Journal of Housing Economics 46 (2019) 101635

Contents lists available at ScienceDirect

Journal of Housing Economics

journal homepage: www.elsevier.com/locate/jhec

The impact of host race and gender on prices on Airbnb☆

Anya Marchenko

Becker Friedman Institute, University of Chicago, United States

ARTICLE INFO

Keywords:
Airbnb
Discrimination
Race
Online marketplace

ABSTRACT

This study investigates the impact of host race and gender on Airbnb property prices. I use an existing dataset of Airbnb listings and visually inspect 70,000 host profile pictures to code host demographics. I estimate that Asian hosts earn 4–5%, and Black male hosts 3%, less than White males for the same type of property. However, controlling for more observables weakens the effects, requiring a cautious interpretation of these point estimates. I use two proxies for the number of bookings a listing has to estimate whether a demand or supply shift is responsible for the price disparity. I find that despite the lower prices they charge for listings, minority hosts face lower demand. These findings are consistent with, but not conclusive of, the presence of discrimination.

1. Introduction

Over $514 billion worth of transactions now occur online (U.S. Census Bureau, 2019). A growing sector of online commerce is peer-to-peer (P2P) platforms, in which two people interact directly as buyers and sellers, rather than making purchases at a brick-and-mortar store or visiting the website of a business. Sharing economy businesses have sprung up as the newest incarnation of the P2P model, allowing members to buy or share certain goods and services through a central platform. Buyers can bike, car, or apartment share by simply renting the product for a short period of time from the seller.

Unlike in traditional markets, participants on these platforms most often do not meet in person before agreeing to a transaction. Moreover, strangers on these platforms are not likely to have the reputational safety of a brick-and-mortar store, and thus have a harder time credibly guaranteeing the quality of their product. These factors contribute to a higher perceived risk of transacting online.

As a result, many sharing economy platforms have instituted measures to mitigate this risk. Most have user profiles and encourage users to post their names and photos, as well as descriptions of themselves and their products, to bolster credibility. Platforms often solicit reviews from previous buyers that are posted alongside the product, or on the seller's profile, to aid future buyers by increasing transparency and thereby encouraging transactions. However, the presence of identifiable information about a person's demographic characteristics gives rise to the possibility of discrimination. Indeed, a growing body of research indicates that P2P market participants discriminate online in much the same ways as they do in traditional markets.

This paper leverages the standardization provided by P2P platforms to measure the effect of host race on listing price on the world's biggest short-term rental platform, Airbnb. Using data scraped from the Airbnb website, I code the demographic information of 70,000 Airbnb hosts throughout the country. I find that non-White hosts, both male and female, have lower prices than White hosts on the platform.

Previous literature has examined discrimination against Airbnb hosts in various small samples (Edelman and Luca, 2014; Wang et al., 2015; Kakar et al., 2018). All three studies have found evidence of a price disparity for minority hosts of different races. The first study of this kind estimated that non-Black hosts have higher prices than Black hosts in New York City (Edelman and Luca, 2014). Wang et al. (2015) and (Kakar et al., 2018) did not find lower prices for Black hosts, but did find evidence of a price disparity for Asian hosts in San Francisco's Bay Area.

The comprehensiveness of mitigating controls varies enormously in this literature. Wang et al. (2015) and (Edelman and Luca, 2014) were limited by the sparse listing information available on the Airbnb website (Airbnb has since added more comprehensive listing details). For instance, Edelman and Luca (2014) only control for a few property characteristics, the quality of the host's reviews, and a measure of the reliability of the host. However, there are still many other observables,

☆ I thank Steven D. Levitt for valuable advice. Kirthi Bellamkonda, Jacob Dorn, Michal Dzitko, Michael Galperin, Sonia Jaffe, Ezra Karger, Sylvia Klosin, Victor Lima, and Kotaro Yoshida provided helpful feedback and comments. Abdallah Aboelela, Melody Jih, and Thiago Resende provided excellent research assistance. I thank Mahathi Ayyagari, Melody Jih, Fong Chai, Lydia Sum, Joseph Day, Leah Umanskiy, Cristian Raygoza, Putter Thepkanjana, Spencer Kang, and Tony Song for coding host demographics. All remaining errors are my own. Data was provided by InsideAirbnb.com, founded and maintained by Murray Cox. Airbnb was not consulted in this research.

E-mail address: amarchenko@uchicago.edu.

https://doi.org/10.1016/j.jhe.2019.101635
Received 10 July 2018; Received in revised form 29 April 2019; Accepted 12 July 2019
Available online 12 August 2019

such as the type of listing, and unobservables, such as the quality of the host themselves, which could differ by race and contribute to variation in prices.

In sum, no study has measured a price disparity in a large, representative, multi-city sample of Airbnb listings. There exists no research conclusively showing whether or not minority hosts face a price disparity in the nationwide Airbnb market, whether the price disparity differs in magnitude by city, or whether hosts of a certain race fare better in some cities and worse in others. This paper seeks to fill these gaps.

I build on the existing literature in the following ways. First, I bring more robust data than existing studies to the question. My sample includes seven large urban centers that cover each geographic region in the United States. Second, I control for a comprehensive set of confounders that includes all information available to a guest on the listing page. To account for unobservable differences in host quality, I use machine learning techniques to analyze the descriptions on listing pages. These controls comprise the rest of the observable information on the listing page that could influence demand for that listing. Third, I code the demographic characteristics of reviewers in Chicago, to examine whether or not minority hosts have worse reviews, and whether review quality differs by the reviewer's race. Fourth, I use predictions from basic microeconomic theory, and proxies for the quantity demanded of a listing, to test hypotheses that could explain the price disparity.

I find a negative effect of minority host race on listing prices on Airbnb. My point estimates indicate that the price disparity between White and non-White hosts is widest for Asian hosts. The prices of Asian hosts of both sexes are roughly 4.5% lower than White male hosts who own a property in a similar location with similar characteristics. The second biggest effect is for Black male hosts, whose listed price is lower by around 3%.[1] While my coefficients suggest an effect of minority host race on price, the point estimates decrease as I add controls. This raises the question of whether controlling for other unobservables would further decrease the point estimates. In Section 4, I use Oster (2017)'s analysis of selection on unobservables to address this concern.

I also explore whether the price disparity holds in individual cities and across different types of listings. I find that the price disparity is driven by lower-priced, newer listings with few reviews. This finding is consistent with the presence of statistical discrimination, as in the absence of a clear signal about a listing's quality, a discriminatory guest might use the host race as a proxy for quality. In the main sample that includes all cities, my estimates are smaller than those measured by the previous studies. However, this fact is unsurprising, as it is unlikely that we would measure the same disparity when there are important structural differences between this paper and previous literature in the sample, sample size, and empirical approach.

There are two theoretical reasons that could explain why minority hosts systemically charge lower prices than White hosts for listings of the same quality. The first is that low prices are a response to lower demand for their listings. Throughout the paper, I refer to this as a demand shift, or a demand effect. Statistical discrimination as laid out by Becker (1957) is one explanation for lower demand. In the Airbnb market, if a guest is discriminating, then given two comparable listings they would be willing to pay more to avoid staying in the listing owned by a minority host. Responding to (or anticipating) this lower demand, minority hosts rationally post a lower price.

The second theoretical reason for a lower price is a higher supply of minority host's listings at each price point. I refer to this as a supply shift, or supply effect. A supply shift could be due to a lower marginal cost of operating the listing for minority hosts relative to White hosts. For instance, minority hosts could charge a lower price for their listing because it is marginally cheaper for them to operate that listing relative to a White host. Since Black and Hispanic workers tend to earn less than their White counterparts, even for the same amount of education, they may have a lower opportunity cost of time (Economic Policy Institute, 2016). They would therefore have a lower marginal cost of managing their listing, and so would set lower prices than White hosts with comparable listings.

Basic microeconomic theory states that the quantity demanded in a market can help differentiate between a demand versus a supply effect. A lower price and lower quantity demanded points to a demand shift, whereas a lower price but higher quantity demanded points to a supply shift. In the context of Airbnb, the true measure of a listing's quantity demanded would be its number of bookings. However, I do not observe this value in my data. Throughout my analysis, I therefore use two proxies for the number of bookings: the number of reviews of a listing and its vacancy rate. I use these two measures of quantity demanded to address whether the price disparity is caused by a demand or supply effect.

My first measure of quantity demanded is the number of reviews. I find that Black hosts, White female hosts, and Asian female hosts have 5–10% fewer reviews than White hosts for a listing that has spent the same amount of time on the market. Effects for all other host categories are slightly negative, but not statistically significant.[2]

One potential explanation for a lower number of reviews is that a minority host might make their listing available to guests less frequently than a White host. A host controls how many days of the month they offer their listing for rent via an availability calendar on the listing page. When a guest books their listing, the booked days disappear from the availability calendar. Therefore, the availability is actually a measure of the vacancy rate for the listing. If a host has a fewer reviews, perhaps this is because they offer their listing for fewer days of the month. To test this, I regress a listing's vacancy rate on host race. Results show that contrary to this hypothesis, the listings of Black hosts stay vacant 1.5–2.5 days per month *longer* than the listings of White hosts. The listings of White female hosts and Asian female hosts, on the other hand, are vacant slightly less frequently than listings owned by White hosts, meaning that higher vacancy rate could be driving lower review numbers.

In sum, no minority host group has a higher number of reviews than White hosts. This fact provides suggestive evidence that minority hosts' quantity demanded on Airbnb is *at worst* not higher than those of White hosts. At least part of the lower number of reviews for Asian females and White females can be explained by differences in availability. Black hosts, on the other hand, have a statistically significant lower number of reviews, but higher availability. Surprisingly, this must mean that

[1] This effect is statistically significant at the $p < .001$ level for Asian female hosts, and the $p < .01$ level for Asian and Black male hosts. For Hispanic hosts and Black female hosts, the effect is small, around 2%, and is not statistically significant. If I regress price on a dummy for the host race, instead of on an interaction of host race and gender, all of the effects become significant. In this regression, the effects are 2.7% for Black hosts, 2.2% for Hispanic hosts, and 4.4% for Asian hosts, and all effects are significant to at least the $p < .01$ level.

[2] This conclusion is only salient if the total number of reviews is a reasonable proxy for the demand of a listing. Yet, one can imagine that if reviewers systematically under-review minority hosts relative to White hosts, these groups would have lower numbers of reviews that do not necessarily represent a lower quantity demanded. There is no way to tell apart these mechanisms in my data. A recent study found that reviews left by hosts on guests' pages can significantly reduce discrimination and render acceptance rates of guests with White-sounding names and African American-sounding names statistically indistinguishable (Cui et al., 2016), but it remains unknown whether or not reviewers discriminate against minorities in leaving reviews (Ye et al., 2017). If reviewers systematically under-review minority hosts, this itself could be evidence of discrimination. My working assumption is that even if not every guest leaves a review, the review proportion is similar across host race, and a lower number of reviews therefore indicates a real difference between quantity demanded of minority hosts and White hosts.

consumers do not respond to the lower prices of Black hosts by increasing demand, as the listings of Black hosts have fewer guests and stay vacant for longer than comparable listings operated by White hosts.

A demand shift could also be caused by differences in review quality. Host race could drive differences in reviews because minority hosts are of worse quality, or because discriminatory guests rate minority hosts lower for the same quality of stay. If minority hosts have worse reviews on their listing page, it might decrease demand by discouraging future guests. To this end, I estimate the effect of host race on review quality. I use the race and gender of the reviewer, and the host, to compare the sentiment (how favorable or unfavorable the review is) of the reviews that guests leave for White and for minority hosts.[3] I find that Black and Hispanic males have reviews that are 0.1–0.2 standard deviations worse than White male hosts, significant at the $p < .01$ level. Lower quality reviews might explain why the listings of Black host are priced lower, but are also less demanded, than the listings of White male hosts.

1.1. About Airbnb

Airbnb is a sharing economy platform founded in 2008 that allows people to rent out their apartment, house, or a single room to short-term lodgers. As of 2017, it had more than 3 million listings, more than Marriott's 1.2 million rooms worldwide (Airbnb, 2017). Just like traditional hotel chains, guests on Airbnb can browse listings by city and property type, and book a stay based on prices, location, past reviews, pictures of the listing, size, and amenities. Unlike traditional hotel chains, however, hosts create a profile for themselves and a page for each listing they are renting. Each listing page includes the name and picture of the host, the reviews left by previous guests, and those guests' profile pictures. Guests can therefore infer demographic information about the host through a host's picture and name, creating the opportunity for discrimination. Figs. 1 and 2 present screenshots of a listing in a Chicago neighborhood, illustrating some of the information that would be available to a potential guest.

1.2. Previous literature

1.2.1. Brief theoretical background

Becker (1957) proposed the idea that discrimination against a group is reflected in the prices that that group charges in a particular market, be it labor or products. In the Airbnb market, Becker's market discrimination would be reflected in the price that the guest (buyer) pays to the host (seller) to stay with them. If the guest is discriminating, then given two comparable listings, they would choose not to stay in the one owned by a minority host. Responding to, or predicting, a lower demand, minority hosts rationally respond by posting a lower price and, despite this, face a lower quantity demanded.

Becker was concerned with discrimination arising from face-to-face interactions between minority and majority groups. Since then, there has been a large amount of research indicating that Becker's theory holds for people participating in online markets for labor, lending, rental, and products. In these cases, participants simply bring their prejudices online and use names and photos to discriminate. Next, I detail the research that explores how the theory of discrimination plays out in both traditional and online markets.

1.2.2. Research on discrimination in traditional housing markets

African–Americans experience pervasively worse outcomes in the housing market as a result of historic and current racial discrimination (Krysan et al., 2014). Even after the gains during the Civil Rights Era, such as the landmark Fair Housing Act of 1968, discrimination in the housing market has been widely documented by social scientists. African-American renters are told that there are 30% fewer available housing units than White renters (Yinger, 1986). African-American families face higher barriers when raising capital to purchase a home (Pope and Sydnor, 2011). E-mails sent to landlords from home-seekers with typically African-American names receive lower response rates than emails sent by those with names commonly associated with White people (Hanson and Hawley, 2011).

Economists have primarily studied discrimination against African-American tenants. There is little research on the other side of the market – when African–Americans are supplying, rather than demanding, housing. Because property ownership cannot be randomized, it is difficult to disentangle true discrimination from systematic differences in the housing owned by African-Americans and White landlords.[4] Some studies have found evidence that African–American homeowners are more likely to be targeted by subprime loans (Rugh and Massey, 2010) or pay more than White people for similar housing (Bayer et al., 2017; Myers, 2004), but few studies have a credible identification strategy to separate discrimination from correlated observables.

1.2.3. Research on discrimination in P2P online commerce

Doleac and Stein (2010) examined the effect of apparent race on market outcomes when selling an iPod on various online marketplaces. In some pictures, a dark-skinned hand was holding the iPod, signaling a Black seller, while in others, a light-skinned hand was holding the iPod, signaling a White seller. Hands which indicate a Black seller received 18% fewer and 11% lower offers than White sellers. Furthermore, bidders were less likely to include their name in offers made to Black sellers. Pope and Sydnor (2011) found that in a P2P lending market (Prosper.com) demographic characteristics conveyed through pictures and text significantly affected loan terms for Black borrowers. Black borrowers were 25%–35% less likely to receive loans than White borrowers with similar credit profiles, and loans received by Black borrowers had an interest rate 60–80 basis points higher than White borrowers.

In sharing economies, a similar pattern occurs. Ge et al. (2016) explored the effect of race on market outcomes in rideshare platforms. These platforms provide rider information such as the first name, photo, and rating to drivers before (Lyft) or at (Uber) the time of ride acceptance. The authors ordered 1500 Uber and Lyft trips to measure the impact of rider race on fare and wait times, varying the apparent race of the rider in user photos, as well as degree to which riders' names were distinctively Black. Uber riders who use distinctively Black names experience up to 35% longer wait times and more frequent cancellations than riders who use distinctively White names, especially for males in low population density areas.

These studies suggest that users of online platforms use visual and textual information to transfer their racial biases from the real world into the online world. Though it is difficult to confirm whether these biases stem from statistical or taste-based discrimination, they provide evidence that P2P market participants like Uber drivers, Prosper lenders, and iPod buyers use racial information in their market decisions.

1.2.4. Research on Airbnb generally

The appeal of easily accessible, affordable, and short-term peer-to-

[3] Since it required hand-coding, demographic information of the reviewers is only available for a randomly-chosen subset of hosts in Chicago.

[4] One would also expect Black landlords to fare worse than White landlords in this area as well. Properties owned by African–Americans tend to be less expensive than those owned by White Americans. The average Black household still has less mean wealth than a White household (Oliver and Shapiro, 2006). Even middle-class Black and Hispanic households still live in neighborhoods with median incomes similar to those of very poor White neighborhoods (Reardon et al., 2015).

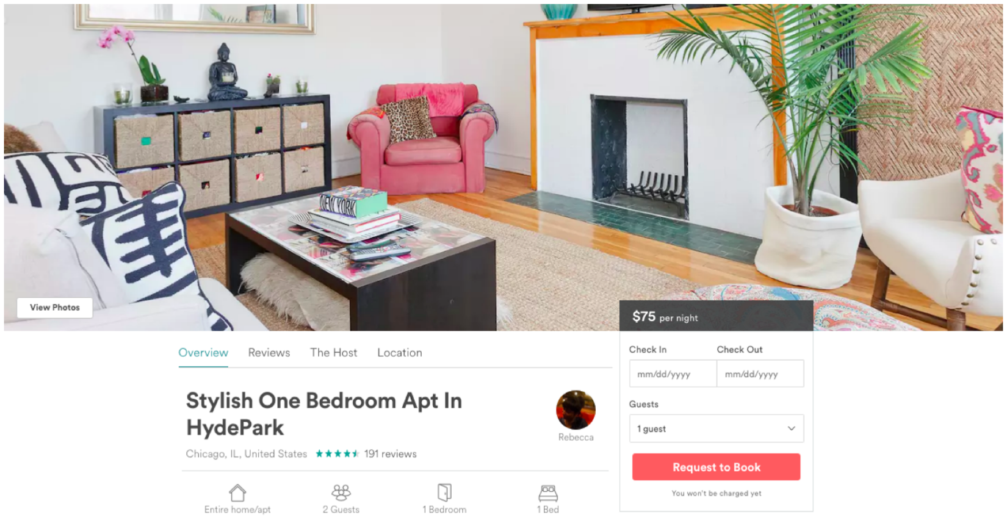

Fig. 1. Sample listing page.

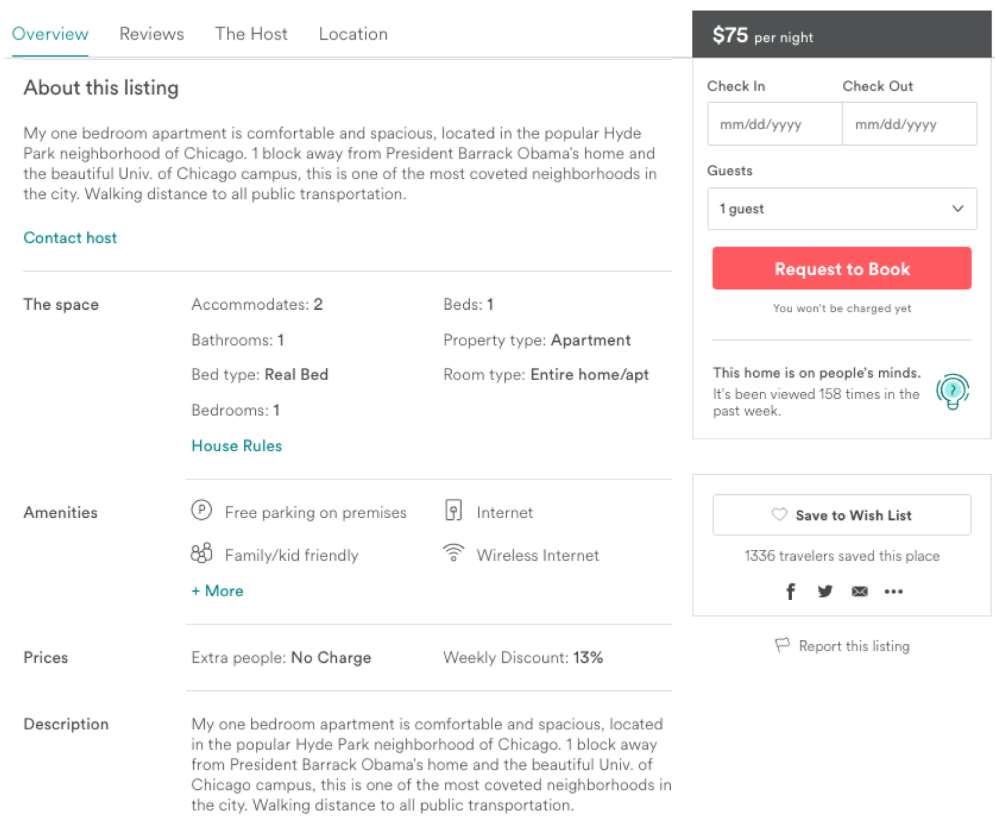

Fig. 2. Listing information.

peer accommodation has led many travelers to opt for Airbnb over a hotel stay under the perception that a number of Airbnb attributes outperform those of traditional hotels (Guttentag and Smith, 2017). This pattern has become increasingly prevalent in popular tourist cities, where Airbnb offers cheaper lodging closer to city centers than hotels (Gutierrez et al., 2017). However, the influx of tourists into residential areas and lack of industry regulation on Airbnb are quickly becoming causes of economic and social concern. Leong and Belzer (2017) investigate the legal ramifications of not subjecting Airbnb to the same regulations as similar establishments. They note that the existing policies prohibiting public accommodations like hotels, restaurants, taxis, and retail businesses from discriminating against customers on the basis of characteristics such as race or religion have not yet evolved to apply to Airbnb, which exposes the platform to the risk of intentional bias and discrimination (Leong and Belzer, 2017).

1.2.5. Research on discrimination in Airbnb

This study builds primarily on research done by Edelman and Luca (2014), Wang et al. (2015), and Kakar et al. (2018).

Edelman and Luca (2014) was the first study to explore the effect of Airbnb host race on the price of a listing using a sample of 3800 New York City hosts. They use Amazon Mechanical Turk workers to identify the race of the host in Airbnb profile pictures. Controlling for a listing's location, its rating on various dimensions, and several variables related to the size of the listing, the researchers find evidence that non-Black hosts on Airbnb charge prices roughly 12% higher than Black hosts. Since their study was conducted using data from 2012, when Airbnb was relatively new, the authors were able to include all of Airbnb's New York City listings in their sample. Edelman and Luca (2014) considered only Black hosts. They did not differentiate impact by gender, and were only able to control for a few property characteristics.

Given these constraints, Wang et al. (2015) and Kakar et al. (2018) build on Edelman and Luca (2014)'s approach and apply it to study discrimination in the context of the San Francisco Bay Area.

Wang et al. (2015) analyzed the impact of Asian host race on Airbnb prices in Oakland and Berkeley by scraping information on 100 hosts from the Airbnb website in those areas. They found that when controlling for the number of bedrooms, bathrooms, and maximum occupancy, Asian hosts earned $90 (20%) less than White hosts with similar rentals.

The fact that a price disparity is estimated for both Black and Asian hosts, albeit in different contexts, is potentially informative about the mechanism behind discrimination in this context. The Black-White wealth gap is well-documented, large, and pervasive (Asante-Muhammed et al., 2016; Heywood and Parent, 2012; Bayard et al., 1999; Mason, 1997; Chiteji, 2010; Altonji et al., 2005; Gittleman and Wolff, 2004). When considering discrimination against Black hosts, it is unclear if guests discriminate because they expect listings of lower quality (statistical discrimination) or if they simply do not want to stay with Black hosts (taste-based discrimination). By contrast, Asians in the United States have the highest incomes out of any ethnicity in the US (U.S. Census Bureau, 2017). It is therefore unlikely that discrimination against Asian hosts would be statistical. In fact, some studies have found evidence that discrimination against Black people in rental markets is statistical in nature, as landlords use race as a proxy for income. One study found that for Blacks who imply that they are of a higher social class when applying for an apartment, discrimination is virtually not present (Hanson and Hawley, 2011).

Kakar et al. (2018) conduct a similar analysis for hosts in San Francisco as (Wang et al., 2015) in Oakland, and (Edelman and Luca, 2014) in New York City. They code the race, gender, self-identified sexual orientation, and whether the host is a couple from the host profile pictures. Since each picture required manual coding, they identify only 800 out of 6000 active listings in San Francisco as of 2015. The researchers find that Asian hosts charge 8% lower prices relative to White hosts for comparable listings, controlling for neighborhood property values from Trulia, area demographics from the Census, and occupancy rates purchased from a third party as proxies for desirability or attractiveness of the locations. The price disparity is 10% for Hispanic hosts, but becomes insignificant when adding a control for occupancy rates. The price discrepancy for Asian hosts is persistent at 8%.

Several audit studies have also examined discrimination on the other side of the market – those who demand, rather than supply, listings on Airbnb. These studies follow the canonical model of the (Bertrand and Mullainathan, 2004) study, which sent identical resumes to employers, varying only whether the names were distinctively White or Black.

Cui et al. (2016) conducted an audit study to measure discrimination against Airbnb guests. They created fake guest accounts with identical profiles but with different names, either distinctively White or distinctively Black. In the first round of their experiment, they sent out requests for accommodation from a set of accounts with no reviews, as well as a set of accounts with one positive review. In the second round, the accounts were modified to have one negative review. They found that guests with White-sounding names were accepted on average 19 percentage points more often than those with African American-sounding names when the guest accounts had no reviews, but that the presence of a single review, whether positive or negative, rendered the acceptance rates statistically indistinguishable (Cui et al., 2016).

In a more recent audit study, Edelman et al. (2017) measured discrimination against Airbnb guests. They created fake guest accounts that differed only by name and inquired about the availability of listings across five cities. They found that requests for reservations by guests with distinctively Black names were 16% less likely to be accepted by hosts than identical guests with distinctively White names. The estimated cost to the median host who rejects a guest on the basis of race was a loss of between $65–$100 of revenue (Edelman et al., 2017).

2. Data

I rely on data from three sources: Inside Airbnb, the US Census, and the results of two machine learning algorithms.

The main data are taken from Inside Airbnb, a project unaffiliated with Airbnb which aggregates cleaned data on Airbnb listings in over 40 cities across the world (Cox, 2017). The data on Inside Airbnb are sourced from a webscrape of publicly available information on the Airbnb website.[5] The scrape of the Airbnb website was conducted throughout 2015 and 2016, and provides a point-in-time snapshot of all of the listings available in a particular city. This includes all of the information that would be available to a potential Airbnb guest, including its price, data on the location of the listing, its property characteristics, and the characteristics of its host.[6]

The location data includes the city, neighborhood, and zip code of each listing.[7] The property data includes all of the characteristics available on the listing page, which can be seen in Fig. 2. These include the property type and room type of the listing, the number of guests the listing accommodates, the number of bathrooms, bedrooms, and beds, the bed type, the amenities, the number of minimum nights, any extra fees, whether the listing can be booked without host approval, and the cancellation policy.

The host data includes all available host information, as seen in Fig. 3. This includes the host's name and profile picture, their response rate and the response time, whether the host is a Superhost, whether their identity is verified by Airbnb, and if the host requires a guest's profile picture or phone to book.

I restrict the sample to hosts who have profile pictures (there are only 2 hosts who do not). I drop extremely high-priced listings, defined as those priced over $800 per night. Lastly, I restrict my sample to hosts who manage fewer than 20 listings to avoid introducing hostels, professional property managers, and other types of businesses into my sample. To see results without this restriction, see Table 11.

2.1. Coding

Airbnb does not provide the demographic information of their hosts,

[5] Airbnb's host profiles and listings are publicly available information, and no private data was accessed in the scrape. The cleaned data is under a Creative Commons Public Domain Dedication.

[6] Inside Airbnb provides some time-series information on prices, but since each listing's price was not scraped daily, there are often week-long or month-long gaps in the time-series price data. A cursory glance at the time-series prices reveals that hosts do not change prices often, and if they do, they often reflect predictable weekend or holiday seasonality. There is therefore reason to believe that the prices posted at the time of the scrape are representative of a listing's price throughout the year. Because of the incompleteness of the time-series data set, I focus on the cross-sectional data for the main analysis.

[7] The data set does not include Airbnb's original neighborhood designations "due to inaccuracies". Instead, Inside Airbnb assigned neighborhoods to each listing by comparing the geographic coordinates of the listing with each city's neighborhood designations. Location information for listings is anonymized by Airbnb, and no exact address is provided for any listing. The location for a listing could be 0–150 m from the actual address. Fig. 4 presents a map of Chicago's neighborhoods to give an example of the granularity of the neighborhood controls.

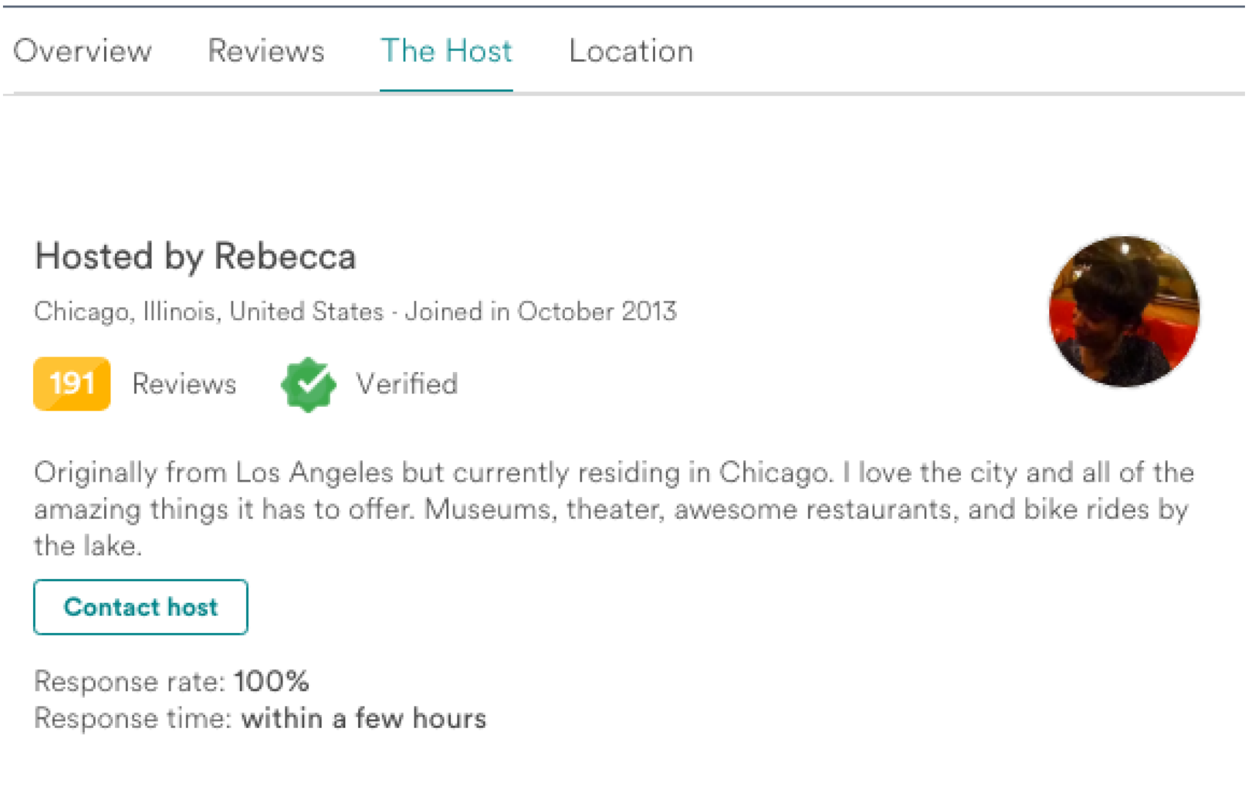

Fig. 3. Host information.

so undergraduate research assistants manually coded the hosts' demographic information. For each listing, assistants identified the race, gender, and age of the host, as well as whether or not there was more than one person in the picture. Only listings with a single person in the profile picture who were identifiably White, Black, Asian, or Hispanic, were included in the main analysis. Listings with couples, groups, children, pictures without a human face, or people of ambiguous race were dropped from the main analysis. Listings that no longer existed at the time of coding were also excluded.[8] After these restrictions I am left with 45,000 observations.

Each research assistant was compensated based on the quantity of the listings they coded. Since this compensation scheme could disincentivize coding accurately, a simple double-checking process was put in place to check codings. Research assistants flagged listings whose picture was ambiguous on any of the dimensions of race, sex, or age. I subsequently coded each flagged listing to check their work.[9]

A total of 70,000 host pictures across seven US cities were coded – Chicago, Los Angeles, New York City, Austin, Washington, D.C., and New Orleans. This sample represents large, racially diverse cities which are geographically dispersed across the United States. For every city but New York, every single Airbnb listing that existed in that city at the time of the scrape was coded. In New York, which had the most listings in the sample, half of the existing 40,000 listings were randomly chosen.[10]

In addition to hosts, research assistants also identified the race, gender, and age of 16,000 reviewers who had stayed at a subset of the listings in Chicago. Chicago's Airbnb market is small enough to be able to code the demographic characteristics of a substantial portion of its Airbnb hosts and their reviewers. A sample size of 16,000 represents about a quarter of the total number of Airbnb guests in my Chicago data. Thus, Chicago was chosen partially because of personpower constraints – since there are on average twenty reviewers for every host, coding all of the reviewers even for a single city remained infeasible. As the third largest city in the US, the results for Chicago are still reasonably generalizable to other major metropolitan areas in the United States. Importantly, Chicago is also racially diverse enough to obtain adequate sample sizes for my analysis. Since reviewer-side analysis requires variation not just in the host race, but also in the reviewer race, conducting the analysis in a smaller but more racially homogenous city would not have provided sufficient statistical power to test hypotheses.

2.2. Census data

In addition to this main data, I merge in zip code-level Census data from the American Community Survey (2016, 5-year estimates). The Census data includes demographic data such as the racial composition of each zip code.

Included are various measures of economic health: the median income, unemployment rate for the total labor force, share of households on Supplemental Security Income, median housing value, and gross

[8] If certain groups of hosts systematically exited the Airbnb market between the time of the scrape and the time of the coding, dropping those listings could bias the results. Unfortunately, there is no way to verify the demographics of the hosts who dropped out, since Airbnb takes down their profile picture.

[9] It is important to note that the coding need not reflect the actual demographics of the host. Rather, it is sufficient that they are coded with the race, sex, and age that the average guest on Airbnb would assume after looking at the profile picture. Since the average University of Chicago undergraduate might not be representative of the average guest on Airbnb, in future research it would be preferable for two people to code each picture, and a third person to resolve disagreements.

[10] This approach limits the applicability of my findings to urban areas, discounting the roughly one-fifth of Airbnb's listings which are located in rural areas. A 2017 report released by Airbnb stated that roughly a fifth of all active listings are located in rural areas, with 138% year-in-year growth in Airbnb guest arrivals at rural listings.

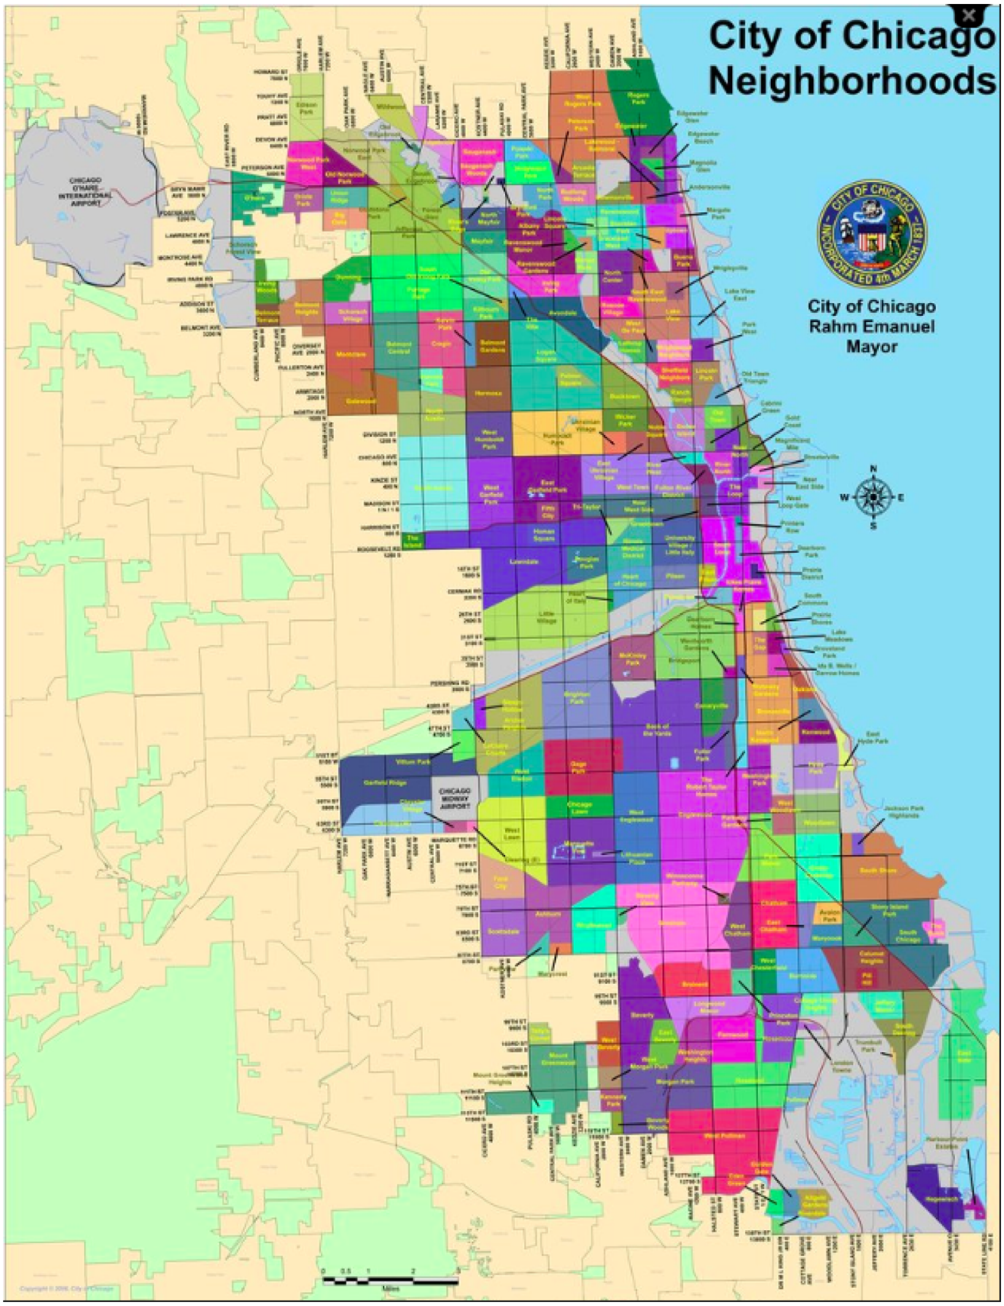

Fig. 4. City of Chicago neighborhoods, showing level of granularity of neighborhood controls.

rent. I add proxies for how far the zip code is from downtown, such as its population density and the commuting time to work. Importantly for this analysis, I have data on the occupancy rate, which I later use to control for variation in the availability of a listing. I standardize all variables, so each zip code has z–scores that are relative to the mean in that city.

2.3. Sentiment analysis

A third source of data are the host effort measures I construct that are not directly observable from the listing page. I use a machine learning algorithm called TextBlob, developed for Python, to analyze the descriptions hosts write of their listings (Loria, 2019). There are several host-written fields on each listing page, the summary, description, space, neighborhood overview, transit, and notes. By filling out these fields, hosts describe their listing and have the opportunity to provide guests with helpful tips and information about the surrounding area. These descriptions can signal host effort, and therefore a listing's potential quality. This signal is important since short, unhelpful, or simply unappealing descriptions might contribute to decreased demand. To capture variation in the quality of these fields, I use natural language processing to assign each of the six fields two scores: subjectivity and polarity. The subjectivity score measures to what extent the text includes words like "I believe" and "I think" rather than more objective sentences, such as "The decoration is contemporary". The polarity score measures how positive or negative the text is in

Table 1
Summary statistics by host race: listing characteristics.

	Regression sample					
	Full data	All	White	Black	Hispanic	Asian
Outcome variables						
Log price	4.81	4.73	4.79	4.51	4.65	4.55
	(0.75)	(0.66)	(0.66)	(0.62)	(0.65)	(0.64)
Log number of reviews	2.20	2.16	2.18	2.14	2.16	2.04
	(1.39)	(1.38)	(1.39)	(1.36)	(1.40)	(1.36)
Covariates						
Property type						
Apartments/Lofts	0.60	0.63	0.63	0.66	0.66	0.62
Townhouses/ Condos	0.04	0.04	0.04	0.04	0.04	0.06
Houses	0.32	0.30	0.30	0.27	0.26	0.30
Others	0.04	0.03	0.03	0.03	0.03	0.03
Room type						
Entire house/ Apartment	0.58	0.55	0.58	0.43	0.51	0.42
Private room	0.38	0.42	0.39	0.50	0.44	0.53
Shared room	0.04	0.04	0.03	0.07	0.05	0.05
Max num. guests	3.44	3.15	3.24	2.94	3.06	2.84
	(2.41)	(2.13)	(2.15)	(2.06)	(2.15)	(2.00)
Bedrooms	1.34	1.26	1.28	1.19	1.21	1.18
	(0.92)	(0.80)	(0.83)	(0.69)	(0.78)	(0.72)
Bathrooms	1.30	1.23	1.24	1.19	1.21	1.19
	(0.69)	(0.55)	(0.56)	(0.49)	(0.52)	(0.53)
Beds	1.82	1.67	1.69	1.60	1.68	1.57
	(1.41)	(1.21)	(1.19)	(1.15)	(1.51)	(1.18)
Cleaning fee	48.94	43.70	46.06	36.20	40.35	36.45
	(59.62)	(48.32)	(49.73)	(43.18)	(45.51)	(42.86)
Extra guests charge	13.74	13.43	13.26	15.13	13.94	12.72
	(23.65)	(22.67)	(23.00)	(22.71)	(22.48)	(20.36)
Minimum nights	3.01	3.03	3.08	2.61	2.86	3.17
	(9.21)	(8.79)	(9.39)	(4.35)	(6.55)	(8.67)
Availability (out of 30 days)	11.54	11.04	10.64	14.19	11.24	10.79
	(10.93)	(10.91)	(10.75)	(11.49)	(10.94)	(11.01)
Number of amenities	0.81	0.79	0.80	0.75	0.79	0.75
	(1.10)	(1.10)	(1.10)	(1.04)	(1.10)	(1.13)
Instantly bookable?	0.15	0.15	0.14	0.21	0.17	0.16
	(0.36)	(0.36)	(0.34)	(0.41)	(0.38)	(0.37)
Year of first review	14.86	14.86	14.83	14.89	14.90	15.03
	(1.22)	(1.22)	(1.22)	(1.30)	(1.21)	(1.17)
Strict cancellation policy	0.43	0.40	0.42	0.42	0.42	0.42
Observations	69010	45076	32934	4354	2913	4875

Note: The values in the table are means and standard deviations of listing-level data in my full sample. Summary statistics for selected covariates are listed in the table. Categorical variables such as room type do not have standard deviations. Property types are explicitly listed if more than 1.5% of listings are that type. Only the most popular cancellation policy type is listed - in the full sample, 99% of listings have strict (43%), flexible (31%) or moderate (25%) cancellation policies. Year of first review is a proxy for the time on the market - 14.86 indicates that the first review of the mean listing in the full sample occurred in October of 2014.

sentiment. Together, these metrics capture variation in how positively the listing is described, and how much ‘flowery’ language is used rather than objective and perhaps more useful language.

For the 16,000 reviewers who left reviews in Chicago listings, I use a second machine learning algorithm called Sentimentr to calculate the sentiment of each review and the mean sentiment of the listing (Rinker, 2019). The algorithm uses a dictionary of positive and negative words to assign each sentence a sentiment score from −1 to 1. In assigning scores, the algorithm considers the number of good or bad words in a sentence, as well as their valence shifters, or words that affect the sentiment-carrying word in the sentence. For example, the

Table 2
Summary statistics by host race: host demographics.

	Regression sample					
	Full data	All	White	Black	Hispanic	Asian
Race						
White	0.64	0.73	1.00	0.00	0.00	0.00
Black	0.07	0.10	0.00	1.00	0.00	0.00
Hispanic	0.05	0.06	0.00	0.00	1.00	0.00
Asian	0.09	0.11	0.00	0.00	0.00	1.00
Unknown	0.15	0.00	0.00	0.00	0.00	0.00
Sex						
Male	0.31	0.45	0.45	0.40	0.50	0.44
Female	0.38	0.55	0.55	0.60	0.50	0.56
Unknown	0.31	0.00	0.00	0.00	0.00	0.00
Age						
Young (< 30)	0.43	0.51	0.49	0.54	0.52	0.61
Middle-aged	0.42	0.46	0.48	0.45	0.47	0.38
Old (> 65)	0.02	0.02	0.03	0.00	0.01	0.01
Unknown	0.13	0.00	0.00	0.00	0.00	0.00
Observations	69000	45076	32934	4354	2913	4875

Note: The values in the table are summaries of host demographics in the host-level data. Column 1 is the summary statistics for the full, unrestricted data set across 7 cities. Columns 2–6 are the restricted data used in the analysis. Column 2 is the full regression sample, and Columns 3–6 break down the regression sample by host race. The “Unknown” category was dropped from the regression and is therefore zero throughout Columns 2–6. White refers only to non-Hispanic Whites.

Table 3
Summary statistics ny host race: host characteristics.

	Regression sample					
	Full data	All	White	Black	Hispanic	Asian
Outcome variables						
Host listings count	6.38	2.23	2.16	2.38	2.49	2.44
	(36.54)	(2.59)	(2.50)	(2.83)	(3.03)	(2.61)
Covariates						
Review scores rating	93.56	93.68	94.18	91.91	92.80	92.26
	(8.13)	(7.90)	(7.33)	(9.44)	(8.71)	(9.27)
Host is a Superhost	0.13	0.13	0.14	0.09	0.11	0.10
	(0.34)	(0.33)	(0.34)	(0.28)	(0.31)	(0.30)
Response rate	0.77	0.76	0.76	0.78	0.76	0.74
	(0.38)	(0.39)	(0.39)	(0.37)	(0.39)	(0.40)
Acceptance rate	0.47	0.45	0.46	0.35	0.49	0.44
	(0.46)	(0.46)	(0.46)	(0.45)	(0.47)	(0.47)
Polarity of summary	0.30	0.30	0.30	0.29	0.30	0.29
	(0.17)	(0.17)	(0.17)	(0.16)	(0.17)	(0.17)
Subjectivity of summary	0.53	0.54	0.54	0.53	0.54	0.53
	(0.15)	(0.15)	(0.15)	(0.15)	(0.15)	(0.15)
Host’s identity verified?	0.70	0.70	0.71	0.66	0.68	0.69
	(0.46)	(0.46)	(0.45)	(0.47)	(0.47)	(0.46)
Guest pic required?	0.04	0.04	0.04	0.06	0.04	0.04
	(0.19)	(0.19)	(0.19)	(0.23)	(0.19)	(0.19)
Guest phone required?	0.05	0.05	0.05	0.06	0.04	0.04
	(0.22)	(0.21)	(0.21)	(0.24)	(0.20)	(0.20)
Response time < 1 h	0.41	0.40	0.39	0.44	0.41	0.41
Observations	69,010	45,076	32,934	4354	2913	4875

Note: The values in the table are means and standard deviations of host-level data in the full sample. Summary statistics for selected covariates are listed in the table. Categorical variables such as response time do not have standard deviations. Statistics for only the most frequent response time (“within an hour”) are included. White refers only to non-Hispanic whites. Polarity of “Summary” and subjectivity of “Summary” refer to the scores from a natural language processing algorithm that measures the sentiment and objectivity of that field. These two measures were also calculated for the description, space, neighborhood overview, notes, and transit fields, but were not included in the table for the sake of clarity and because they follow a similar pattern as the “Summary” field.

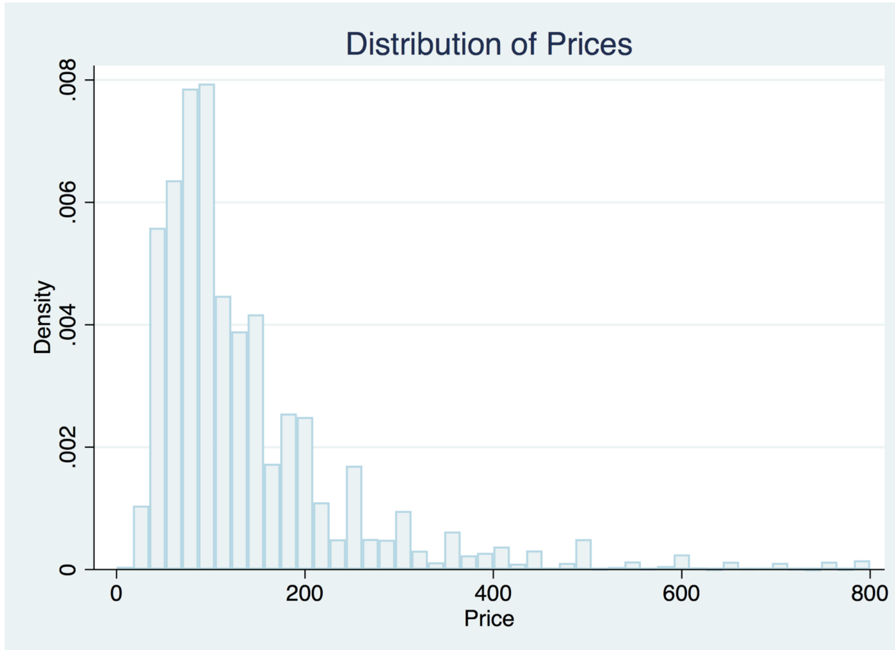

Fig. 5. Distribution of listing prices in the sample.

algorithm assigns "I like the listing," "I *really* like the listing," and "I like the listing, *but*... " different scores because of the presence of valence-shifting words like "really" and "but". One limitation of conducting sentiment analysis is that not every sentence that a human would consider bad or good carries a sentiment word that the algorithm recognizes. For example, "The apartment had cockroaches" is certainly a horrible review, but would be given a score of 0 because it contains no emotion-laden words.

2.4. Data summary

Summary statistics of listing characteristics, host demographics, and host characteristics are displayed in Tables 1–3. Histograms of price, number of reviews, and review sentiment are included in Figs. 5–7. There is significant variation in both sex and race of the hosts on Airbnb. Roughly a third of the sample are single females (38%), and a third are single males (31%), and the rest are couples or groups (31%).

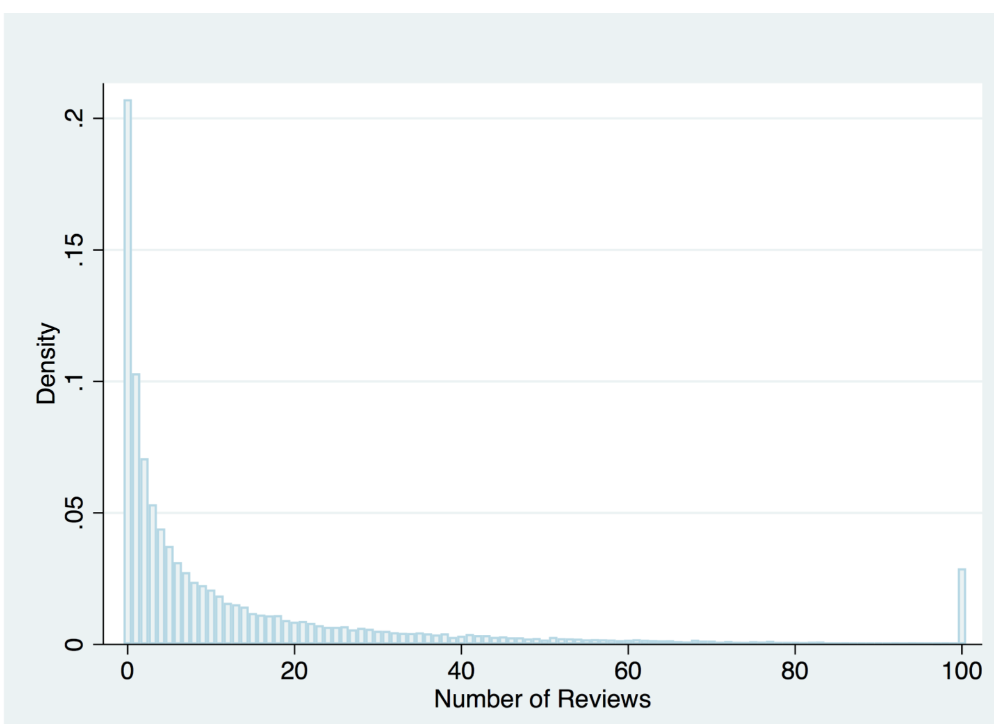

Fig. 6. Distribution of the number of reviews per listing.

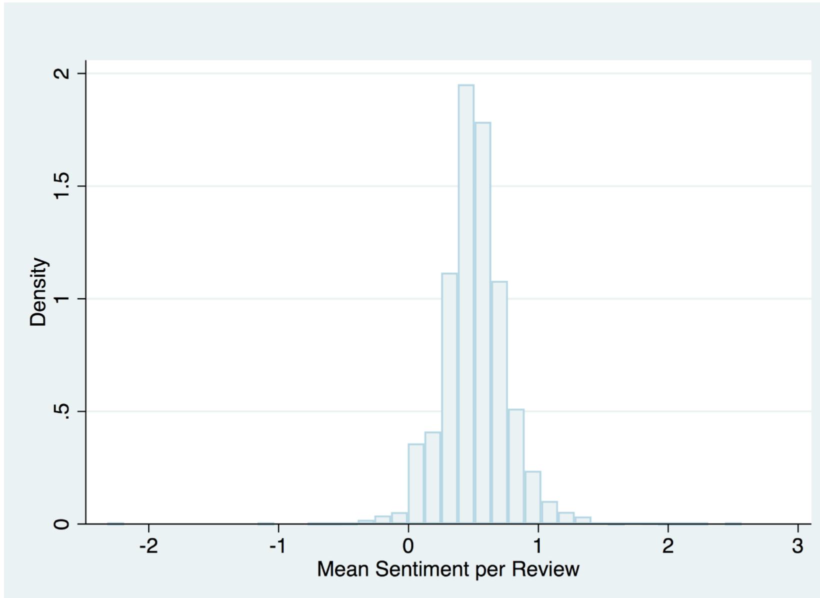

Fig. 7. Distribution of average review sentiment per listing.

About two-thirds of the hosts are White (64%), and less than a tenth are Black (7%), Hispanic (5%), or Asian (9%).[11]

The prices of listings owned by White hosts are higher than those of other hosts. The mean price per night of a listing is $178 per night for White, $125 for Black, $160 for Hispanic, and $131 for Asian hosts. Minority hosts also have lower median prices and lower standard deviations, indicating that not only do minority hosts own cheaper listings on average, but their listings are more concentrated around the lower mean.[12]

Consistent with previous research, differences in property characteristics explain a substantial part of the price disparity. Table 1 shows that White hosts own the most costly property type (houses) and the fewest cheaper property types (apartments or lofts), and their properties have the most features correlated with higher prices (such as more bedrooms, bathrooms, beds). Though White hosts' listings are better on every single metric of immutable property quality, this is not the case for host characteristics. Minority hosts do well in categories where the host can personally influence their desirability. For example, Black hosts have the fastest response time, fewest minimum nights restrictions, make their listing available for more days, and make their property available to book instantly than any other group.

Reviewer characteristics are displayed in Table 4. The reviewers have similar gender diversity as the overall host population but significantly less racial diversity. The measure of review quality assigned by the machine learning algorithm to the review text correlates well with the numeric review score. While all hosts have on average very positive reviews, White hosts have the most positive, and Black hosts the least positive, review sentiment.

Table 4
Summary statistics by race: reviewer characteristics.

		Reviewer race in Chicago data				
	Full data	All	White	Black	Hispanic	Asian
Reviewer race	1.00	1.00	0.66	0.03	0.04	0.11
Host race						
White	0.73	0.83	0.84	0.70	0.75	0.75
Black	0.06	0.06	0.05	0.17	0.07	0.06
Hispanic	0.04	0.05	0.05	0.06	0.10	0.08
Asian	0.05	0.05	0.05	0.08	0.08	0.11
Unknown	0.12	0.00	0.00	0.00	0.00	0.00
Review sentiment	0.51	0.51	0.51	0.50	0.47	0.53
	(0.26)	(0.26)	(0.25)	(0.23)	(0.30)	(0.25)
Listing sentiment	0.51	0.51	0.51	0.50	0.50	0.51
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.09)
Observations	17,050	10,573	6929	319	402	1153

Note: The values in this table are means and standard deviations of reviewer-level data who left reviews for a randomly chosen set of hosts in Chicago. Column 1 presents the means for the full reviewer data. Column 2 presents the means of the sample used in Table 7. Columns 3–6 partition Column 2 by reviewer race. Row 1, "reviewer race" indicates the proportion of the different reviewer races in the data coded. Row 2, "host race" indicates the marginal probability of a host race given a reviewer race. The review sentiment is the sentiment of each review, the listing sentiment is the average sentiment per listing. Observations in Columns 2–5 do not add up to 17,050 because multiracial or unidentifiable reviewer pictures are excluded. White refers only to non-Hispanic Whites.

3. Empirical approach

I use OLS to estimate the impact of host race on listing price on Airbnb, controlling for location, property, and host characteristics. My main specification is of the form:

$$log\,(\text{Price}_{i,j}) = \alpha_i + \beta_1 \text{Race}_i \times \text{Sex}_i + \beta_2 \text{Age}_i + \beta_3 \text{Location} + \beta_4 x_{i,j}$$

Where $log(\text{Price}_{i,j})$ is the log of host i's price from their Airbnb listing j. For hosts with multiple listings, each listing is a separate observation. $\text{Race}_i \times \text{Sex}_i$ is the interaction of the race and sex of the host, where White males are the omitted category throughout. Age_i is a dummy variable for whether the host is young, middle-aged, or senior. $x_{i,j}$ is a vector of property and host controls that grows additively in each model. Together, I control for all features of the listing that are available to a potential guest, as well as additional metrics that aim to capture unobservable differences between hosts. Each column of Table 5 controls for all covariates in the previous columns, plus a new set of covariates, as detailed below. All other tables control only for the full specification, Model 4. Standard errors are clustered by neighborhood throughout.

Table 5 presents OLS estimates of the effect of host race and gender

[11] The rest of the profile pictures were pictures of groups, pictures without a human face, or multiracial couples.

[12] The median price of a listing owned by a White hosts is $115 per night, $90 for Black hosts, $99 for Hispanic hosts, and $90 for Asian hosts.

Table 5
Main result: estimates of effect of host race and gender on listing price.

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4
White female	−0.0236*	−0.0138	0.00201	0.00298
	(0.0106)	(0.00854)	(0.00496)	(0.00484)
Black male	−0.276***	−0.0828**	−0.0360**	−0.0328**
	(0.0315)	(0.0259)	(0.0123)	(0.0123)
Black female	−0.299***	−0.0586**	−0.0196	−0.0167
	(0.0296)	(0.0188)	(0.0102)	(0.00996)
Hispanic male	−0.153***	−0.0521**	−0.0233*	−0.0200
	(0.0259)	(0.0191)	(0.0113)	(0.0113)
Hispanic female	−0.150***	−0.0653**	−0.0196	−0.0202
	(0.0280)	(0.0202)	(0.0115)	(0.0114)
Asian male	−0.221***	−0.0987***	−0.0425**	−0.0446***
	(0.0336)	(0.0225)	(0.0134)	(0.0135)
Asian female	−0.283***	−0.131***	−0.0409***	−0.0396***
	(0.0299)	(0.0161)	(0.00874)	(0.00893)
Constant	4.802***	4.979***	3.891***	4.003***
	(0.0300)	(0.398)	(0.343)	(0.344)
Location controls		Yes	Yes	Yes
Property controls			Yes	Yes
Host controls				Yes
Observations	45,073	45,073	45,073	45,073
Adjusted R2	0.0263	0.246	0.716	0.720

Standard errors in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table presents the impact of host race on the price of a listing. The dependent variable is the log price. The omitted category is White males. The unit of observation is a listing. The sample is listings across seven US cities, whose prices are no more than $800 per night, and whose hosts own no more than twenty properties. Model 1 is the baseline effect of host demographics on price. Model 2 includes fixed effects for the neighborhood of the listing and Census demographic, economic health characteristics, and occupancy rates on the zipcode-level. Model 3 adds listing characteristics such as the property type and size. Model 4 adds host characteristics such as response and acceptance rates, and measures of host effort.

on the listing price according to the following models.

1. *Model 1* presents the raw effect of host race and sex on the price of a listing. These coefficients are consistent with the mean listing prices by race presented in Table 1, but further broken down by male and female hosts within each racial category.
2. *Model 2* adds city and neighborhood fixed effects, as well as zip code-level Census information on demographics and various indicators of economic health.
   Adding neighborhood fixed effects removes any variation in prices that are due to a property's location in a neighborhood with better amenities or proximity downtown. However, if Airbnb listings are clustered in certain areas of each city, neighborhood controls might not be very informative. To this end, I add data on the property values, proxies for the desirability of the neighborhood, including the unemployment rate, the occupancy rate, and other measures of poverty. See full details in Section 2. I also include population density and commuting time to work as proxies for distance to downtown.
3. *Model 3* adds controls for listing-specific characteristics, as detailed in Section 2. See Table 1 for a full list of property controls. I also control for the listing's duration on the market by proxying with fixed effects for the month and year of the listing's first review.
4. *Model 4* represents my full specification. I add all remaining host-dependent fields on the listing page, such as Superhost status, the host's response time, and their cancellation policy. I also include my constructed host quality controls in the form of sentiment analysis of the text on the listing page to control for hosts who write more objective descriptions of more positive valence. See Table 3 for a full list of these controls and Section 2 for details of construction of host effort variables.

4. Results

4.1. Baseline analysis: estimating the effect of race on listing price

Model 1 in Table 5 is a naive regression and presented solely to show baseline effects of host race and gender on price. Model 2 adds the property's location. Since a listing's price is strongly correlated with its location, it is unsurprising that a large amount of variation in Airbnb prices between racial groups can be explained by location controls. This is consistent with the idea that individuals of a particular race are not randomly distributed across neighborhoods of a given city for a variety of reasons, including persistent urban segregation or in-group preferences. For example, Black hosts had the largest drop in point estimate with the addition of location controls, and it is well-documented that Blacks in urban populations are nearly four times more likely than Whites to live in neighborhoods where the poverty rate is 40% or higher (Firebaugh and Farrell, 2016).

Moving to Model 3, property characteristics have the most explanatory power in accounting for price differences, as the R^2 increases from around 0.25–0.72 with the addition of listing controls. As in all rental markets, the type of property, the amount of rooms a guest is renting, and other property size characteristics are all significant drivers of price differences. In Model 4, the majority of my results are stable to the addition of host characteristics while still clustering standard errors at the neighborhood level. However, the inclusion of host characteristics does wipe out a significant price disparity for Black females and Hispanic males. Since they do not improve the fit of the model substantially, increasing the R^2 by only 0.004, it is unlikely that adding even more host quality controls would explain significant variation outside of an experimental setting.

The final estimates in Model 4 indicate that Asian hosts, both male and female, earn lower prices from their Airbnb listing than White male hosts. These effects are 4% for Asian females and 4.5% for Asian male hosts. The second biggest effect is for Black males, whose prices are roughly 3.5% lower.[13] The estimates are negative but not statistically significant for Hispanic and Black female hosts. Moving from left to right in the table, the estimated coefficients on host race become less negative. This raises the question of whether or not a more saturated model that includes more covariates would further reduce the estimates.

I follow the logic of Oster (2017) and Altonji et al. (2005) to explore this question. I use the results from Models 3 and 4 in my calculations. The change in estimates from Models 1 to 2 is largely not relevant to the question of whether there is a price disparity on Airbnb, as the location controls account for a different kind of variation – the fact rental prices are highly correlated with location, and cities are highly segregated by race. I therefore skip Model 1 and use Models 2 and 3 as my comparison. Under the most extreme assumption that all of the variation in the outcome variable can be attributed to unobservables, I can no longer reject the Null hypothesis of no price disparity. However, as Oster (2017) notes, this is a very high standard. In her review of papers published in the American Economic Review, Journal of Political Economy, and the Quarterly Journal of Economics, only 40% of the positive results for non-randomized studies survive this standard, and 30% of randomized results.

Using the change in coefficients going from Models 2 to 3, I lose statistical significance at Oster's maximum R^2 of 0.95. It is worth noting that the pattern of decreasing coefficients observed in my paper is remarkably similar to (Kakar et al., 2018). For example, comparing Columns 3–6 in Table 3 of their paper, their point estimate implied by Oster's analysis is 0.

These effects of Asian host race on listing price are lower than

[13] This effect is statistically significant at the $p < .001$ level for Asian female and male hosts, and the $p < .01$ level for Black male hosts.

Table 6
Effect of host race on two proxies of a listing's number of bookings.

	(1) Number of reviews (log)	(2) Number of vacant days (out of 30)
White female	−0.0584***	−0.888***
	(0.0169)	(0.110)
Black male	−0.0448	2.338***
	(0.0321)	(0.258)
Black female	−0.0929***	1.775***
	(0.0280)	(0.227)
Hispanic male	−0.0525	−0.185
	(0.0320)	(0.272)
Hispanic Female	0.0224	−0.0908
	(0.0368)	(0.274)
Asian male	−0.0148	−0.130
	(0.0329)	(0.234)
Asian female	−0.0567*	−1.106***
	(0.0251)	(0.215)
Constant	4.033***	−11.46
	(0.446)	(11.12)
Location controls	Yes	Yes
Property controls	Yes	Yes
Host controls	Yes	Yes
Observations	35,734	45,076
Adjusted R2	0.549	0.228

Standard errors in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table presents the effect of host race on two proxies for the quantity demanded of a listing: its number of reviews and its availability out of 30 days. The availability metric represents the number of days out of the total days available for booking that a listing is vacant. The omitted category is White males. I control for the specification in Table 5, Model 4.

measured by Wang et al. (2015) (a 20% disparity) and Kakar et al. (2018) (an 8% disparity). This fact is unsurprising, as it is unlikely that we would measure the same disparity when there are important structural differences between this paper and previous literature. My sample is national, while both previous papers considered Airbnb listings in the Bay Area only. There are also vast difference in our sample sizes: Wang et al. (2015) has a sample size of 101 observations, while (Kakar et al., 2018) has 715 observations. The effects of Black host race that I measure are smaller than those of Edelman and Luca (2014), who measured a price disparity of 12%. In Table 10, I confirm their 12% result using my data and their controls, evidence that the smaller point estimate is due at least in part to better controls. For similar reasons, however, this comparison is uninformative.

4.2. Using quantity demanded to test hypotheses

I now turn to testing different hypotheses to explain this price disparity. First, I consider whether lower prices are due to a demand shift or a supply shift. Theoretically, estimating the quantity demanded in a market is one way to distinguish between a demand shift versus a supply shift (see Section 1, pages 3 and 4 for a full discussion). In the context of Airbnb, the true measure of a listing's quantity demanded would be its number of bookings. In the absence of this data, I instead consider two different measures of quantity demanded to investigate whether the properties of minority hosts are under-booked or over-booked compared to Whites.

4.2.1. Number of reviews as a proxy for quantity demanded

I regress the number of reviews and the vacancy rate on host race, controlling for Model 4 in Table 5. To account for the fact that older listings mechanically have more reviews, I control for the listing's time on the market. The results of this analysis are in Table 6, Column 1. I find that minority hosts have either the same or lower review numbers than White hosts for a listing of the same age. The significant effects are for the women – Black females have 10% lower review numbers, and Asian and White female hosts 5% lower review numbers, than White males. Coefficients are imprecisely estimated zeros, or negative and approaching significance, for all other hosts.

4.2.2. Vacancy rate as a proxy for quantity demanded

One worry is that minority hosts may have fewer reviews because they list their property as available for booking fewer days of the month, not because they face lower demand. One way to address this concern is to consider the number of days the property was available over the last month. The host can mark certain days as available for bookings on their listing page. Potential guests can then see on which days the listing is available and book accordingly. When a guest books a day, that day is removed from the availability calendar. Therefore, the availability out of 30 days is a measure of a listing's vacancy rate. Throughout the analysis, I control for the occupancy rate in the listing's zip code. A unit is considered vacant by the Census if the owners have a usual residence somewhere else. These controls help account for race-based differences in variables which impact vacancy rates.

The results of the regression of this vacancy rate on host race are presented in Table 6, Column 2. I find that the listings of Black hosts spend about 20% more time vacant on the market than the listings of White males. The effect is statistically significant, and amounts to about 2–3 days per month in real units. Contrary to Black hosts, the vacancy rate of White and Asian women is actually lower than White male hosts by 1 day per month. These findings tell a nuanced story. For Black hosts, both measures of quantity demanded in Table 6 suggest that even though Black hosts offer their listings for more days and charge lower prices, fewer guests stay with them. Female Asian and female White hosts, on the other hand, have lower vacancies than White hosts. Lower availability is therefore a possible explanation for why these groups have a lower number of reviews.

The availability of a listing is the result of booking and endogenous choices by hosts of whether to make their listing available. If there is no systemic difference across racial groups in the endogenous component, then a regression of availability on host race can shed light on whether race affects consumer demand for listings. However, there may be good reason to believe that there are supply-side differences by race in the propensity of hosts to make their listing available. For example, if White hosts are less likely to live in the properties that they list, perhaps because they own a second property, the availability of their listings may be high regardless of their actual demand. This would bias any coefficients on minority host race downwards relative to White hosts. The higher vacancy rate for Black hosts and lower vacancy rate for White and Asian female hosts are therefore likely a lower bound for the true vacancy rates. A bias of this kind would therefore not be problematic for my analysis, as it would mean that I underestimate the gap in availability between White hosts and minority hosts.

4.2.3. Review quality of minority hosts

Good reviews are essential to establishing the credibility of a host on Airbnb, as well as for transacting in the wider P2P market. Previous analyses, including Edelman and Luca (2014), control for the numeric review score of the listing as a proxy for listing quality. However, since there is little variation in the numeric review score, these measures could be uninformative for potential guests in inferring listing quality.[14] For this reason I use review text instead of the numeric score in my analysis.[15] I conduct sentiment analysis on those reviews to give each review a polarity and subjectivity score, allowing me to include controls

[14] This is the case for most online marketplaces. Fradkin et al. (2017) study the determinants of review informativeness on Airbnb and find that most reviews, both numeric and text, are positive. However, the written reviews tend to reflect real experience of the user.

[15] A low share of guests who review may be a more accurate proxy for low quality, because many users prefer to leave no review rather than a negative review. Review share information, however, is not available.

Table 7
Estimates of effect of host demographics on review sentiment, by reviewer demographics.

	Reviewers								
	(1) Full sample	(2) White M	(3) White F	(4) Black M	(5) Black F	(6) Hispanic M	(7) Hispanic F	(8) Asian M	(9) Asian F
White female	−0.00268	−0.0246	0.0640	−0.0718	2.215***	−1.032***	−0.898***	4.660***	−1.823***
	(0.0468)	(0.0758)	(0.0575)	(0.0902)	(2.03e−11)	(2.01e−10)	(2.73e−13)	(0.416)	(2.77e−08)
Black male	−0.172**	−0.176	−0.155	−0.168	−15.79***	22.87***	0.186***	−4.616	−18.00***
	(0.0621)	(0.192)	(0.312)	(0.459)	(1.20e−10)	(1.80e−09)	(5.36e−13)	(3.232)	(0.000000824)
Black female	0.104	−0.0509	0.0863	0.144	−2.253***	6.886***	0.345***	−4.269***	3.430***
	(0.0685)	(0.185)	(0.126)	(0.195)	(3.76e−11)	(4.07e−10)	(1.60e−12)	(0.883)	(0.000000275)
Hispanic male	−0.115**	−0.0477	−0.0130	−0.436***	6.431***	19.07***	−0.512***	−7.871***	6.453***
	(0.0411)	(0.0967)	(0.108)	(0.113)	(6.66e−11)	(1.11e−09)	(8.00e−13)	(1.597)	(0.000000210)
Hispanic female	0.0711	0.00719	0.0117	0.0858	−37.98***	85.18***	−2.929***	6.073***	−4.928***
	(0.0951)	(0.401)	(0.139)	(0.187)	(2.99e−10)	(4.27e−09)	(8.01e−13)	(0.888)	(0.000000109)
Asian male	0.0219	−0.281	−0.168	0.182	6.200***	−21.97***	0.792***	8.107***	11.59***
	(0.162)	(0.231)	(0.141)	(0.271)	(1.05e−10)	(1.28e−09)	(8.01e−13)	(0.736)	(0.000000331)
Asian female	−0.147	−0.224	−0.321	−0.0774	−8.758***	−11.16***	−0.993***	7.884***	−2.325***
	(0.0882)	(0.201)	(0.162)	(0.307)	(6.98e−11)	(7.24e−10)	(1.60e−12)	(0.797)	(8.11e−08)
Location controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Host controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,573	2665	2527	1737	121	171	27	198	142
Adjusted R2	0.0238	0.0548	0.0525	0.0922	0.838	0.719	0.970	0.642	0.786

Standard errors in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table presents the quality of reviews that reviewers leave for hosts in Chicago. The columns are the demographics of the reviewers (male is "M", female is "F"), and the rows are the demographics of the host, consistent with previous tables. The outcome variable is the standardized sentiment of the review, as assigned by a machine learning algorithm. Reviews that are numerically positive are of positive sentiment and numerically negative are negative sentiment, relative to the mean sentiment score for each host type. The unit of observation is a single review. The data is a subsample of the Chicago hosts and their reviewers. I control for the specification in Table 5, Model 4.

Table 8
Main results by city.

	(1) LA	(2) NYC	(3) Austin	(4) Chicago	(5) New Orleans	(6) DC	(7) Nashville
White female	−0.00890	0.0106	0.0181	−0.0195	0.0299	0.0222	0.0277
	(0.00765)	(0.00662)	(0.0160)	(0.0193)	(0.0208)	(0.0127)	(0.0164)
Black male	−0.0514*	−0.00649	−0.0487	−0.00837	−0.0474	−0.0881*	0.0210
	(0.0212)	(0.0116)	(0.0952)	(0.0404)	(0.0458)	(0.0427)	(0.0750)
Black female	−0.0191	0.0108	−0.0838	−0.0689*	−0.0373	−0.00692	−0.0649
	(0.0152)	(0.0128)	(0.0847)	(0.0284)	(0.0536)	(0.0408)	(0.0694)
Hispanic male	−0.0178	−0.0388	−0.0331	−0.0347	−0.0128	−0.0137	−0.152*
	(0.0143)	(0.0210)	(0.0352)	(0.0339)	(0.0903)	(0.0392)	(0.0683)
Hispanic female	−0.0513**	0.0226	0.0992*	−0.0865**	0.0308	0.0337	−0.104
	(0.0159)	(0.0186)	(0.0488)	(0.0311)	(0.0781)	(0.0433)	(0.0675)
Asian male	−0.0575***	−0.0130	−0.0245	−0.121**	−0.0399	−0.0730*	−0.00731
	(0.0165)	(0.0211)	(0.0502)	(0.0437)	(0.107)	(0.0321)	(0.0824)
Asian female	−0.0171	−0.0379**	−0.197**	−0.110**	0.0535	−0.00646	−0.0647
	(0.0112)	(0.0130)	(0.0656)	(0.0349)	(0.0653)	(0.0208)	(0.0768)
Fixed effects:							
Location controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Host controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	16,824	14,765	3635	3255	2562	2285	1747
Adjusted R2	0.754	0.736	0.721	0.730	0.676	0.675	0.774

Standard errors in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table estimates the main results in Table 5 separately across the 7 cities in the sample. Each set of coefficients represents the coefficient on host race in a regression with log price as the outcome variable. Low number of observations for Black, Hispanic, and Asian hosts contribute to imprecise estimates in smaller cities (New Orleans, Nashville have less than 100 Hispanic and Asian hosts; DC and Austin have less than 200 such hosts). The omitted category is White males. I control for the specification in Table 5, Model 4.

that better approximate the listing selection process for potential Airbnb guests.

For each sentence of each review, a sentiment-analysis algorithm evaluated how positive or negative the sentence is. In Table 7, I regress this sentiment score on the host race, controlling for my preferred specification from Table 5, Model 4. Each coefficient indicates the standardized review quality, relative to white males, that a reviewer of demographic A gave a host of demographic B. I break up my regressions by the race and sex of the reviewer, varying across the columns of Table 7. The race and sex of the host varies by row.

Column 1 of Table 7 pools the sample across all reviewers. In this pooled sample, I find that Black and Hispanic males have reviews that are 0.1–0.2 standard deviations worse than White male hosts, significant at the $p < .01$ level. Lower quality reviews might therefore explain why the listings of Black host are priced lower, but are also less demanded, than the listings of White male hosts. When I break down the results by the race and gender of the reviewer, no clear pattern emerges in the results. White reviewers show little evidence of systematic bias against minority hosts. There are some anomalies: Black male guests rate Asian hosts almost 4–8 standard deviations above the

Table 9
Main results by listing characteristics.

	> $800/night	All prices	≥ 5 reviews	Older Listings	Newer Listings	Apartments	Condos	Houses
colrule White female	0.0533	0.00537	0.0127*	0.00722	0.0130*	−0.00601	−0.0108	0.0163
	(0.0563)	(0.00550)	(0.00533)	(0.00824)	(0.00538)	(0.00562)	(0.0301)	(0.00884)
Black male	−0.0595	−0.0352**	−0.0273*	−0.0374	−0.0388**	−0.0376**	0.0387	−0.0505*
	(0.158)	(0.0133)	(0.0120)	(0.0271)	(0.0119)	(0.0121)	(0.0525)	(0.0251)
Black female	−0.0673	−0.0134	−0.0131	−0.0228	−0.0192	−0.0169	−0.0396	−0.0243
	(0.237)	(0.0106)	(0.00961)	(0.0183)	(0.0111)	(0.0110)	(0.0531)	(0.0217)
Hispanic male	0.0359	−0.0183	−0.0246	−0.0470	−0.0140	−0.0296*	0.00989	−0.0190
	(0.107)	(0.0115)	(0.0152)	(0.0242)	(0.0139)	(0.0146)	(0.0612)	(0.0258)
Hispanic female	−0.0445	−0.0200	0.00437	−0.0193	−0.00425	−0.0192	−0.0827	−0.0352
	(0.117)	(0.0126)	(0.0143)	(0.0232)	(0.0134)	(0.0138)	(0.0617)	(0.0222)
Asian male	−0.107	−0.0498***	−0.0262	−0.0225	−0.0418*	−0.0470**	−0.0359	−0.0507*
	(0.144)	(0.0137)	(0.0154)	(0.0174)	(0.0173)	(0.0162)	(0.0512)	(0.0217)
Asian female	0.260	−0.0404***	−0.0106	−0.0206	−0.0267*	−0.0378***	−0.0878	−0.0276
	(0.315)	(0.00949)	(0.0109)	(0.0156)	(0.0113)	(0.0108)	(0.0527)	(0.0175)
Constant	6.515***	4.069***	3.576***	4.287***	3.691***	3.742***	3.918***	4.959***
	(0.556)	(0.335)	(0.0827)	(0.111)	(0.272)	(0.402)	(0.305)	(0.166)
Percentage								
Observations	703	45,776	23,507	9846	25,882	28,408	1854	13,509
Adjusted R2	0.587	0.731	0.790	0.769	0.763	0.683	0.787	0.795

Standard errors in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table estimates the main results in Table 5 separately for listings of different price points, review numbers, age, and property types. The categories, from left to right, are: listings whose price is above the cutoff price in the original sample, listings of all prices, listings with more than 5 reviews, listings who have been on the market for no more than 2 years versus no more than 8 years, and listings of different property types, including apartments (includes apartments and lofts), condos (includes condos and townhouse), and houses. The omitted category is White males. I control for the specification in Table 5, Model 4. The outcome variable is the log price of the listing.

Table 10
Robustness check with controls from Edelman and Luca (2014).

	(1) Price per night
Black	−0.117***
	(0.0107)
Accommodates	0.0684***
	(0.00288)
Bedrooms	0.129***
	(0.00724)
Review scores location	−0.488***
	(0.0434)
Review scores location squared	0.0363***
	(0.00249)
Review scores checkin	−0.000735
	(0.00683)
Review scores communication	−0.00366
	(0.00718)
Review scores cleanliness	0.0230***
	(0.00417)
Review scores accuracy	−0.0186**
	(0.00574)
Host's identity verified?	0.0233**
	(0.00801)
Private room	−0.627***
	(0.00826)
Shared room	−1.123***
	(0.0183)
Location controls	Yes
Property controls	Yes
Host controls	Yes
Observations	11,999
Adjusted R2	0.619

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table presents the effect on log price of controlling for (Edelman and Luca, 2014)'s full specification using my NYC data. The omitted category for race is White hosts. The omitted category for room type is Entire Apartment. I could not control for host social media accounts as a proxy for host reliability like (Edelman and Luca, 2014) did, because Airbnb no longer provides this information. Instead, I control for "host verified", a dummy for whether Airbnb has the host's phone number and email. I similarly can not control for "picture quality", but picture quality did not significantly influence price in (Edelman and Luca, 2014)'s regression.

mean, but rate Black female hosts 3 standard deviations lower than the mean. Across all sub-demographic splits, there is not enough evidence to substantiate that minority hosts have systematically lower review quality that can explain lower prices.

4.3. Robustness to different samples

In this section, I explore patterns in price disparities across cities. I also address the sensitivity of my point estimates to various samples, including sensitizing by price, number of reviews, and property type.

4.3.1. Effects by city

In Table 8, I break up the effects of host race on listing price by city. I find that the effects are either statistically significant and negative, or roughly zero, depending on the city. Los Angeles has the most precisely estimated price disparities, as its large sample size provides enough power to precisely measure effects. Los Angeles' effects are between 5–6% for Black males, Asian males, and Hispanic females. New York, with a similarly large sample, does not exhibit the same price disparity as Los Angeles. Only Asian female hosts have a price disparity of 3.7% in New York.

By contrast, minority hosts in Chicago have large price disparities relative to White male hosts in Chicago. Asian hosts of both sexes, Hispanic females, and Black females have 7–12% lower prices in Chicago relative to White males, with Asian hosts being the worst off. These point estimates are comparable in magnitude to the effects estimated by Edelman and Luca (2014) for Black hosts in New York City. In smaller cities such as New Orleans, DC, and Nashville, there are fewer significant effects, but the effects that are statistically significant are larger: I measure an 8% disparity in DC for Asian male hosts, and a 15% disparity for Hispanic males in Nashville. There are no significant effects of host race on listing price in New Orleans.

In sum, no single city is driving the between-race variation in prices. The effects on price are mostly negative for minority hosts, with a few zero coefficients in cities with fewer observations, and one positive coefficient for Hispanic female hosts in Austin. Consistent with my pooled results, Asian hosts fare the worst, with significant price disparities in 3 of the 7 cities in my sample.

Table 11
Estimates of effect of host race on price, by host's listing count.

	Cutoff for the number of listings operated by a host					
	≤ 1	≤ 2	≤ 5	≤ 10	≤ 20	All
White female	−0.0116*	−0.00950	−0.00144	0.000385	0.00350	0.0452
	(0.00559)	(0.00492)	(0.00475)	(0.00494)	(0.00485)	(0.364)
Black male	−0.0383**	−0.0416***	−0.0372***	−0.0304**	−0.0347**	0.0103
	(0.0128)	(0.0116)	(0.0106)	(0.0110)	(0.0122)	(0.364)
Black female	−0.0479***	−0.0327**	−0.0269**	−0.0211*	−0.0180	0.0268
	(0.0123)	(0.0109)	(0.0101)	(0.00967)	(0.00999)	(0.364)
Hispanic male	−0.0110	−0.0214	−0.0180	−0.0197	−0.0202	0.0230
	(0.0132)	(0.0111)	(0.0107)	(0.0107)	(0.0114)	(0.364)
Hispanic female	−0.0228	−0.0207	−0.0277*	−0.0230*	−0.0206	0.0231
	(0.0132)	(0.0120)	(0.0108)	(0.0112)	(0.0113)	(0.364)
Asian male	−0.0519***	−0.0529***	−0.0555***	−0.0562***	−0.0445**	−0.00131
	(0.0140)	(0.0126)	(0.0123)	(0.0132)	(0.0136)	(0.365)
Asian female	−0.0530***	−0.0522***	−0.0520***	−0.0462***	-0.0394***	0.00733
	(0.0107)	(0.00994)	(0.00921)	(0.00881)	(0.00893)	(0.366)
Constant	4.044***	3.902***	3.904***	3.945***	3.912***	3.544***
	(0.381)	(0.360)	(0.347)	(0.348)	(0.348)	(0.466)
Location controls	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes
Host controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	26,788	34,920	41,626	43,953	45,073	47,388
Adjusted R2	0.693	0.705	0.712	0.717	0.720	0.732

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ *Note:* This table estimates the main results in Table 5 separately for hosts who operate different numbers of listings. The categories, from left to right, are: listings operated by hosts who own 1 property on Airbnb, listings operated by hosts who have 2 or fewer properties on Airbnb, etc. The omitted category is White males. I control for the specification in Table 5, Model 4. The outcome variable is the log price of the listing.

4.3.2. *Effects by listing type*

I also break up my sample by various listing characteristics, such as price point, age, and property type.

In the main analysis, I exclude high-priced listings over $800. Columns 1 and 2 of Table 9 present the effects of host race on price for listings over $800, and for the whole sample. The results indicate that the price disparity is more pronounced for minority hosts who own cheaper properties as opposed to expensive ones. In Column 3, I restrict the sample to listings with more than 5 reviews, the median in the data. All point estimates lose significance except for Black male hosts (negative effect of 2.8%) and White female hosts (positive effect of 1.3%). In Columns 4 and 5, I estimate the results separately for old and new listings. These columns show that price disparities are stronger for newer listings as opposed to old listings (old listings are defined as those which have been on the market for more than two years), suggesting that price disparities could be erased the longer a listing is on the market. Lastly, Columns 7–9 break price disparities up by property type, with little pattern across property types.

Statistical discrimination is one possible, but by no means definitive, hypothesis that is consistent with all of these results. In the absence of a clear signal about the listing's quality, guests could be using the host's race as a proxy for quality. A guest who is statistically discriminating would avoid the listings of minority hosts that they have little information about, and instead be willing to pay more for listings operated by White hosts. This would explain why I measure a price disparity for listings with a low, but not a high, number of reviews. Similarly, I measure no price disparity for older listings, presumably because older listings have had time to accumulate sufficient reviews.

5. Conclusion and further work

In this paper, I measure a price disparity between minority and White landlords in the P2P short-term housing market of Airbnb. I find that Asian hosts earn roughly 5% less per day, and Black hosts 3%, for comparable properties as White hosts. Since prices only matter to hosts to the extent that they affect revenue, I calculate the impact of the price disparity on annual revenue, multiplying the price per day by the number of reviews. The biggest yearly revenue loss in the entire sample is for Black females, who could expect to earn around $350 per year less than a White male operating the same listing, and for Black males, who would lose about $300.

Airbnb itself can do much to address issues of discrimination on the platform. In response to media outcry about allegations of discrimination, Airbnb updated its Discrimination Policy in September 2016, increasing opportunities for guests to book without waiting for host approval and making host profile pictures smaller. Evaluating Airbnb's efforts to address discrimination is therefore a relevant extension of this research. Since InsideAirbnb.com is continually being updated, there is now data available from webscrapes of listings after Airbnb's new discrimination policy took effect in September 2016. Future work can explore whether the policy helped curb discrimination on the platform by measuring the extent of discrimination before and after the policy took effect. If better design of user profile on P2P platforms can mitigate discriminatory behavior, then the prices of listings owned by minority and White hosts should start converging after 2016.

References

Airbnb, 2017. About Us. Web.

Altonji, J., Elder, T., Taber, C., 2005. Selection on observed and unobserved variables: assessing the effectiveness of catholic schools. J. Political Econ. 113 (1), 151–184.

Asante-Muhammed, D., Collins, C., Hoxie, J., Nieves, E., 2016. The Ever-Growing Gap. Web.

Bayard, K., Hellerstein, J., Neumark, D., Troske, K., 1999. Why are Racial and Ethnic Wage Gaps Larger for Men than for Women? Exploring the Role of Segregation. National Bureau of Economic Research. Working Paper 6997.

Bayer, P., Casey, M., Ferreira, F., McMillan, R., 2017. Racial and ethnic price differentials in the housing market. J. Urban Econo. 102 (C), 91–105.

Becker, G.S., 1957. The Economics of Discrimination. University of Chicago.

Bertrand, M., Mullainathan, S., 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. Am. Econ. Rev. 94.4, 991–1013.

Chiteji, N., 2010. The racial wealth gap and the borrower's dilemma. J. Black Stud. 41 (2), 351–366.

Cox, M., 2017. Inside Airbnb. Adding data to the debate. Inside Airbnb.

Cui, R., Li, J., Zhang, D., 2016. Discrimination with incomplete information in the sharing economy: evidence from field experiments on Airbnb.

Doleac, J., Stein, L., 2010. Race has a hand in determining market outcomes.

Economic Policy Institute, 2016. Wages by Education. State of Working American Data Library.

Edelman, B., Luca, M., 2014. Digital discrimination: the case of Airbnb.com. Working Paper 14-054.

Edelman, B., Luca, M., Svirsky, D., 2017. Racial discrimination in the sharing economy: evidence from a field experiment. Am. Econ. J. Appl. Econ. 9 (2), 1–22.
Firebaugh, G., Farrell, C.R., 2016. Still large, but narrowing: the sizable decline in racial neighborhood inequality in metropolitan aMerica, 1980–2010. Demography 53 (1), 139–164.
Fradkin, A., Grewal, E., Holtz, D., 2017. The determinants of online review informativeness: evidence from field experiments on Airbnb.
Ge, Y., Knittel, C.R., MacKenzie, D., Zoepf, S., 2016. Racial and Gender Discrimination in Transportation Network Companies. National Bureau of Economic Research. Working Paper 22776.
Gittleman, M., Wolff, E.N., 2004. Racial differences in patterns of wealth accumulation. J. Hum. Resour. 39 (1), 193–227.
Gutierrez, J., Garcia-Palomares, J.C., Romanillos, G., Salas-Olmedo, M.H., 2017. The eruption of Airbnb in tourist cities: comparing spatial patterns of hotels and peer-to-peer accommodation in Barcelona. Tour. Manag. 62, 278–291.
Guttentag, D.A., Smith, S.L., 2017. Assessing Airbnb as a disruptive innovation relative to hotels: substitution and comparative performance expectations. Int. J. Hosp. Manag. 64, 1–10.
Hanson, A., Hawley, Z., 2011. Do landlords discriminate in the rental housing market? Evidence from an internet field experiment in US cities. J. Urban Econ. 70 (2), 99–114.
Heywood, J.S., Parent, D., 2012. Performance pay and the white-black wage gap. J. Labor Econ. 30 (2), 249–290.
Kakar, V., Voelz, J., Wu, J., Franco, J., 2018. The visible host: does race guide Airbnb rental rates in san francisco? J. Hous. Econ. 40, 25–40.
Krysan, M., Crowder, K., Bader, M.D., 2014. Pathways to residential segregation. Choos. Homes Choos. Schools 27–63.
Leong, N., Belzer, A., 2017. The new public accommodations: race discrimination in the platform economy. Georget. Law J. 105:1271, 1271–1322.
Loria, S., 2019. TextBlob. Python package release v0.15.2. https://textblob.readthedocs.io/en/dev/.
Mason, P.L., 1997. Race, culture and skill: interracial wage differences among african Americans, Latinos, and Whites. Rev. Black Political Econ. 25 (3), 5–39.
Myers, C., 2004. Discrimination and neighborhood effects: understanding racial differentials in US housing prices. J. Urban Econ. 56 (2), 279–302.
Oliver, M.L., Shapiro, T.M., 2006. Black Wealth, White Wealth: A New Perspective on Racial Inequality. Taylor & Francis.
Oster, E., 2017. Unobservable selection and coefficient stability: theory and evidence. J. Bus. Econ. Stat. 0 (0), 1–18. https://doi.org/10.1080/07350015.2016.1227711.
Pope, D.G., Sydnor, J.R., 2011. What's in a picture? Evidence of discrimination from prosper.com. J. Hum. Resour. 46.1, 53–92.
Reardon, S.F., Fox, L., Townsend, J., 2015. Neighborhood income composition by household race and income, 1990–2009. Ann. Am. Acad. Political Soc. Sci. 660 (1), 78–97. https://doi.org/10.1177/0002716215576104.
Rinker, T., 2019. Sentimentr. R package version 2.7.1. https://CRAN.R-project.org/package=sentimentr.
Rugh, J.S., Massey, D.S., 2010. Racial segregation and the American foreclosure crisis. Am. Sociol. Rev. 75 (5), 629–651.
U.S. Census Bureau, 2017. Current population survey: annual social and economic supplement.
U.S. Census Bureau, 2019. Quarterly retail e commerce sales. https://www.census.gov/retail/mrts/www/data/pdf/eccurrent.pdf.
Wang, D., Xi, S., Gilheany, J., 2015. The model minority? Not on Airbnb.com: a hedonic pricing model to quantify racial bias against Asian Americans. Technol. Sci. https://techscience.org/a/2015090104/#Citation.
Ye, T., Pierce, C., Alahmad, R., Robert, L.P., 2017. Race and rating on sharing economy platforms: the effect of race similarity and reputation on trust and booking intention in Airbnb. Transforming Society with Digital Innovation. Association for Information Systems.
Yinger, J., 1986. Measuring racial discrimination with fair housing audits: caught in the act. Am. Econ. Rev. 76 (5), 881–893.