions. Subsequently, we include the frequency of each n-gram for each listing description as a feature. Second, we obtain the text for the previous 5 reviews written for each listing as of the observation period and utilize Latent Dirichlet Allocation (LDA) topic modelling. We include the results of this process as an additional feature.[26] Third, we include the *Image Aesthetic Score* for the first five images of each listing. We also incorporate a city by period fixed effects for each unique combination of city and period. This captures any potential city level temporal variation.

6.2. Heterogeneity of Race

To estimate the treatment effect of Superhost status, we evaluated the performance of various ensemble machine learning methods and found that random forest performed optimally for both

[24] For valid inference, half of the data is used to fit the model and the other half is used to estimate the residuals

[25] As previously described, another advantage of focusing on the marginal hosts is that the we are able to use Mechanical Turk to obtain host race predictions. In this section, we must obtain the host race of all hosts, not just those on the margins. Since this scale of prediction is not financially feasible using Mechanical Turk, we use machine learning to predict the race of each host. Importantly, an advantage we have is that we use the Mechanical Turk prediction as the ground truth. This procedure is outline in the Online Appendix.

[26] We provide more detailed explanations of this process in the online appendix.

predicting the treatment (Superhost status) and the outcome variables.[27] Table 11 and 12 report the results of the this procedure for various subsamples. The Tables report the average treatment effect (ATE) as well as heterogeneous treatment effects for different groups. Column 1 of both tables shows the results of the full sample of observations, investigating the heterogeneity related to Black hosts. The results indicate that the Superhost impact for Black hosts' number of reservations is almost double that for all other hosts. The confidence intervals indicate that this difference is generally statistically significant, with a few cases having slight overlap. A similar relationship is shown in column 2, which is based on a sample of only Black and White hosts.

Columns 3 and 4 display the results examining the heterogeneity of hosts without profile pictures. The results indicate that these hosts do not have a larger benefit from attaining Superhost status when compared to all other hosts (column 3) nor when when compared to White hosts (column 4). The confidence intervals have significant overlaps for these the coefficeint associated with these hosts. Overall, the results from this analysis are consistent with those reported in Section 5, indicating that Black hosts benefit substantially more than others from attaining Superhost status.

6.3. Heterogeneity with Reviews

Our results suggest that Superhost status provides information about hosts that potential visitors incorporate into their decision-making. An important factor in fully understanding the role of platform sponsorship is whether, and in what cases, other sources of information, such as user generated feedback render platform sponsorship inconsequential. More importantly, what is the heterogeneity of this relationship as it relates to host race. The double machine learning approach enables us to examine this issue. Specifically, we evaluate the heterogeneity of both host race and a host's reviews as they relate to the gain obtained from Superhost status. Figure 3 shows the results of this analysis with log(reservations) as the outcome variable. As expected, the results indicate that, for all hosts, as the number of reviews increases the impact of Superhost status diminishes.

[27] We explored the performance of the following methods from scikit-learn, all using the leave one out validation criteria: GradientBoostingRegressor, RandomForestRegressor, AdaBoostRegressor, BaggingRegressor, and SupportVectorRegressor.

However, the decrease is more pronounced for White hosts and hosts without a profile picture, with the impact statistically indifferent from zero for hosts with more than 100 reviews. On the other hand, Black hosts do witness a decrease, but the Superhost effect retains a higher magnitude, even for Black hosts with more than 100 reviews. These results provide further evidence as to the statistical discrimination facing Black hosts, especially without the platform sponsorship program. The results demonstrate that a Black host requires significantly more reviews than a White host for demand side participants to consider the information they have about the host adequate. That is, the results suggest that Airbnb visitors still utilize the additional information made available by platform Sponsorship for Black hosts with more than 100 reviews, but they do not require this information for other hosts with more than 100 reviews.

Extending the analysis to the *Revenue* and *Booking Window* produces similar findings. Figures 4 and 5 show the results related to these outcome variables and find that, unlike Black hosts, the impact of Superhost status is diminished after hosts accumulate reviews. Therefore, the findings indicate that, for all outcome variables we consider, Airbnb visitors do not find the information in reviews sufficient to mitigate statistical discrimination targeting Black hosts. The results indicate consistently that Black hosts with 100 or more accumulated reviews continue to benefit from gaining Superhost status while other hosts do not.

7. Conclusion

Digital marketplaces remove many traditional entry barriers, especially for minorities, and facilitate micro-entrepreneurship irrespective of the personal characteristics of entrepreneurs – race, origin, ethnicity, etc. The trust mechanism in such marketplaces, however, often require identifying information of the entrepreneurs (e.g., hosts on AirBnb, Taskers on TaskRabbit, etc.), accentuating concerns about discrimination against minority entrepreneurs. Using detailed performance data on the largest homesharing platform in the world, Airbnb, we examine discrimination facing supply side participants (hosts) on this platform. We utilize a unique quasi-random design that exploits the discontinuity in the criteria to an Airbnb Superhost and find that an average Superhost enjoys growth in reservations of approximately 10% from the platform certification. Consistent with the prediction of statistical discrimination theories, borderline Black Superhosts benefit disproportionately from the information shock—a 20% increase in reservations. While the documented inequity faced by the black hosts on AirBnb is troubling, we are encouraged by the fact that platform certification programs can be effective in mitigating the discrimination faced by minority entrepreneurs on digital marketplaces.

References

Aigner, Dennis J., Glen G. Cain. 1977. Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review* **30**(2) 175. doi:10.2307/2522871.

Arrow, Kenneth J. 1998. What Has Economics to Say about Racial Discrimination? *Journal of Economic Perspectives* **12**(2) 91–100. doi:10.1257/jep.12.2.91.

Blanchflower, David G., Phillip B. Levine, David J. Zimmerman. 2003. Discrimination in the small-business credit market. *Review of Economics and Statistics* **85**(4) 930–943. doi:10.1162/003465303772815835.

Bohren, J. Aislinn, Alex Imas, Michael Rosenberg. 2019. The dynamics of discrimination: Theory and evidence. *American Economic Review* **109**(10) 3395–3436. doi:10.1257/aer.20171829.

Bolton, Gary, Ben Greiner, Axel Ockenfels. 2013. Engineering trust: Reciprocity in the production of reputation information. *Management Science* **59**(2) 265–285. doi:10.1287/mnsc.1120.1609.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey. 2017. Double/debiased/neyman machine learning of treatment effects. *American Economic Review* **107**(5) 261–65.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, James Robins. 2018. Double/debiased machine learning for treatment and structural parameters.

Chetty, Raj, Nathaniel Hendren, Maggie R. Jones, Sonya R. Porter. 2020. Race and Economic Opportunity in the United States: an. *American Educational Research Association* **135**(2) 711–783.

Cui, Ruomeng, Jun Li, Dennis J. Zhang. 2020. Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on airbnb. *Management Science* **66**(3) 1071–1094. doi:10.1287/mnsc.2018.3273.

Dellarocas, Chrysanthos, Charles A. Wood. 2008. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Science* **54**(3) 460–476. doi:10.1287/mnsc.1070.0747.

Doleac, Jennifer L., Luke C.D. Stein. 2013. The visible hand: Race and online market outcomes. *Economic Journal* **123**(572) 469–492. doi:10.1111/ecoj.12082.

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).

Edelman, Benjamin G., Michael Luca. 2014. Digital Discrimination: The Case of Airbnb.com. *SSRN Electronic Journal* doi:10.2139/ssrn.2377353.

Elfenbein, Daniel W, Raymond Fisman, Brian McManus. 2015. Market structure, reputation, and the value of quality certification. *American Economic Journal: Microeconomics* **7**(4) 83–108.

Fairlie, Robert W., Alicia M. Robb. 2007. Why are black-owned businesses less successful than white-owned businesses? The role of families, inheritances, and business human capital. *Journal of Labor Economics* **25**(2) 289–323. doi:10.1086/510763.

Fradkin, Andrey, Elena Grewal, Dave Holtz, Matthew Pearson. 2015. Bias and Reciprocity in Online Reviews 641–641doi:10.1145/2764468.2764528.

Freeland, Robert E., Lisa A. Keister. 2016. How Does Race and Ethnicity Affect Persistence in Immature Ventures? *Journal of Small Business Management* **54**(1) 210–228. doi:10.1111/jsbm.12138.

Ge, Yanbo, Christopher R. Knittel, Don MacKenzie, Stephen Zoepf. 2020. Racial discrimination in transportation network companies. *Journal of Public Economics* **190** 104205. doi:10.1016/j.jpubeco.2020.104205. URL `https://doi.org/10.1016/j.jpubeco.2020.104205`.

Horton, John J. 2019. Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. *Management Science* **65**(8) 3518–3540. doi:10.1287/mnsc.2018.3116.

Jin, Ginger Zhe, Andrew Kato. 2006. Price, quality, and reputation: Evidence from an online field experiment. *RAND Journal of Economics* **37**(4) 983–1005. doi:10.1111/j.1756-2171.2006.tb00067.x.

Lang, Kevin, Ariella Kahn-lang Spitzer. 2020. Race Discrimination: An Economic Perspective. *Journal of Economic Perspectives* **34**(2) 68–89.

Laouénan, Morgane, Roland Rathelot. 2020. Can Information Reduce Ethnic Discrimination? Evidence from Airbnb .

Lee, David S, Thomas Lemieux. 2010. in Economics. *Journal of economic literature* **20**(1) 281–355.

Lennan, Christopher, Hao Nguyen, Dat Tran. 2018. Image quality assessment. `https://github.com/idealo/image-quality-assessment`.

Luca, Michael, Georgios Zervas. 2016. Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science* **62**(12) 3412–3427. doi:10.1287/mnsc.2015.2304.

Mayzlin, Dina, Yaniv Dover, Judith Chevalier. 2014. *Promotional reviews: An empirical investigation of online review manipulation*, vol. 104. doi:10.1257/aer.104.8.2421.

Mohammed, A.R. Shariq. 2020. Designing for Racial Impartiality: The Impact of Relocating Host Photos on the Airbnb Website. *Working Paper* (October). URL `https://static1.squarespace.com/static/59dc19e003596ed02a2348e3/t/59f123a0f9a61ef6a799a1fb/1508975527058/2017_10_25_Airbnb.pdf`.

Nosko, Chris;, Steven Tadelis. 2015. the Limits of Reputation in Platform Markets :. *National Bureau of Economic Research working paper* (20830).

Pope, Devin G, Justin R Sydnor, Devin G Pope, Justin R Sydnor. 2017. What ' s in a Picture ?: Evidence of Discrimination from Prosper . com What ' s in a Picture ? Evidence of Discrimination from Prosper . com **46**(1) 53–92.

Saeedi, Maryam. 2019. Reputation and adverse selection: theory and evidence from eBay. *The RAND Journal of Economics* **50**(4) 822–853.

Tadelis, Steven. 2016. Reputation and feedback systems in online platform markets. *Annual Review of Economics* **8** 321–340.

Younkin, Peter, Venkat Kuppuswamy. 2018. The colorblind crowd? Founder race and performance in crowdfunding. *Management Science* **64**(7) 3269–3287. doi:10.1287/mnsc.2017.2774.

Zervas, Georgios, Davide Proserpio, John W. Byers. 2020. A first look at online reputation on Airbnb, where every stay is above average. *Marketing Letters* doi:10.1007/s11002-020-09546-4.

Zhang, Shunyuan, Nitin Mehta, Param Vir Singh, Kannan Srinivasan. 2021. Can an ai algorithm mitigate racial economic inequality? an analysis in the context of airbnb. *Marketing Science* .

Figure 1 Airbnb Listings Sample

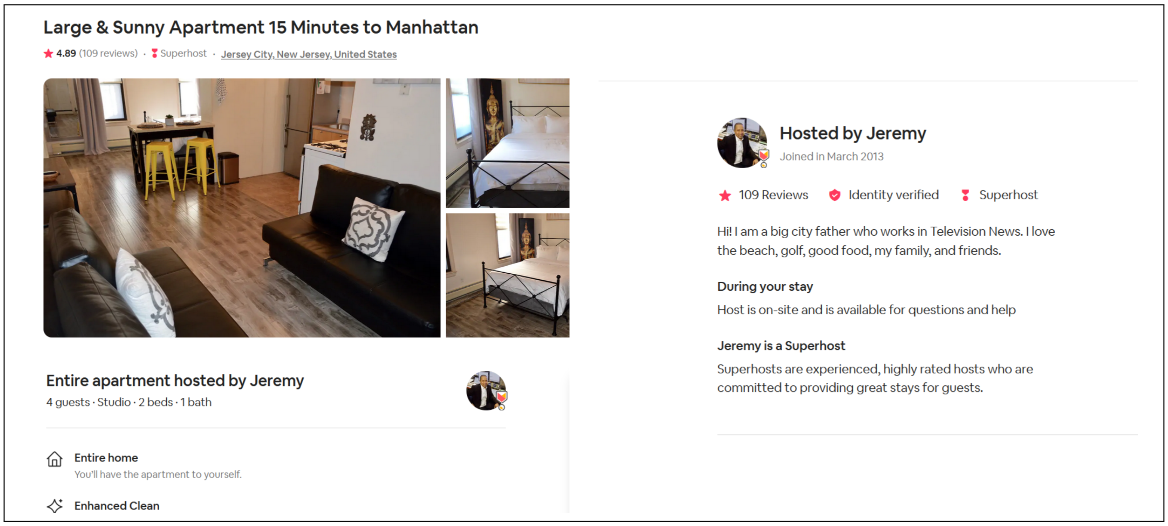

Note: This figure shows a sample Airbnb listing.

Figure 2 Airbnb Superhost: Additional 5-Star Reviews Needed

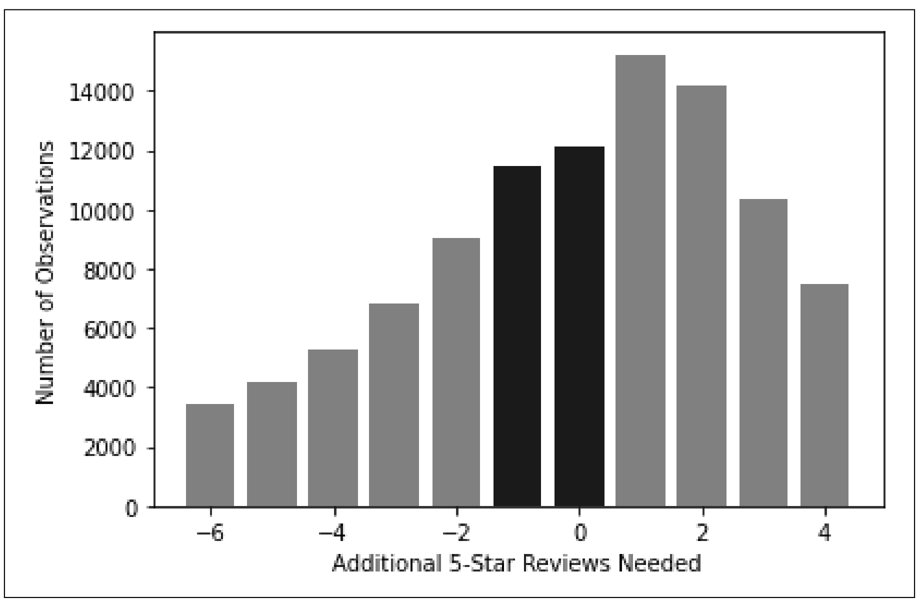

Figure 3 **Double Machine Learning Reviews Heterogeneity Results: log(Reservations)**

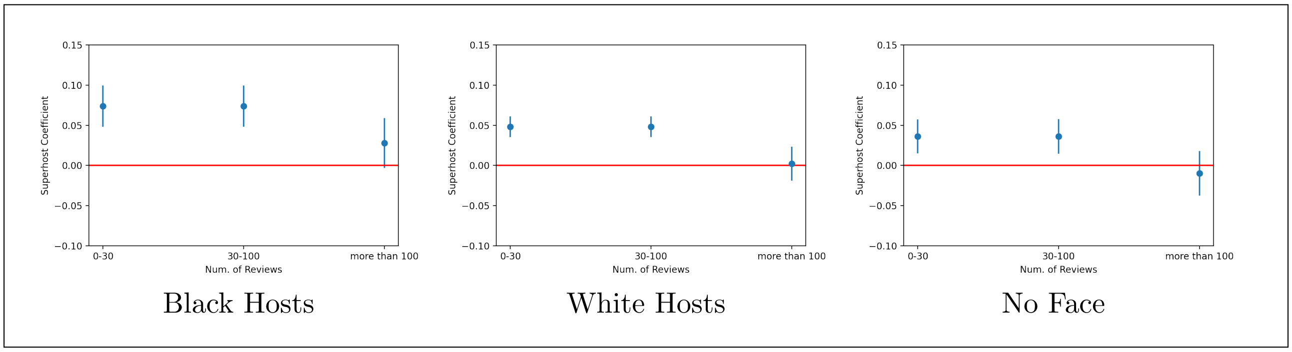

Figure 4 **Double Machine Learning Reviews Heterogeneity Results: log(Revenue)**

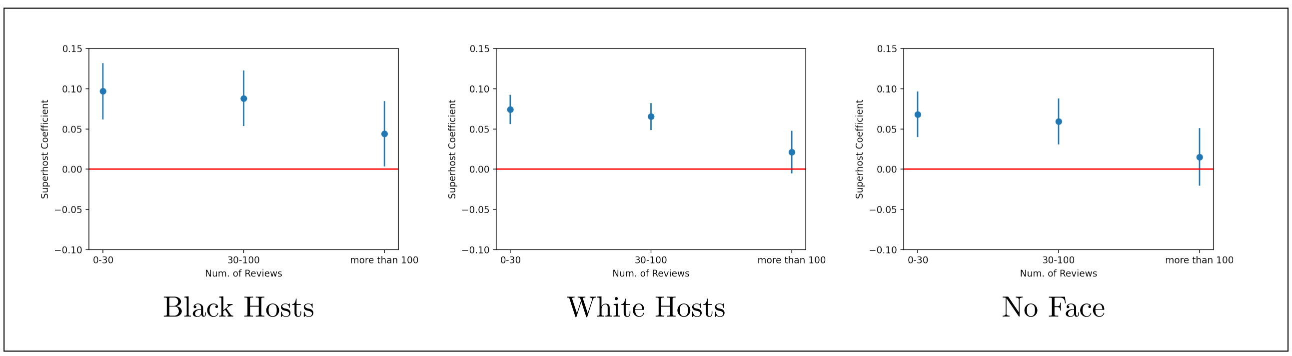

Figure 5 **Double Machine Learning Reviews Heterogeneity Results: log(Booking Window)**

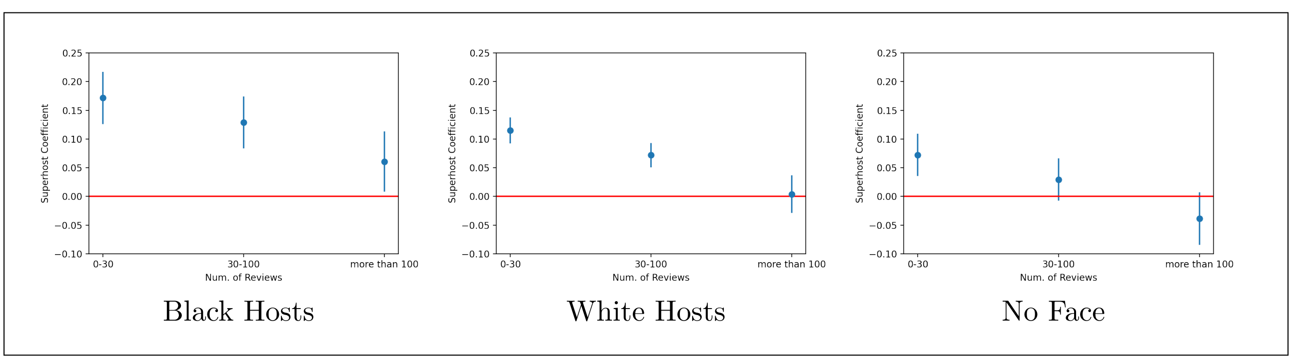

Table 1 Summary Statistics: Superhosts

	All Hosts			Hosts on Margin of Superhost Status		
	All (1)	Superhosts (2)	Not Superhosts (3)	All (4)	Superhosts (5)	Not Superhosts (6)
Observations	154,836	89,472	65,364	22,326	11,203	11,123
Reservations	7.83	7.94	7.69	5.96	6.13	5.79
Revenue ($)	3,680	3,852	3,445	3,135	3,282	2,987
Booking Window	40.37	41.66	38.57	39.47	40.22	38.70
Price ($)	126.78	131.53	120.27	126.13	125.16	127.10
Number of Reviews	51.58	54.43	47.68	34.27	34.92	33.61
Average Rating	94.84	97.28	91.50	94.86	95.74	93.98
Instant Bookable (1/0)	0.39	0.36	0.42	0.32	0.33	0.31
Shared Listing (1/0)	0.40	0.39	0.42	0.41	0.40	0.42
Num. of Bathrooms	1.24	1.24	1.23	1.19	1.20	1.19
Num. of Bedrooms	1.28	1.29	1.26	1.27	1.26	1.27
Zipcode Revenue Share	4.7%	4.9%	4.4%	4.2%	4.2%	4.1%
Image Aesthetic Score	5.05	5.07	5.02	5.02	5.04	5.01
Prop. Black Hosts	8.5%	8.1%	9.1%	8.6%	8.6%	8.5%
Prop. White Hosts	53.2%	55.4%	50.6%	52.3%	53.8%	50.9%

Table 2 Summary Statistics: Marginal Superhosts by Race

	Black Hosts			White Hosts		
	All (1)	Superhosts (2)	Not Superhosts (3)	All (4)	Superhosts (5)	Not Superhosts (6)
Observations	1,910	967	943	11,679	6,022	5,657
Reservations	5.91	6.33	5.47	5.70	5.87	5.53
Revenue ($)	2,424	2,598	2,245	3,222	3,342	3,094
Booking Window	38.08	39.18	36.85	40.96	41.27	40.63
Price ($)	97.90	96.62	99.22	131.74	129.83	133.77
Number of Reviews	33.25	34.13	32.34	36.39	36.72	36.03
Average Rating	94.52	95.51	93.50	94.96	95.82	94.05
Instant Bookable (1/0)	0.39	0.41	0.36	0.28	0.29	0.27
Shared Listing (1/0)	0.50	0.49	0.50	0.39	0.38	0.40
Num. of Bathrooms	1.17	1.17	1.18	1.18	1.19	1.17
Num. of Bedrooms	1.24	1.22	1.25	1.25	1.25	1.25
Zipcode Revenue Share	2.9%	3.0%	2.8%	4.5%	4.5%	4.4%
Image Aesthetic Score	4.95	4.95	4.94	5.03	5.04	5.02

Table 3 Summary Statistics: Marginal Superhosts by Race; Majority Black Zipcodes

	Black Hosts			White Hosts		
	All (1)	Superhosts (2)	Not Superhosts (3)	All (4)	Superhosts (5)	Not Superhosts (6)
Observations	809	405	404	790	438	352
Reservations	5.49	6.09	4.89	5.55	5.63	5.45
Revenue ($)	2,045	2,223	1,866	2,196	2,269	2,106
Booking Window	38.07	39.13	36.87	45.43	47.94	42.49
Price ($)	91.74	87.12	96.37	89.61	90.38	88.64
Number of Reviews	33.16	32.49	33.83	36.09	34.74	37.77
Average Rating	94.07	95.47	92.68	94.94	95.60	94.12
Instant Bookable (1/0)	0.38	0.39	0.37	0.29	0.29	0.28
Shared Listing (1/0)	0.50	0.50	0.50	0.54	0.55	0.54
Num. of Bathrooms	1.18	1.17	1.19	1.18	1.22	1.14
Num. of Bedrooms	1.28	1.27	1.28	1.27	1.30	1.24
Zipcode Revenue Share	1.1%	1.1%	1.2%	1.5%	1.6%	1.4%
Image Aesthetic Score	4.93	4.95	4.91	5.02	5.02	5.02

Table 4 Summary Statistics: Marginal Superhosts by Race; Majority White Zipcodes

	Black Hosts			White Hosts		
	All (1)	Superhosts (2)	Not Superhosts (3)	All (4)	Superhosts (5)	Not Superhosts (6)
Observations	288	134	154	5,312	2,731	2,581
Reservations	6.58	7.28	5.97	5.75	6.01	5.47
Revenue ($)	3,292	3,611	3,015	3,905	4,135	3,663
Booking Window	38.14	42.83	33.82	42.29	42.44	42.12
Price ($)	123.76	125.69	122.08	157.74	157.09	158.42
Number of Reviews	34.90	39.51	30.88	37.89	38.35	37.41
Average Rating	94.91	95.40	94.49	95.00	95.80	94.16
Instant Bookable (1/0)	0.31	0.36	0.27	0.26	0.28	0.25
Shared Listing (1/0)	0.43	0.38	0.47	0.33	0.31	0.36
Num. of Bathrooms	1.19	1.17	1.21	1.18	1.19	1.17
Num. of Bedrooms	1.15	1.13	1.17	1.28	1.29	1.27
Zipcode Revenue Share	4.7%	4.9%	4.6%	4.8%	4.9%	4.7%
Image Aesthetic Score	4.95	4.97	4.93	5.06	5.08	5.04

Table 5 Comparing Averages Between Marginal Superhosts and Marginal Not Superhosts

	Not Superhosts		Superhosts			
	Mean	**St. Dev.**	**Mean**	**St. Dev.**	**T-Stat.**	**P-Value**
Previous Period Posted Price	168.97	189.29	165.79	133.14	1.32	0.19
Previous Period Number of Reviews	26.26	34.40	26.73	34.54	-0.93	0.35
Number of Bathrooms	1.19	0.49	1.20	0.54	-2.29	0.02
Num. of Bedrooms	1.27	0.81	1.26	0.79	0.33	0.74
Previous Period Zipcode Revenue Share	0.04	0.05	0.04	0.05	-1.02	0.31

Table 6 Comparing Averages Between Superhosts and Not Superhosts: 3 reviews away from Margin

	Not Superhosts		Superhosts			
	Mean	**St. Dev.**	**Mean**	**St. Dev.**	**T-Stat.**	**P-Value**
Previous Period Posted Price	167.47	134.38	179.96	181.57	-4.61	0.00
Previous Period Number of Reviews	32.97	38.93	30.67	37.82	3.73	0.00
Number of Bathrooms	1.18	0.48	1.23	0.56	-6.37	0.00
Num. of Bedrooms	1.28	0.83	1.31	0.85	-2.35	0.02
Previous Period Zipcode Revenue Share	0.04	0.05	0.05	0.05	-5.63	0.00

Table 7 Main Superhost Impact: All Dependent Variables

	(1) log(Reservations)	(2) log(Revenue)	(3) log(Booking Window)
Superhost	**0.095***** (0.013)	**0.208***** (0.042)	**0.077***** (0.019)
log(Number of Reviews)	0.276*** (0.007)	0.712*** (0.023)	0.171*** (0.010)
Listing Rating	0.008*** (0.001)	0.036*** (0.004)	0.011*** (0.003)
Instant Bookable (1/0)	0.481*** (0.013)	0.942*** (0.041)	0.068*** (0.019)
Shared Listings (1/0)	-0.277*** (0.017)	-1.022*** (0.059)	-0.417*** (0.026)
Num. of Bedrooms	0.065*** (0.010)	0.240*** (0.033)	0.113*** (0.014)
Num. of Bathrooms	0.035** (0.014)	0.055 (0.047)	0.025 (0.023)
log(Available Days Listing Price)	-0.245*** (0.016)	0.004 (0.060)	-0.051** (0.025)
log(Available Days)	0.428*** (0.007)	1.276*** (0.029)	0.010 (0.018)
Zipcode Revenue Share	1.008*** (0.148)	2.713*** (0.492)	1.280*** (0.203)
Image Aesthetic Score	0.123*** (0.015)	0.270*** (0.050)	0.085*** (0.022)
Period Fixed Effects	Yes	Yes	Yes
City Fixed Effects	Yes	Yes	Yes
Constant	-1.494*** (0.154)	-6.417*** (0.567)	1.633*** (0.309)
Observations	15,362	15,362	12,480
R-squared	0.352	0.294	0.135

*Notes: Standard Errors are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.*

Table 8 Superhost Heterogeneity Results by Race: Reservations

Dep. Variable: log(Reservations)	(1)	(2)	(3)	(4)
Superhost	**0.087***** (0.013)	**0.092***** (0.018)	**0.089***** (0.013)	**0.089***** (0.018)
Black Host (1/0)	-0.148*** (0.029)	-0.122*** (0.031)		
Superhost x Black Host (1/0)	0.096** (0.044)	0.095** (0.046)		
No Face Host (1/0)			0.054** (0.027)	0.067** (0.028)
Superhost x No Face Host (1/0)			0.056 (0.040)	0.056 (0.042)
log(Number of Reviews)	0.276*** (0.007)	0.279*** (0.009)	0.278*** (0.007)	0.273*** (0.008)
Listing Rating	0.008*** (0.001)	0.008*** (0.001)	0.008*** (0.001)	0.008*** (0.001)
Instant Bookable (1/0)	0.483*** (0.013)	0.430*** (0.018)	0.479*** (0.013)	0.450*** (0.017)
Shared Listings (1/0)	-0.279*** (0.017)	-0.278*** (0.022)	-0.275*** (0.017)	-0.271*** (0.022)
Num. of Bedrooms	0.066*** (0.010)	0.079*** (0.013)	0.066*** (0.010)	0.073*** (0.012)
Num. of Bathrooms	0.036** (0.014)	0.025 (0.023)	0.031** (0.014)	0.037** (0.015)
log(Available Days Listing Price)	-0.251*** (0.016)	-0.245*** (0.020)	-0.247*** (0.016)	-0.250*** (0.020)
log(Available Days)	0.430*** (0.007)	0.429*** (0.009)	0.426*** (0.007)	0.434*** (0.009)
Zipcode Revenue Share	0.966*** (0.148)	0.988*** (0.196)	0.987*** (0.148)	0.900*** (0.176)
Image Aesthetic Score	0.120*** (0.015)	0.106*** (0.019)	0.122*** (0.015)	0.117*** (0.019)
Period Fixed Effects	Yes	Yes	Yes	Yes
City Fixed Effects	Yes	Yes	Yes	Yes
Constant	-1.441*** (0.154)	-1.446*** (0.197)	-1.479*** (0.153)	-1.519*** (0.198)
Observations	15,362	9,079	15,362	9,419
R-squared	0.353	0.353	0.353	0.356

*Notes: Standard Errors are reported in parentheses. * p < .10, ** p < .05, *** p < .01.*

Table 9 Superhost Heterogeneity Results by Race: Revenue

Dep. Variable: log(Revenue)	(1)	(2)	(3)	(4)
Superhost	**0.181******	**0.177******	**0.195******	**0.172******
	(0.044)	(0.059)	(0.044)	(0.059)
Black Host (1/0)	-0.411***	-0.411***		
	(0.098)	(0.103)		
Superhost x Black Host (1/0)	0.326**	0.328**		
	(0.142)	(0.147)		
No Face Host (1/0)			0.040	0.027
			(0.090)	(0.095)
Superhost x No Face Host (1/0)			0.114	0.120
			(0.130)	(0.136)
log(Number of Reviews)	0.712***	0.723***	0.714***	0.712***
	(0.023)	(0.031)	(0.023)	(0.029)
Listing Rating	0.036***	0.038***	0.036***	0.039***
	(0.004)	(0.005)	(0.004)	(0.005)
Instant Bookable (1/0)	0.948***	0.868***	0.939***	0.925***
	(0.041)	(0.054)	(0.041)	(0.053)
Shared Listings (1/0)	-1.027***	-1.089***	-1.020***	-1.082***
	(0.059)	(0.078)	(0.059)	(0.077)
Num. of Bedrooms	0.242***	0.265***	0.242***	0.276***
	(0.033)	(0.045)	(0.033)	(0.041)
Num. of Bathrooms	0.057	0.120	0.049	0.084
	(0.047)	(0.079)	(0.048)	(0.053)
log(Available Days Listing Price)	-0.009	-0.050	0.003	-0.057
	(0.060)	(0.079)	(0.060)	(0.077)
log(Available Days)	1.282***	1.278***	1.274***	1.294***
	(0.029)	(0.038)	(0.029)	(0.038)
Zipcode Revenue Share	2.603***	2.460***	2.689***	2.486***
	(0.493)	(0.676)	(0.493)	(0.593)
Image Aesthetic Score	0.262***	0.236***	0.270***	0.238***
	(0.050)	(0.066)	(0.050)	(0.065)
Period Fixed Effects	Yes	Yes	Yes	Yes
City Fixed Effects	Yes	Yes	Yes	Yes
Constant	-6.281***	-6.362***	-6.400***	-6.593***
	(0.569)	(0.743)	(0.567)	(0.734)
Observations	15,362	9,079	15,362	9,419
R-squared	0.295	0.296	0.294	0.294

*Notes: Standard Errors are reported in parentheses. * p < .10, ** p < .05, *** p < .01.*

Table 10 Superhost Heterogeneity Results by Race: Booking Window

Dep. Variable: log(Booking Window)	(1)	(2)	(3)	(4)
Superhost	**0.069*****	**0.032**	**0.082*****	**0.038**
	(0.019)	(0.026)	(0.020)	(0.025)
Black Host (1/0)	-0.156***	-0.212***		
	(0.048)	(0.050)		
Superhost x Black Host (1/0)	0.105	0.134*		
	(0.068)	(0.070)		
No Face Host (1/0)			-0.043	-0.086**
			(0.040)	(0.042)
Superhost x No Face Host (1/0)			-0.045	-0.006
			(0.059)	(0.061)
log(Number of Reviews)	0.170***	0.157***	0.169***	0.162***
	(0.010)	(0.013)	(0.010)	(0.013)
Listing Rating	0.011***	0.014***	0.011***	0.009**
	(0.003)	(0.004)	(0.003)	(0.004)
Instant Bookable (1/0)	0.070***	0.092***	0.070***	0.080***
	(0.019)	(0.025)	(0.019)	(0.024)
Shared Listings (1/0)	-0.421***	-0.442***	-0.418***	-0.448***
	(0.026)	(0.034)	(0.026)	(0.033)
Num. of Bedrooms	0.114***	0.091***	0.111***	0.120***
	(0.014)	(0.018)	(0.014)	(0.017)
Num. of Bathrooms	0.027	0.130***	0.029	0.027
	(0.023)	(0.031)	(0.023)	(0.027)
log(Available Days Listing Price)	-0.059**	-0.121***	-0.051**	-0.094***
	(0.025)	(0.033)	(0.025)	(0.031)
log(Available Days)	0.013	0.051**	0.012	0.047**
	(0.018)	(0.022)	(0.018)	(0.022)
Zipcode Revenue Share	1.250***	1.145***	1.302***	1.036***
	(0.203)	(0.270)	(0.204)	(0.245)
Image Aesthetic Score	0.082***	0.107***	0.085***	0.074***
	(0.022)	(0.029)	(0.022)	(0.028)
Period Fixed Effects	Yes	Yes	Yes	Yes
City Fixed Effects	Yes	Yes	Yes	Yes
Constant	1.695***	1.513***	1.625***	2.070***
	(0.308)	(0.472)	(0.310)	(0.419)
Observations	12,480	7,332	12,480	7,662
R-squared	0.136	0.136	0.136	0.130

*Notes: Standard Errors are reported in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.*

Table 11 Double Machine Learning Results: Reservation and Revenue

Dependent Variable		(1) Full Sample	(2) Black and White Hosts	(3) Full Sample	(4) No Face and White Hosts
log(Reservations)	ATE	0.047 (0.006) [0.037 , 0.057]	0.058 (0.008) [0.046 , 0.071]	0.052 (0.007) [0.041 , 0.062]	0.049 (0.008) [0.036 , 0.061]
	Black Hosts	0.082 (0.015) [0.057 , 0.106]	0.094 (0.015) [0.07 , 0.119]		
	All Other Hosts	0.044 (0.005) [0.036 , 0.052]		0.054 (0.005) [0.046 , 0.062]	
	White Hosts		0.053 (0.006) [0.044 , 0.063]		0.051 (0.006) [0.041 , 0.06]
	No Face Hosts			0.039 (0.012) [0.019 , 0.059]	0.04 (0.012) [0.02 , 0.061]
log(Revenue)	ATE	0.072 (0.009) [0.058 , 0.086]	0.071 (0.01) [0.054 , 0.088]	0.075 (0.009) [0.061 , 0.09]	0.065 (0.01) [0.048 , 0.082]
	Black Hosts	0.107 (0.02) [0.074 , 0.141]	0.119 (0.021) [0.085 , 0.153]		
	All Other Hosts	0.069 (0.007) [0.058 , 0.079]		0.076 (0.007) [0.065 , 0.088]	
	White Hosts		0.064 (0.008) [0.051 , 0.078]		0.069 (0.008) [0.056 , 0.082]
	No Face Hosts			0.068 (0.016) [0.042 , 0.095]	0.049 (0.016) [0.023 , 0.076]

Notes: Standards errors are reported in parentheses and 95% confidence intervals are reports in brackets.

Table 12 Double Machine Learning Results: Booking Window

Dependent Variable		(1) Full Sample	(2) Black and White Hosts	(3) Full Sample	(4) No Face and White Hosts
	ATE	0.094 (0.011) [0.076 , 0.112]	0.096 (0.013) [0.075 , 0.118]	0.096 (0.011) [0.077 , 0.114]	0.075 (0.013) [0.054 , 0.095]
log(Booking Window)	Black Hosts	0.145 (0.026) [0.102 , 0.188]	0.135 (0.027) [0.091 , 0.179]		
	All Other Hosts	0.089 (0.008) [0.076 , 0.103]		0.105 (0.009) [0.091 , 0.119]	
	White Hosts		0.091 (0.01) [0.075 , 0.107]		0.087 (0.01) [0.071 , 0.104]
	No Face Hosts			0.043 (0.021) [0.009 , 0.077]	0.021 (0.02) [-0.012 , 0.055]

Notes: Standards errors are reported in parentheses and 95% confidence intervals are reported in brackets.