American Economic Journal: Applied Economics 2022, 14(1): 107–132
https://doi.org/10.1257/app.20190188

Can Information Reduce Ethnic Discrimination? Evidence from Airbnb†

By Morgane Laouénan and Roland Rathelot*

We use data from Airbnb to identify the mechanisms underlying discrimination against ethnic minority hosts. Within the same neighborhood, hosts from minority groups charge 3.2 percent less for comparable listings. Since ratings provide guests with increasingly rich information about a listing's quality, we can measure the contribution of statistical discrimination, building upon Altonji and Pierret (2001). We find that statistical discrimination can account for the whole ethnic price gap: ethnic gaps would disappear if all unobservables were revealed. Also, three-quarters (2.5 points) of the initial ethnic gap can be attributed to inaccurate beliefs of potential guests about hosts' average group quality. (*JEL* D83, J15, L84)

Ethnic discrimination is a pervasive phenomenon, and understanding which mechanisms are at work is needed to design effective policies. In their recent reviews, Charles and Guryan (2011) and Lang and Lehmann (2012) stress that empirical attempts to uncover these mechanisms are still inconclusive. This paper takes advantage of the features of Airbnb, a major online marketplace for short-term rentals, to measure to what extent information can influence ethnic price gaps.

Airbnb hosts list their property, set the daily price, and provide information about themselves (at least first name and picture) and their properties (precise location, equipment, local amenities, pictures, etc.). Potential guests book properties on given dates at the price set by the hosts. In this paper, we study the differential between

* Laouénan: CNRS, Centre d'Economie de la Sorbonne and Sciences-Po LIEPP (email: Morgane.Laouenan@univ-paris1.fr); Rathelot: Department of Economics, University of Warwick and CEPR (email: r.rathelot@warwick.ac.uk). Alexandre Mas was coeditor for this article. We would like to thank David Autor, Manuel Bagues, Pat Bayer, Leah Boustan, Clément de Chaisemartin, Raj Chetty, Tatiana Coutto, Bruno Decreuse, Christian Dustmann, Andrey Fradkin, Lucie Gadenne, Ingrid Gould Ellen, Zack Hawley, Nathan Hendren, Simon Jäger, Larry Katz, Kevin Lang, Victor Lavy, Thomas Le Barbanchon, Attila Lindner, Andrea Moro, David Neumark, Claudia Olivetti, Barbara Petrongolo, Imran Rasul, Yona Rubinstein, Dan Svirsky, Fabian Waldinger, Etienne Wasmer, Natalia Zinovyeva, and participants in seminars at UC Irvine, UCSB, UCLA (CCPR), IEB Barcelona, Umeå, Sorbonne, PSE, Warwick, DIW, Sciences Po, Louvain, IZA, TSE, LSE, Tinbergen, Howard, Bristol, OECD, BC, Duke, NYU-AD, Turing Institute, and several workshops and conferences (Aix-Marseille 2014, Nizza Monferrato 2015, SOLE 2016, IZA-SOLE TAM 2016, CEPR Labour Symposium 2016, IAAE 2018, UEA 2018) for fruitful discussions and comments. Earlier versions of this paper were circulated under the title "Ethnic Discrimination on an Online Marketplace of Vacation Rentals." We gratefully acknowledge financial support from CAGE (Warwick), from Sciences-Po LIEPP (ANR-11-LABX-0091, ANR-11-IDEX-000502) and from Labex OSE (Opening Economics). Airbnb was not involved in the research or publication of the paper and does not endorse its findings or conclusions. All errors are our own.

† Go to https://doi.org/10.1257/app.20190188 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

prices set by hosts who belong to an ethnic minority and those set by majority hosts. We ask whether this ethnic price gap that remains unexplained by differential in observable characteristics is driven by statistical discrimination or other factors.

While taste-based discrimination stems from the existence of racial preferences or an aversion toward cross-racial interaction (Becker 1957), statistical discrimination is the result of imperfect information and ethnic differences in the mean or the variance of unobservable characteristics (Phelps 1972, Arrow 1973, Aigner and Cain 1977). The most direct approach to distinguish statistical discrimination from other mechanisms is to measure how the ethnic gap varies with the amount of information about a service (Farber and Gibbons 1996, Altonji and Pierret 2001).

We adapt the Altonji and Pierret (2001) approach to our setting, in which we observe a measure of the quantity and quality of information about a property available to potential guests. In contrast with labor markets, the short-term rental market is well suited for testing statistical discrimination because (i) transactions happen frequently compared to changes in the quality of the property, (ii) ratings and the number of reviews can be observed, and (iii) large sample and longitudinal data are available. The profiles of new properties contain only self-reported information. After their stay, guests are allowed to leave a quantitative rating and a qualitative assessment of both the property and the host. As the number of reviews grows, more information becomes available to potential guests.

We rely on a simple conceptual framework where the quality of the properties is partially unobservable. Initially, a property has no reviews, and potential guests can only infer unobservable quality using a host's ethnicity, conditional on other observables. As a property accumulates reviews, potential guests aggregate the content of reviews and the host's ethnicity to form the best possible guess about the property's unobservable quality. From this model, we derive a first test for the existence of statistical discrimination that relies on the longitudinal nature of the data. In the presence of statistical discrimination, the price gap should decrease with the number of reviews and tend to zero, conditional on observables and the measure of quality provided by the reviews. If, instead, the price gap is due to taste-based discrimination or ethnic differentials in variables that are not observable to the econometrician but observable to potential guests, the price gap should remain stable with the number of reviews.

Guests' beliefs about unobservable quality do not need to be accurate (Bordalo et al. 2016). If potential guests believe that properties belonging to an ethnic minority are, on average, worse than they actually are, an ethnic price gap will emerge. We categorize this phenomenon as statistical discrimination, as the gap will disappear when more information about quality becomes available. We account for inaccurate beliefs in our conceptual framework and provide an additional empirical prediction that allows us to measure its contribution to the statistical discrimination component of the ethnic price gap.

As an illustration, suppose that properties held by minority and majority hosts have the same average quality but potential guests believe that minority listings are worth 10 percent less. When properties have no reviews, minority-held properties will be priced 10 percent lower. When the number of reviews grows to infinity, the average price will be identical in both groups. Conversely, imagine now that the average quality is indeed 10 percent lower in the minority group, i.e., beliefs are accurate. In this case, there will be a 10 percent price gap between the two sets of properties whether there are few or many reviews. However, if we follow two properties of the same quality—one held by a minority host, the other one by a majority host—there will be an initial price gap of 10 percent that will converge to 0 as they accumulate reviews. Formally, we will use the cross-group differential slope of prices with respect to the number of reviews to quantify (i) the ethnic price gap due to statistical discrimination (when we control for a proxy of the quality of the listing) and (ii) the part of statistical discrimination that is due to inaccurate beliefs (when we don't control for a proxy of the quality of the listing).

Our dataset includes daily prices, characteristics of hosts and properties, and associated reviews. We collected data relating to around 670,000 properties, corresponding to apartments to rent in 19 cities in North America and Europe. In total, 21 waves of data, collected between June 2014 and November 2017, form an unbalanced panel of 3.8 million observations. The ethnic minority groups we consider are hosts with Arabic or Muslim first names and hosts categorized as Black based on their profile pictures.

We find that the within-city raw ethnic price gap is around 16 percent. The set of observable characteristics about the property (including its location) is rich and explains more than 67 percent of the variance of the price. When the heterogeneity in observable characteristics is accounted for, the ethnic price gap is reduced to a significant 3.2 percent.[1] This figure may look small, but a price gap of 3.2 percent represents a gap of 17 percent of the hosts' surplus, which is substantial.[2] We show that prices increase faster with the number of reviews when the host belongs to an ethnic minority, conditional on the average rating based on reviews received by the listing over the whole observation period. We find that 3.4 percentage points of the price gap (i.e., the whole gap) are accounted for by statistical discrimination. Of these 3.4 percentage points, 2.5 are due to inaccurate beliefs—that is, the fact that potential guests underestimate the average unobservable quality of minority properties compared to majority ones. The difference, a statistically significant 0.9 percentage points, is due to the true difference in average unobservable quality between the two groups.

Our paper contributes to the growing but largely inconclusive literature on the sources of discrimination. Altonji and Pierret (2001) find little evidence for statistical discrimination in wages on the basis of ethnicity in the US labor market. A strand of literature uses the fact that the relevant outcome is perfectly observed ex post. Knowles, Persico, and Todd (2001) show that vehicles of African Americans are more often searched by the police and that statistical discrimination explains more than the observed gap.[3]

[1] Edelman and Luca (2014) are the first to document the existence of significant ethnic price gaps on Airbnb, focusing on the Black-White price gap in New York City.

[2] We use the estimates from Farronato and Fradkin (2018) for the hosts surplus and average price. See Section IC for details.

[3] Using data from a peer-to-peer lending website, Pope and Sydnor (2011) find that African Americans are likely to be subject to statistical discrimination. Using data from television game shows, Anwar (2012) finds that White contestants believe that African Americans have lower skill levels, while Levitt (2004) and Antonovics, Arcidiacono, and Walsh (2005) find no evidence of discrimination.

The amount and nature of information available to discriminatory agents can also be manipulated experimentally. In the online rental apartment market, Ewens, Tomlin, and Wang (2014) find that the response to differential quality varies in a way that is consistent with statistical discrimination.[4] Cui, Li, and Zhang (2020), in a paper developed independently from ours, send Airbnb accommodation requests expressed by African-American-sounding-name and White-sounding-name guests in three American cities. They compare requests by guests who have no reviews to those with one review and find that both positive and negative reviews reduce the ethnic acceptance gap by hosts. Experimental evidence can be complemented by lab games to separate discrimination mechanisms. In the case of the sports card market, List (2004) finds that the lower offers received by minorities are mainly explained by statistical discrimination.[5]

Other approaches have been used to separate sources of discrimination. Wozniak (2015) shows how a policy (drug-testing legislation) that affects a relevant dimension of the unobservables (drug use) can provide evidence of statistical discrimination against low-skilled African American men. The heterogeneity in agents' prejudice, whether revealed or assumed, is sometimes used to infer which source of discrimination is more prevalent. Bayer et al. (2017) show that the minority home buyers pay higher prices in the US housing market regardless of the sellers' ethnicity, suggesting statistical discrimination. Zussman (2013) finds that the discrimination toward Arabs in an online market for used cars in Israel is not related to sellers' revealed attitudes toward Arabs. Doleac and Stein (2013) show that online iPod ads featuring dark-skinned hands receive fewer offers, with poorer outcomes in thin markets and those with higher racial isolation and crime.[6]

Following Bordalo et al. (2016), a recent literature has attempted to go beyond the dichotomy between taste-based and statistical discrimination. Customers may have inaccurate beliefs about sellers' quality, which itself could be due to stereotypes. To our knowledge, few papers have tried to isolate this source of differentials from other discriminatory mechanisms. Arnold, Dobbie, and Yang (2018) and Dobbie et al. (2018) isolate "inaccurate stereotyping" from racial animus in the context of bail decisions and consumer lending, and they find that this mechanism explains a large part of the racial bias. In our setting, we classify ethnic gaps coming from inaccurate beliefs as statistical discrimination because new information will reduce these gaps. Bohren, Imas, and Rosenberg (2019) design a randomized experiment on a math forum to measure the dynamics of discrimination against women, allowing for "belief-based" discrimination.

[4] Conversely, in their correspondence studies on the US and Canadian labor markets, Bertrand and Mullainathan (2004) and Oreopoulos (2011) find that adding information or enhancing résumés do not benefit minority applicants. Heckman (1998) and Neumark (2018) list some of the challenges associated with the current use of experimental methods for discrimination.

[5] See also Fershtman and Gneezy (2001) and Castillo and Petrie (2010) for papers using lab experiments for this purpose.

[6] Taking the opposite approach, Charles and Guryan (2008) introduce an indirect test of the Becker (1957) prejudice model based on associations between prejudice and wages and find that around one-quarter of the unconditional racial wage gap is due to prejudice while the three other quarters can be due to differences in unobservables or other forms of discrimination.

We also contribute to the growing literature on the role of information provided by online market intermediaries on markets' outcomes.[7] Our paper is related to Autor and Scarborough (2008), who show that while minorities perform poorly on job tests, introducing job testing in a large retail firm has no impact on minority hiring. We contribute to the study of ethnic discrimination on the rental market with the unprecedented scale of our data, covering 19 cities in 8 countries in both Europe and North America. The online marketplace Airbnb (2014–2017) is relevant in itself from an economic point of view: launched in 2008, the website offers more than 7 million listings in 220 different countries and claims to have served over 750 million guests.[8]

Section I presents the context, the data, and the first empirical evidence about ethnic price gaps. Section II introduces our conceptual framework. Section III presents our empirical strategy and main results. Section IV provides additional results and discusses alternative explanations. Section V concludes.

I. Context and Data

A. *Description of the Platform*

Airbnb connects hosts looking for opportunities to let their properties with potential guests looking for a place to stay. Both types of users have to register and provide a large set of information about themselves. Hosts also have to provide information about their properties. In practical terms, potential guests usually start by typing the city where and time period when they want to stay into the search engine. They can filter the results of the search according to the price or other characteristics (e.g., accommodation capacity, room type, property type, number of bedrooms). At that stage, potential guests obtain a list of results with basic information, among which is the daily price, a picture of the property, a thumbnail photo of the host, and the overall rating (presented in stars and defined as the average rating over the reviews of the listing). When they click on one of the listings, they have access to more detailed information, notably the first name of the host, a detailed description of the property, a standardized list of the offered amenities, more pictures, and detailed reviews from previous guests.[9]

Hosts can revise the price of their properties at any moment. The potential guest decides which place she prefers among those available during the period selected and commits by clicking on the "Book It" button. The decision is then in the hands of the host. She can accept or reject the guest without any justification.[10] A guest who gets rejected receives an email encouraging her to look for another place. The rejection is not reported on her profile. If the host accepts the guest, the deal is

[7] See, e.g., Autor (2001, 2009); Bagues and Labini (2009); Pallais (2014); Horton (2017); Pallais and Sands (2016); Brown, Setren, and Topa (2016); and Stanton and Thomas (2018).

[8] https://news.airbnb.com/fast-facts/.

[9] See Figure A1 for a screenshot of a listing corresponding to the period of the data we use.

[10] Rejections are frequent; see Fradkin (2017).

concluded and there is no way to modify its terms.[11] The potential guest may decide to cancel her booking. In this case, the terms of the cancellation policy (specified on the listing) apply: depending on the flexibility of the policy, penalties of different amounts are charged. The host may also decide to cancel the deal. In this case, there is no financial penalty, but there is a reputation cost: the website records on the host's profile that she has canceled a deal.

We consider hereafter that potential guests are price-takers. Using a simple model of supply and demand, we consider that the existence of discrimination toward hosts, which triggers a shift in demand, should translate into lower prices. We formalize this idea in the section dedicated to the conceptual framework.

B. *Data*

We collect data from publicly available web pages of the marketplace. We store all information visible on the first page of the listing: the price asked by the host, the characteristics of the listing, the characteristics of the host, and the last ten reviews and ratings. We focus on the 19 cities in North America and Europe with the highest number of listings.[12] We repeat the collection process every two to three weeks between June 2014 and June 2015 and add a last wave in November 2017, obtaining 21 waves in total.[13] Our sample includes 663,090 distinct properties. The panel is unbalanced: some properties enter the system, and others exit.[14]

We restrict our analysis to the subsample of listings that have gained at least one review over the observation period. The motivation behind this restriction is to work with active listings where there have been established transactions and feedback from the guests. This restriction reduces the sample size from 663,090 to 220,939: most Airbnb listings do not get any reviews during that period (or they exit before they do). The distribution of the number of waves during which we observe each property is in online Appendix Figure A2: 13 percent of listings are observed in at least 20 waves, and half of the listings are observed in at least 11 waves.

For each property, we know the main characteristics: the type of property, size, type of bed, amenities, services, and rules. Most properties are apartments, and the entire place is let in 70 percent of cases. Properties are rather small, with 1.2 bedrooms on average, and they can host, on average, 3 guests. Some properties add a cleaning fee and charge for additional people. We count the cleaning fee directly into the price in order to obtain the final price paid by the guest.[15] We also obtain some information about the hosts on their profile pages. Aside from the first name, a

[11] While the acceptance/rejection decision would in itself be of interest as regards discrimination, we do not have the necessary data to study that side of the market. See Edelman, Luca, and Svirsky (2017) for a field study about discrimination against potential guests.

[12] The cities are London, Paris, Madrid, Barcelona, Rome, Milan, Florence, Amsterdam, Berlin, Marseille, Vancouver, Toronto, Montreal, Boston, New York City, Miami, Chicago, San Francisco, and Los Angeles. See online Appendix Table A1 for the number of observations and listings by city.

[13] See the collection dates of each wave in online Appendix Table A2. The last wave was added because we wanted to increase the longitudinal depth of our dataset.

[14] We check the possibility of differential attrition between ethnic groups. In online Appendix Section C, we show that the probability to leave the market is the same for minority and majority groups after controlling for property characteristics, ratings, and neighborhood fixed effects.

[15] We assume guests stay on average six days and add a sixth of the fee to the price.

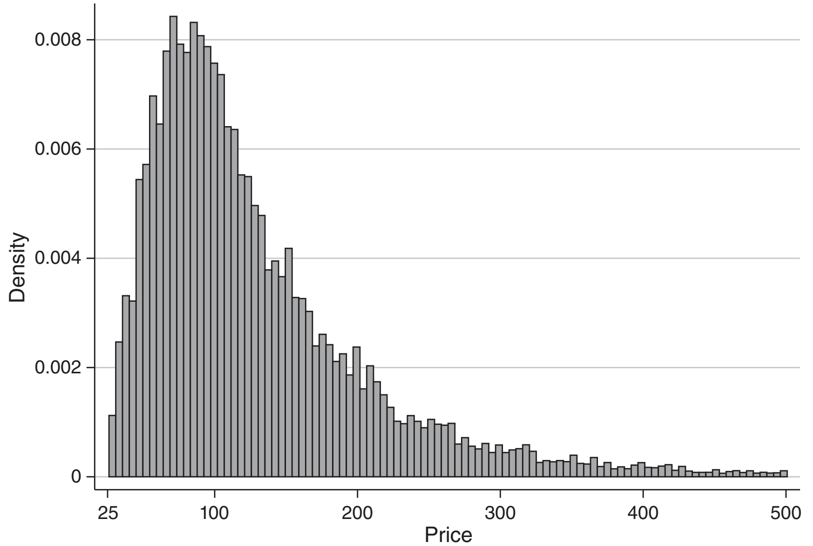

Figure 1. Distribution of Daily Price

Notes: This figure shows the distribution of the final price (cleaning fees included) represented through 100 bins. The sample is restricted to listings that have gained at least one review over the observation period. The figure is right truncated with a maximum of $500.

picture, and a free-text description, potential guests know whether hosts have other properties and when they joined the platform. Most hosts have only one property and have joined the platform recently. See the full list of characteristics of properties and hosts in online Appendix Table A3.

Figure 1 shows the distribution of daily prices. There is much variation in prices across properties. To reduce the influence of outliers, we drop 1 percent of the observations at the top and bottom of the price distribution. The first quarter is $75, the median $107, and the third quarter $160 per night. The skewness of the distribution implies that the mean price is $130. The daily price varies across cities and according to the amenities of the listing (number of accommodates, bedrooms, bathrooms, etc.). Online Appendix Table A4 provides details on how amenities affect the price.

In order to identify statistical discrimination, we need to have enough variability in the number of reviews, and we need reviews to be informative about listings' quality. Online Appendix Figure A3 displays the distribution of reviews across the observations of our sample (left panel) and the variation of the number of reviews between the last and first observations (right panel) and shows that the sample offers a decent amount of heterogeneity in the number of reviews.

For each property, we use the last observed rating, which represents the average of all ratings received over its lifetime on Airbnb. Ratings can vary between one and five stars (with half-star increments), and the distribution is skewed toward good ratings, as documented in Fradkin, Grewal, and Holtz (2018). If we consider the last rating observed for each property of our sample, 44 percent of observations have 5

Table 1—Raw Price Gaps by Ethnic Groups

	Sample size	Share (percent)	Within-city-wave gap (percent)
Majority	2,320,285	93.8	—
Blacks (US/Canada)	49,706	2.0	31.3
Blacks (Europe)	21,365	0.9	26.3
Arabic/Muslim (US/Canada)	31,145	1.3	4.7
Arabic/Muslim (Europe)	52,050	2.1	6.8

Note: The within-city-wave gaps are obtained as the coefficients on the dummies of each group in a linear regression of the log price that includes dummies for the interaction of each city and each wave.

stars and 39 percent 4.5 stars. By contrast, only 4 percent have 3.5 stars or less (see online Appendix Table A5).

C. *Ethnic Groups and Gaps*

We consider two groups of ethnic minorities. First, we consider Blacks, which we identify using the pictures provided on their host profile.[16] Second, we consider hosts that have a first name associated with Arabic, Muslim, or Sub-Saharan African ethnicity (labeled Arabic/Muslim hereafter).[17] We use two different sources to obtain a complete list of names: Jouniaux (2001) and Hawramani (2015).[18] Table 1 displays the share of ethnic groups in the sample and the price gap, controlling for interacted dummies for the city and the wave of observation (i.e., within-city-wave price gap). Blacks living in North America represent 2 percent of the observations, and those living in Europe 0.9 percent. Hosts with Arabic/Muslim names in North America represent 1.3 percent of the sample, and those in Europe 2.1 percent of the sample. Compared to their share in total population (both in North America and Europe), ethnic minorities seem to be underrepresented on the website. A possible explanation is that only those with a fairly good-quality property may attempt to rent on Airbnb, which would induce a positive selection. Overall, the share of minorities is 6.2 percent, but this share varies across cities. New York City has 9.8 percent of African American and 3.7 percent of Arabic/Muslim observations. London and Paris both have around 5 percent of Arabic/Muslim observations, while this group represents less than 1 percent of the observations in Milan and Rome. The raw price gap for Arabic/Muslim hosts, controlling only for the heterogeneity across cities and waves, is around 5 percent in North America and 7 percent in Europe. For Blacks, the raw gap reaches 31 percent in North America and 26 percent in Europe.

[16] Specifically, pictures were coded by workers specialized in this picture-coding task. Workers were asked to code each picture in three categories: (i) whether they thought that at least one person in the picture was African American, (ii) whether nobody in the picture was African American, or (iii) whether it was impossible to say anything about the ethnicity of anyone in the picture, or the picture was not showing any human being (pictures of flats, pets, furniture, landscape, etc.). We created one dummy variable equal to one in the first case. In order to check their results, we selected random samples and found mistakes at a rate below 5 percent for this dummy variable. In online Appendix Section D, we provide suggestive evidence that minority hosts do not seem to strategically obfuscate their skin color.

[17] See Rubinstein and Brenner (2014) for an example of discrimination based on names.

[18] The list of Arabic/Muslim names we used is available in the online data Appendix.

Table 2—Ethnic Price Gap, by Specification

	log daily rate			
	(1)	(2)	(3)	(4)
Minority	−0.169	−0.111	−0.067	−0.032
	(0.008)	(0.005)	(0.009)	(0.006)
City-wave fixed effects	Yes	Yes	Yes	Yes
Neighborhood fixed effects	No	No	Yes	Yes
Block fixed effects	No	No	Yes	Yes
Property characteristics	No	Yes	No	Yes
Adjusted R^2	0.15	0.63	0.36	0.73
Observations	2,474,551	2,474,551	2,474,551	2,474,551

Notes: OLS regression of the daily log price on the minority dummy, controlling city-wave fixed effects. See the list of all property characteristics in online Appendix Table A4. Robust standard errors clustered at the property level.

Table 2 shows the ethnic price differential for several specifications. The first column displays within-city-wave raw differential in daily log prices: only differences in cities and waves are taken into account, no differences in characteristics. The raw ethnic gap is large (17 percent) and highly significant. Accounting for ethnic disparities in property observable characteristics reduces the gap to 11 percent (column 2), which shows that ethnic minorities have, on average, properties of lower observable quality. Characteristics include all information provided by the host concerning her listing and her profile. The overall number of pictures and the number of pictures taken by professionals are also taken into account in our estimation.[19] Observable characteristics explain a large part of the variance: the adjusted R^2 jumps to 0.63 in the second column.

A major source of heterogeneity across listings is their location. Airbnb does not publicize the exact coordinates of a given listing but rather a 0.3-mile-radius circle. We build a grid of blocks that are 0.6 miles large for all cities and assign each listing to the block where the centroid of its circle is located. On top of this, Airbnb assigns listings to the neighborhood they belong to. In total, we work with 6,700 squared blocks and 1,500 neighborhoods. Throughout the paper, controlling for the listing's location means that we control for both the block and the neighborhood where the listing is located. Online Appendix Table A6 shows the number of neighborhoods and blocks per city. The ratio of blocks per neighborhood mainly depends on the area and density of the city.

Including neighborhood and block fixed effects reduces the ethnic price gap from 17 percent to 7 percent (column 3), and the adjusted R^2 increases from 0.15 to 0.36. Finally, in the fourth column, both location and property characteristics are included in the regression: the residual ethnic price gap is reduced to 3.2 percent but is still very significant. The adjusted R^2 is high in this last specification, equal to 0.73. Compared to the unexplained ethnic wage gaps found on labor markets, a figure of 3.2 percent may look small. To make sense of it, one has to compare it to the

[19] We identify the number of "verified photos" on each listing. Verified photos mean a professional Airbnb photographer visited the listing and captured and uploaded the photos. Airbnb contracts the photographers, and the photography service is free for hosts. More information can be found at https://airbnb.com/info/photography.

Table 3—Ethnic Price Gap, for Several Segments of the Number of Reviews

	log daily rate				
	(1)	(2)	(3)	(4)	(5)
Minority	−0.034	−0.026	−0.023	−0.022	−0.021
	(0.009)	(0.007)	(0.009)	(0.011)	(0.015)
Number of reviews	0	1–4	5–19	20–49	50+
Minority share (percent)	6.1	6.2	6.3	6.3	6.1
Adjusted R^2	0.72	0.75	0.78	0.79	0.80
Observations	351,631	808,000	789,798	352,906	172,216

Notes: OLS regressions of the daily log price on the minority dummy, controlling for neighborhood fixed effects, block fixed effects, property characteristics, and ratings (for properties with at least one review). See the list of all property and host characteristics in online Appendix Table A4. Robust standard errors clustered at the property level.

average surplus that hosts realize on Airbnb. Working with the 50 largest US cities, Farronato and Fradkin (2018) find that hosts enjoy an average of \$26 in surplus per night booked for an average price of \$136. A 3.2 percent ethnic price gap represents a loss of \$4.4 per night, i.e., a 17 percent ethnic differential in surplus.[20]

Table 3 shows the coefficient associated to the ethnic minority dummy in a regression of the log price on property characteristics, neighborhood dummies, and ratings on several subsamples defined by the number of reviews. We find that the point estimates differ across subsamples: from 3.4 percent for listings with no reviews to an insignificant 2 percent for listings with more than 49 reviews. These results are suggestive of the existence of statistical discrimination if reviews bring information that helps offset the ethnic price gap. However, there are two caveats about this interpretation. First, we don't have the statistical power to reject the null hypothesis that all five coefficients are equal. Second, there is a potential sample bias: properties with no reviews are likely to be different from those with more than 49 reviews. In the remainder of the paper, we introduce a conceptual framework leading to an empirical test of statistical discrimination that leverages the longitudinal dimension of our data.

II. Conceptual Framework

In this section, we present a simple conceptual framework where ethnic price gaps can be due to statistical discrimination, taste-based discrimination, ethnic differentials in characteristics unobserved by the econometrician but observed by potential guests, and ethnic differentials in outside options.

A. *Prices and Demand as a Function of Quality*

At each period (say, a week), a host shares her working time between two activities: renting her property (looking for guests, communicating with guests, cleaning

[20] Our framework (*see infra*) allows the remaining 3.2 percent gap to be explained by the uneven distribution of unobservables across ethnic groups. Thus, we refrain from using a test à la Altonji, Elder, and Taber (2005).

up) or working at a regular job. Here, L is the amount of labor dedicated to renting and $1 - L$ the amount dedicated to the regular job. Renting the property is assumed to have decreasing returns to scale: the number of nights supplied is equal to $L^{\tilde{\alpha}}$, with $\tilde{\alpha} \in (0,1)$. The regular job has constant returns to scale. Given the price of a night P and the wage of the regular job W, the revenue of the host over the period is $PL^{\tilde{\alpha}} + W(1 - L)$.

From the point of view of potential guests in a particular market, properties differ in three dimensions: quality Q, price P, and the ethnicity of the host m (equal to 1 if the host belongs to an ethnic minority, 0 otherwise). Demand D for a particular property is assumed to increase with Q and decrease with P. Taste-based discrimination is embedded in this framework: demand is assumed to be divided by $\Gamma > 1$ when $m = 1$, relatively to $m = 0$. Assuming β and κ are strictly positive, we write demand as

$$D = \frac{Q^{\beta}}{P^{\kappa}\Gamma^{m}}.$$

Taking Q and m as given, hosts can set the price P and the effort L they dedicate to renting to maximize their profit under the demand constraint

$$\max_{P} PD(P) + \left(1 - D^{1/\tilde{\alpha}}(P)\right)W \qquad \text{with} \quad D(P) = \frac{Q^{\beta}}{P^{\kappa}\Gamma^{m}}.$$

Solving the program, hosts will set the log price such that

$$p = p_0 + \lambda\alpha w + \lambda\beta q - \lambda\gamma m, \tag{1}$$

where $p = \log P$, $w = \log W$, $q = \log Q$, $\gamma = \log\Gamma$, $\alpha = \tilde{\alpha}/(1-\tilde{\alpha})$, $\lambda = (\kappa+\alpha)^{-1}$, and $p_0 = \lambda\alpha\log(\tilde{\alpha}(\kappa-1)/\kappa)$.

B. *Imperfectly Observed Quality*

Potential guests cannot observe quality perfectly. They have an information set that contains everything that the website displays about the listing (description, pictures, host ethnicity, and ratings, if any). We assume that quality q is the sum of two components orthogonal from each other: $q = \zeta + \nu$. Note that ζ is immediately observable in the listing by potential guests, while ν is unobservable when the listing has no reviews but perfectly observable when it has an infinite number of reviews.

We assume that the distribution of the quality component inferred from reviews conditional on ethnicity $\nu|m$ is a $\mathcal{N}(\bar{\nu}_m, \sigma_\nu^2)$.[21] Each review transmits a signal, which is a random draw around ν in a normal distribution, the error on a single review being of variance σ^2.[22] Potential guests observe r, the average signal transmitted by the set of K existing reviews, which is distributed as a $\mathcal{N}(\nu, \sigma^2/K)$. Denoting

[21] In online Appendix Section E, we show that we can obtain a similar expression for the expectation of the price when we assume, more realistically, that ν follows a nonnormal prior distribution (beta distribution).

[22] This assumption is not obvious. Reviews may depend on the quality but also on prices. We abstract from this aspect to simplify.

$\rho = \sigma^2/\sigma_\nu^2$, the expected ν for a listing with average r, K reviews, and host ethnicity m is the weighted average between the prior $\bar{\nu}_m$ and the signal r:

$$E(\nu|r,K,m) = \frac{Kr + \rho\,\bar{\nu}_m}{K+\rho}.$$

From the point of view of potential guests, the expected quality of a listing with K reviews, a signal r, a host ethnicity m, and observable characteristics ζ is

$$E(q|\zeta,r,K,m) = \zeta + \frac{Kr + \rho\,\bar{\nu}_m}{K+\rho}.$$

In a context where quality is not perfectly observed, the host will combine the expected quality conditional on the information set of potential guests with equation (1) to form the price-setting rule

$$p = p_0 - \lambda\gamma m + \lambda\alpha w + \lambda\beta\zeta + \lambda\beta\frac{Kr + \rho\,\bar{\nu}_m}{K+\rho}. \tag{2}$$

III. Empirical Strategy and Results

In this section, we first derive empirical predictions from the theoretical framework in order to identify statistical discrimination from the other mechanisms. Second, we show how we can separate the part of the statistical discrimination ethnic price gap that corresponds to differences in true average unobservable quality from the part that corresponds to inaccurate beliefs. Finally, we present the estimation results.

A. *Identification Strategy When Beliefs Are Accurate*

We assume that the econometrician observes for each listing i a sequence of prices p_{it} at different dates t, the associated number of reviews K_{it}, listing characteristics X_{it}, the host's ethnicity m_i, and the last-known average rating $\bar{r}_i$. We assume that, conditional on listing fixed effects and X_{it}, the variability in prices over time does not come from variations in features ζ or the outside option w.

Prediction 0 (Accurate Beliefs).—Under the previous set of assumptions, our main empirical prediction is that the nonlinear regression with listing fixed effects specified in equation (3) will allow the econometrician to identify $\beta_m = \lambda\beta(\bar{\nu}_1 - \bar{\nu}_0)$, the ethnic price gap that can be attributed to statistical discrimination, as well as ρ, the number of reviews that are necessary to make up for half of the gap due to statistical discrimination. Proofs are in online Appendix Section F. Formally,

$$p_{it} = \sum_{\bar{r}\in \mathrm{supp}(\bar{r})} \beta_{\bar{r}} \mathbf{1}\{\bar{r}_i = \bar{r}\}\frac{K_{it}}{K_{it}+\rho} - \beta_m m_i \frac{K_{it}}{K_{it}+\rho} + \mu_i + X_{it}\beta_x + \varepsilon_{it}. \tag{3}$$

Once we control for the time evolution of prices that corresponds to listings of quality $\bar{r}_i$ (where $\mathrm{supp}(\bar{r})$ is the set of all possible values of $\bar{r}_i$), the specific time evolution of the prices of the listings of minority hosts reveals the extent of statistical discrimination. If minority hosts have, on average, listings that have worse unobservables than majority hosts, $\bar{\nu}_1 < \bar{\nu}_0$, we have $\beta_m < 0$. Intuitively, all minority hosts have to post lower prices initially to compensate for lower expectations from the demand side. Within bins of listings of the same quality, the price of listings belonging to minority hosts will increase faster with the number of reviews than will those belonging to majority hosts. As information about ν becomes more accurate, the price of minority-host listings will catch up and converge toward the price of their majority-host counterparts.

Note that within this framework, we cannot disentangle the other possible channels causing ethnic price gaps. Differences in unobservables that do not evolve with reviews (ζ), differences in outside option (w), and taste-based discrimination (γ) are pooled together and absorbed by the listing fixed effects.

B. *Identification Strategy When Beliefs Are Inaccurate*

So far, we have assumed that potential guests have accurate beliefs and that statistical discrimination exists because the average quality of the listings proposed by minorities is lower than those proposed by the majority ($\bar{\nu}_1 < \bar{\nu}_0$). Here, we relax the assumption that guests have accurate beliefs about the average quality $\bar{\nu}$ in each group. For simplicity, let us assume that potential guests make no mistake on the average quality $\bar{\nu}_0$ of listings held by majority hosts. However, their prior on the average quality $\tilde{\nu}_1$ might differ from the true average quality $\bar{\nu}_1$. For instance, guests might wrongly believe that minority listings are worse on average than they actually are ($\bar{\nu}_1 - \tilde{\nu}_1 > 0$).

When beliefs are allowed to be inaccurate, we can decompose into two components the term $\bar{\nu}_0 - \tilde{\nu}_1$ that we attribute to statistical discrimination. The first one, $\bar{\nu}_0 - \bar{\nu}_1$, is due to the difference in the true average unobservable quality across groups. The second one, $\bar{\nu}_1 - \tilde{\nu}_1$, is due to the difference between the average true unobservable quality and the (potentially inaccurate) beliefs that potential guests hold about it.

Predictions 1 and 2 (*Inaccurate Beliefs*).—Under this new, more general set of assumptions, our first empirical prediction is that the regression specified in equation (3) will allow the econometrician to identify $\beta_m = \lambda\beta(\tilde{\nu}_1 - \bar{\nu}_0)$, the ethnic price gap that can be attributed to statistical discrimination, as well as ρ.

Our second empirical prediction is that the nonlinear regression with listing fixed effects in which we do not include the interaction terms between ratings dummies $\mathbf{1}\{\bar{r}_i = \bar{r}\}$ and the evolution in the number of reviews $K_{it}/(K_{it} + \rho)$, as specified in equation (4), will allow the econometrician to identify $\tilde{\beta}_m = \lambda\beta(\tilde{\nu}_1 - \bar{\nu}_1)$, the ethnic price gap that can be attributed to inaccurate beliefs, as well as ρ. Proofs are in online Appendix Section F. Now,

$$(4) \qquad p_{it} = \beta_k \frac{K_{it}}{K_{it} + \rho} - \tilde{\beta}_m m_i \frac{K_{it}}{K_{it} + \rho} + \mu_i + X_{it}\beta_x + \varepsilon_{it}.$$

Whenever beliefs are correct ($\tilde{\nu}_1 = \bar{\nu}_1$), the estimate of $\tilde{\beta}_m$ in equation (4) should be equal to 0, while the estimate of β_m in equation (3) will be equal to $\bar{\nu}_1 - \bar{\nu}_0$. When $\tilde{\nu}_1 = \bar{\nu}_1$, potential guests are right, on average, about the property quality in each group. When we do not control by the price evolution specific to ratings' levels, the prices of minority-owned listings will evolve at the same pace as those of the majority.

On the contrary, when both groups have the same true average quality ($\bar{\nu}_1 = \bar{\nu}_0$) but potential guests have inaccurate beliefs ($\tilde{\nu}_1 < \bar{\nu}_1$), the estimate of β_m in equation (3) and the estimate of $\tilde{\beta}_m$ in equation (4) should be equal to each other and strictly positive. In the empirical subsection below, we will report β_m, the total statistical discrimination gap; $\tilde{\beta}_m$, the ethnic gap due to inaccurate beliefs; and $\beta_m - \tilde{\beta}_m$, the ethnic gap due to differences in the true average quality.

Where could inaccurate beliefs come from? Listings on Airbnb are a selected subset from all homes. Most likely, hosts self-select into Airbnb based on the quality of their homes, and it is possible that minority listings are even more selected, given that ethnic minorities tend to live in areas and properties that are less valued by guests. This differential selection may induce a gap between the guests' beliefs about unobservables and actual quality for minority listings. This hypothesis is consistent with the fact that the share of minorities on Airbnb is smaller than their share in the whole population. Another way to explain why potential guests, who are primarily from the majority group, have unduly low beliefs about the quality of minority hosts' listing is provided by the model of stereotypes in Bordalo et al. (2016).

C. *Main Empirical Results*

We estimate regressions (3) and (4) using 4 values for the support of the last observed average rating (5, 4.5, 4, and $\leq$3.5 stars), including listing fixed effects. We use all property characteristics as well as city dummies interacted with the wave in which the listing appears. We estimate the main parameters of interest: β_m, $\tilde{\beta}_m$, and ρ. For inference, we bootstrap at the property level.

We present the estimation results in Table 4. In the first column, we show the results of regression (3). The point estimate for the total ethnic gap corresponding to statistical discrimination is 3.4 percent. This figure is similar to the ethnic price gap observed in the subset of listings with no reviews (3.4 percent; see Table 3, column 1). This point estimate suggests that the whole initial price gap can be accounted for by statistical discrimination. In other words, when the number of reviews tends to infinity, the price gap between a property held by a minority host and one of the same quality held by a majority will converge to zero.

In the second column of Table 4, we show the results of regression (4). The component of statistical discrimination corresponding to inaccurate beliefs is estimated to be equal to 2.5 percent. We interpret this result as evidence that roughly three-quarters (i.e., 2.5/3.4) of the gap due to statistical discrimination is driven by inaccurate beliefs and one-quarter (0.9/3.4) is driven by differences in average unobservable quality. Potential guests may be either overestimating the average quality of listings by majority hosts or underestimating the average quality of

Table 4—Nonlinear Model of log Prices as a Function of the Number of Reviews

	(1)	(2)
5 stars $\times f(K)$	0.115	
	(0.002)	
4.5 stars $\times f(K)$	0.062	
	(0.002)	
4 stars $\times f(K)$	−0.005	
	(0.003)	
$\leq$3.5 stars $\times f(K)$	−0.030	
	(0.007)	
$f(K)$		0.075
		(0.001)
Minority $\times f(K)$	0.034	0.025
	(0.006)	(0.006)
ρ	13.7	13.7
	(0.30)	(0.30)
Observations	2,474,551	2,474,551

Notes: Estimations by nonlinear least squares of equations (3) and (4). The outcome is the daily log price. Stars represent the last-known average rating for a listing. Minority is an indicator that identifies the minority host; i.e., $m_i = 1. f(K_{it}) = K_{it}/(K_{it} + \rho)$, where K_{it} is the number of reviews for listing i at time t and ρ is the number of reviews that are necessary to make up for half of the gap due to statistical discrimination. Values in row Minority $\times f(K)$ are estimates of the coefficients on the term $m_i\left(K_{it}/(K_{it} + \rho)\right)$. Under our assumptions, the interaction $m \times f(K)$ is an estimate of $-\beta_m$ (the total ethnic gap due to statistical discrimination) in column 1 and of $-\tilde{\beta}_m$ (the part of statistical discrimination due to inaccurate beliefs) in column 2. On top of covariates included in the table, we include neighborhood fixed effects, block fixed effects, and property/host characteristics. See the list of all property and host characteristics in online Appendix Table A4. Inference by block bootstrap at the listing level.

those held by minority hosts. The true average unobservable quality of minority and majority listings is very similar and creates a price gap of less than 1 percent, while the inaccurate beliefs of potential guests are responsible for most of the gap.[23] These inaccurate beliefs are corrected by new information about the quality of listings coming from reviews, which is, in practice, very different from what taste-based discrimination would generate.

We find that ρ is equal to 14. ρ can be interpreted as the number of reviews necessary to reveal half of the relevant information about the unobservables of a listing. If $\underline{p}$ is the price of a property in the absence of reviews and $\overline{p}$ the price when all the information is revealed, the price $\left(\underline{p} + \overline{p}\right)/2$ is reached in expectation after ρ reviews. On average, 14 reviews are required to correct the ethnic gap for half of the component due to statistical discrimination.

IV. Additional Results

In this section, we first present two additional pieces of evidence in support of our main empirical strategy. We show that our results are robust to more flexible or

[23] Block bootstrapping the estimation, we find that the 95 percent confidence interval (CI) of β_m is $[0.014, 0.054]$, the CI of $\tilde{\beta}_m$ is $[0.005, 0.046]$, and the CI of $\beta_m - \tilde{\beta}_m$ is $[0.006, 0.010]$.

different functional form assumptions on the relationship between log prices and the number of reviews. Second, we present results by subsamples. Finally, we provide empirical elements that lead us to argue against alternative stories that could explain why minority prices increase faster than majority ones in the absence of statistical discrimination.

A. *Robustness*

In this subsection, we present additional results that do not rely on imposing the $K/(K+\rho)$ functional form on the relationship between the number of reviews and prices. We estimate a within-listing price model where the number of reviews enters as a linear or a quadratic function:

$$(5) \qquad p_{it} = \sum_{\bar{r} \in \mathrm{supp}(\bar{r})} \mathbf{1}\{\bar{r}_i = \bar{r}\}\left(\beta_{\bar{r},1} K_{it} + \beta_{\bar{r},2} K_{it}^2\right) - m_i\left(\beta_{m,1} K_{it} + \beta_{m,2} K_{it}^2\right) + \mu_i + X_{it}\beta_x + \varepsilon_{it}.$$

If reviews matter and ratings provide some information about unobserved quality, we should have $\beta_r > \beta_{r'}$ if $r > r'$, what we have checked above with a more flexible specification. In the presence of statistical discrimination, we should have $\beta_{m,1} > 0$. The $K/(K+\rho)$ functional form also implies that the relationship between the number of reviews and prices is concave, so that $\beta_{m,2} < 0$.

Table 5 presents the results of the estimation of this model. Columns 1 and 2 show the estimation results for a linear specification in which we restrict the sample to observations with less than 40 and 60 reviews. In column 3, we present the results for the quadratic specification. The results are all consistent with those of the previous section. The higher the final rating, the faster prices grow with the number of reviews. The slope of the relationship is higher for hosts belonging to the minority group. The quadratic specification shows that the relationship is indeed concave.

B. *The Relationship between Prices and Reviews: Nonparametric Estimation*

Another way to support our empirical strategy is to show that the relationship between prices and reviews, irrespective of hosts' ethnicity, is compatible with the function $K/(K+\rho)$. Do we observe such a pattern in our data? Restricting our sample to properties held by majority hosts, we regress the log price on splines of the number of reviews interacted with the last rating (5, 4.5, 4, and 3.5 stars and less) and the full set of characteristics of the properties. The spline specification allows us to flexibly accommodate any form of the relationship between prices and the number of reviews. Formally,

$$(6) \qquad p_{it} = \sum_{\bar{r}=3.5}^{5} \mathbf{1}\{\bar{r}_i = \bar{r}\}\, s_r(K_{it}) + \mu_i + X_{it}\beta_x + \varepsilon_{it},$$

where p_{it} is the log price of property i at wave t, K is the number of reviews, X are observable characteristics of the property and the host, $s_{\bar{r}}(\cdot)$ are piecewise-linear

TABLE 5—ROBUSTNESS: LINEAR AND QUADRATIC MODELS OF PRICE WITH LISTING FIXED EFFECTS

	log price		
	(1)	(2)	(3)
3.5 stars × $K/100$	−0.145	−0.137	−0.168
	(0.058)	(0.052)	(0.085)
4 stars × $K/100$	−0.133	−0.132	−0.134
	(0.023)	(0.019)	(0.035)
4.5 stars × $K/100$	0.048	0.014	0.133
	(0.010)	(0.008)	(0.014)
5 stars × $K/100$	0.185	0.114	0.295
	(0.011)	(0.008)	(0.015)
Minority × $K/100$	0.090	0.060	0.120
	(0.036)	(0.027)	(0.045)
3.5 stars × $(K/100)^2$			0.154
			(0.154)
4 stars × $(K/100)^2$			0.082
			(0.053)
4.5 stars × $(K/100)^2$			−0.193
			(0.018)
5 stars × $(K/100)^2$			−0.315
			(0.018)
Minority × $(K/100)^2$			−0.122
			(0.053)
Samples	$K < 40$	$K < 60$	$K < 80$
Observations	1,883,500	1,996,554	2,051,820

Notes: OLS regressions with listing fixed effects. Stars represent the last-known average ratings, and K is the number of reviews. Aside from those mentioned in the table, controls include city-wave fixed effects and property characteristics (see online Appendix Table A4). Robust standard errors clustered at the property level.

splines that are specific to each level of the last rating $\bar{r}$, and μ are property fixed effects. The results of the estimation are displayed in Figure 2.

The figure shows that, depending on the last rating, prices diverge in a way that is close to the functional form predicted by our conceptual framework, displayed in online Appendix Figure A4. This result supports our assumptions that (i) reviews provide information to potential guests, (ii) hosts use reviews and information to update their prices, and (iii) the functional form between log prices and the number of reviews conditional on the last rating looks like $K/(K+\rho)$.

C. *Heterogeneity*

In Table 6, we perform the main analysis on several subsamples according to the ethnic minority group (African American versus Arabic/Muslim), the continent (North America versus Europe), and the nature of the listing (entire property versus shared property). For each sample or specification, we report in panel A the estimates of β_m and ρ from equation (3). Panel B shows the unexplained price gap on the sample of properties with no reviews.

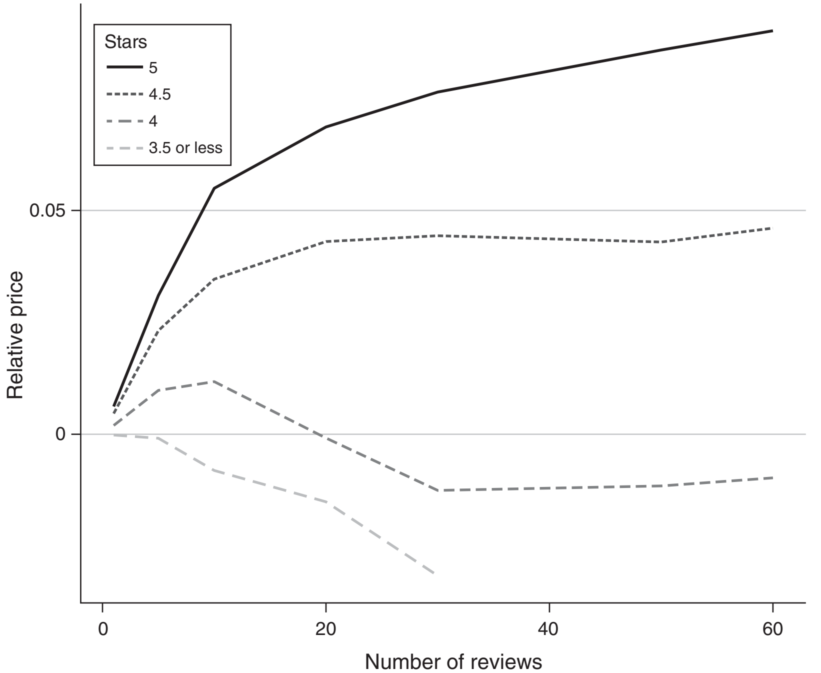

FIGURE 2. ESTIMATED PRICES WITH THE NUMBER OF REVIEWS, STRATIFIED BY THE MOST RECENT AVERAGE RATING

Notes: Equation (6) was estimated by linear regression with property fixed effects. We use linear splines with knots at 5, 10, 20, 30, and 50 reviews. The sample is restricted to listings with majority hosts. We plot the estimates $\hat{s}_r(\cdot)$ for all values of r, with the normalization $\hat{s}_r(0) = 0$. The number of observations of properties with ratings 3.5 or lower is very small when the number of reviews are higher than 30, and we do not report the corresponding estimates.

In most cases, the point estimate of β_m is of the same magnitude as the ethnic price gap for nonreviewed listings. According to our model, the ethnic price gap is maximum at zero reviews and decreases once information is revealed. Statistical discrimination seems to be higher for Black hosts than for Arabic/Muslim hosts. There is no significant difference in the extent of statistical discrimination between Europe and North America. Comparing shared flats with entire flats is intuitively interesting. A possible hypothesis is that shared flats involve a more substantial amount of interaction between hosts and guests than entire flats (where, sometimes, hosts and guests hardly meet). Our analysis shows that shared flats tend to have higher statistical discrimination than entire flats. We also find that information is more difficult to collect for shared flats (ρ being roughly twice larger) than for entire flats. This is consistent with the fact that the set of observables is larger (including how friendly the host is, for instance).

D. *Do Ethnic Groups Compete on the Same Market?*

So far, we have made the implicit assumption that minority and majority hosts compete in the same market. In this section, we investigate whether markets are also segmented: minority hosts receiving almost only guests of their own ethnicities.

Table 6—Results on Subsamples

	Full sample (1)	Arabic Muslims (2)	Blacks (3)	US Canada (4)	Europe (5)	Shared flat (6)	Entire flat (7)
Panel A. Estimation of the main model							
Minority $\times f(K)$	0.034	0.019	0.049	0.029	0.037	0.082	0.013
	(0.006)	(0.007)	(0.008)	(0.007)	(0.008)	(0.012)	(0.006)
ρ	14	14	12	10	17	22	12
	(0.3)	(0.3)	(0.3)	(0.4)	(0.4)	(0.9)	(0.3)
Panel B. Unexplained ethnic price gap (nonreviewed listings)							
Minority	−0.034	−0.036	−0.25	−0.020	−0.049	−0.034	−0.038
	(0.009)	(0.010)	(0.013)	(0.012)	(0.012)	(0.015)	(0.010)
Adjusted R^2	0.72	0.72	0.73	0.76	0.70	0.59	0.68
Minority share (percent)	6.1	3.8	3.4	8.7	4.8	7.8	5.5
Observations	351,631	342,988	270,896	119,506	232,125	99,087	252,544

Notes: In panel A, estimations by nonlinear least squares following the specification adopted in Table 4, column 1. Minority is an indicator that identifies the minority host; i.e., $m_i = 1$. $f(K_{it}) = K_{it}/(K_{it} + \rho)$, where K_{it} is the number of reviews for listing i at time t and ρ is the number of reviews that are necessary to make up for half of the gap due to statistical discrimination. Values in row Minority $\times f(K)$ are estimates of the coefficients on the term $m_i\left(K_{it}/(K_{it} + \rho)\right)$. Under our assumptions, the interaction $m \times f(K)$ is an estimate of $-\beta_m$ (the total ethnic gap due to statistical discrimination). On top of covariates included in the table, we include neighborhood fixed effects, block fixed effects and property/host characteristics. See the list of all property and host characteristics in online Appendix Table A4. Inference by block bootstrap at the listing level. In panel B, OLS regressions following the specification adopted in Table 3, column 1: daily log price on the minority dummy when the number of reviews are null, controlling for neighborhood fixed effects and block fixed effects property characteristics.

We first have to extract information about guests' ethnicities. On the website, we observe the first name of the last ten guests leaving reviews on each listing and each wave. Since we do not use the pictures provided on each guest profile, we are not able to identify Black guests. To keep a consistent definition for both hosts and guests, we restrict our analysis to the Arabic/Muslim minority group.

For each listing, we regress the share of reviews written by guests with an Arabic/Muslim first name on a dummy for the host ethnicity, controlling for the location and observable characteristics of the listing. In Table 7, we find evidence for some ethnic matching: a host with an Arabic/Muslim first name is 1 percentage point more likely to have a review from a guest with an Arabic/Muslim first name. While minority hosts seem to receive more minority guests, the magnitude of the difference shows that markets are far from being segregated.

E. *Are Reviews Ethnically Biased?*

Another way to explain our empirical results would involve the combination of taste-based discrimination and ethnically biased reviews. In this scenario, the initial ethnic gap (among listings with few reviews) would reflect taste-based discrimination. If reviews are ethnically biased, minorities would receive overall lower ratings and worse reviews than majority listings with the same quality. Therefore, minority listings with the same observables and the same ratings would be of higher quality

Table 7—Ethnic Matching between Arabic/Muslim Hosts and Arabic/Muslim Guests

	Share of Arabic/Muslim guests
Arabic/Muslim host	0.0081
	(0.0014)
Adjusted R^2	0.016
Observations	220,126

Notes: OLS regression. Aside from the dummy Arabic/Muslim Host, controls include city-wave fixed effects, neighborhood fixed effects, block fixed effects, property characteristics (see online Appendix Table A4), log price, number of reviews, and ratings. Standard errors are clustered at the property level.

than majority listings. Prices of listings owned by minorities conditional on observable characteristics and ratings would increase faster than prices of majority listings.

A key ingredient of this scenario is that reviews are ethnically biased. In this subsection, we show that minority hosts do not receive significantly better or worse reviews from minority guests than they do from majority guests. We read this result as an argument against the hypothesis that reviews are biased. To investigate this question, we must build, for each listing i and wave t, the ratings corresponding to the new reviews between t and $t-1$. This step is necessary because the rating we observe at date t, $\bar{r}_{it}$, is the average rating over all the reviews obtained by the listing until date t. We infer $\tilde{r}_{it}$, the average rating over reviews obtained between $t-1$ and t, from $\bar{r}_{it}$, $\bar{r}_{it-1}$, K_{it} (the total number of reviews at t), and K_{it-1}:

$$\tilde{r}_{it} = \frac{\bar{r}_{it} \cdot K_{it} - \bar{r}_{i,t-1} \cdot K_{i,t-1}}{K_{it} - K_{i,t-1}}.$$

We then estimate

$$\tilde{r}_{it} = \alpha \tilde{g}_{it}^{m} + \gamma m_i \tilde{g}_{it}^{m} + X_{it}\beta + \mu_i + \varepsilon_{it},$$

where $\tilde{g}_{it}^{m}$ is the share of guests between $t-1$ and t that belong to the minority group and μ_i is a listing-specific fixed effect. As in Section IVD, we exclude Blacks from the analysis because we are not able to identify them among the guests. In this regression, γ can be interpreted as the difference between the ratings given by minority and majority guests to minority listings. Restricting the sample to observations with new guests between waves, Table 8 shows that the coefficient of the interaction term is nonsignificant and small in magnitude: minority guests do not seem to give better reviews to minority hosts.

F. *Ethnic Differences in Property Upgrading*

Minority hosts might react to lower demand by improving the quality of their listing to a larger extent than majority hosts. In this case, we would also observe that minority prices increase faster than majority ones. Hosts can upgrade their property through both observable and unobservable characteristics.

Table 8—Average Rating, Depending on Hosts' and Guests' Ethnicity

Average rating over reviews received between $t-1$ and t	
Share of minority among new guests	0.000
	(0.025)
Minority host × share of minority among new guests	0.008
	(0.007)
Adjusted R^2	0.072
Observations	954,361

Notes: OLS regressions with listings fixed effects. The outcome is $\tilde{r}_{it}$, the average rating over reviews obtained between $t-1$ and t. Aside from those mentioned in the table, controls include city-wave fixed effects and property characteristics (see online Appendix Table A4). Robust standard errors clustered at the property level.

We exploit the information about observable characteristics of a listing and test whether minority hosts tend to change these observables in a way that improves the perceived quality of their listing. First, we estimate a hedonic price regression: we regress the log price on property characteristics, controlling for location and city-wave fixed effects, on the majority population. We use the estimated coefficients of this regression to predict the log price corresponding to all properties for each period as a function of the observables, $\hat{p}(X_{it})$. Second, we use the predicted price $\hat{p}(X_{it})$ as the outcome in the following model:

$$\hat{p}(X_{it}) = \sum_{\bar{r}=3.5}^{5} b_r \mathbf{1}\{\bar{r}_i = \bar{r}\} \frac{K_{it}}{K_{it}+\rho} - b_m m_i \frac{K_{it}}{K_{it}+\rho} + \mu_i + \varepsilon_{it}. \tag{7}$$

If minorities upgrade their properties more and sooner than their majority counterparts do, b_m should be negative.

Upgrading may also come from characteristics that are not directly observable by the econometrician. We test for this by looking at the differential evolution of the word count of the listing description, the number of pictures displayed on the listing's page, and the number of pictures taken by a professional photographer. We run a regression very close to the previous one, except that we now control for observables on the right-hand side:

$$Y_{it} = \sum_{\bar{r}=3.5}^{5} b_r \mathbf{1}\{\bar{r}_i = \bar{r}\} \frac{K_{it}}{K_{it}+\rho} - b_m m_i \frac{K_{it}}{K_{it}+\rho} + \mu_i + X_{it} b_x + \varepsilon_{it}, \tag{8}$$

where Y_{it} is the word count or the number of pictures of listing i at date t. Again, we expect b_m to be negative if minority hosts upgrade their properties more than majority hosts.

Table 9 shows the results of the estimations of equations (7) and (8). In column 1, the point estimate of b_m is significantly negative but very small. This suggests that upgrading on the observable characteristics plays a negligible role in explaining our main results. In all other columns, the point estimates of b_m are small and insignificant. Minority hosts do not seem to have upgraded their listings differentially from majority hosts in terms of the listing description or pictures.

Table 9—Differential Upgrading

	Pred. log price	Word count	Pictures	Pro. pictures
5 stars $\times f(K)$	0.007	92.368	7.516	5.867
	(0.000)	(3.028)	(0.105)	(0.102)
4.5 stars $\times f(K)$	0.007	80.547	6.400	4.788
	(0.000)	(3.183)	(0.097)	(0.094)
4 stars $\times f(K)$	0.005	55.212	5.743	2.714
	(0.000)	(5.635)	(0.183)	(0.166)
$\leq$3.5 stars $\times f(K)$	0.003	30.629	5.359	−0.096
	(0.001)	(8.808)	(0.325)	(0.318)
Minority $\times f(K)$	0.001	7.503	−0.150	0.283
	(0.001)	(8.419)	(0.299)	(0.287)
Adjusted R^2	0.995	0.218	0.082	0.206
Observations	2,474,551	2,474,551	2,474,551	2,474,551

Notes: OLS regressions of equations (7) and (8), using the estimated $\hat{\rho} = 13.6$ and $f(K) = K/(K + \hat{\rho})$, where K is the number of reviews and ρ is the number of reviews that are necessary to make up for half of the gap due to statistical discrimination. Stars represent the last-known average ratings. Minority is an indicator that identifies the minority host; i.e., $m_i = 1$. Values in row Minority $\times f(K)$ are estimates of the coefficients on the term $m_i\left(K_{it}/(K_{it} + \rho)\right)$. Under our assumptions, the interaction $m \times f(K)$ is an estimate of $-b_m$. In column 1, the outcome is the predicted log price based on observable characteristics of the listing. In column 2, the outcome is the word count of the listing description. In column 3, the outcome is the number of pictures on the listing profile. In column 4, the outcome is the number of pictures taken by professionals on the listing profile. On top of the covariates included in the table, we include property/host characteristics (except in column 1). Robust standard errors clustered at the listing level.

G. *Do Minority Hosts Set Prices That Are Too Low Initially?*

Another way to rationalize our results is that minority hosts might be less familiar with this market than majority ones are. If minorities are initially more pessimistic about the potential of their listings than majority hosts, they will set up lower initial prices. As information about their quality comes back to them (through reviews or the number of transactions), they would revise up their prices more quickly than the majority, which would generate differential dynamics in prices. This would explain our empirical results without the existence of statistical discrimination.

This story entails that minority hosts should get more transactions initially and that demand for ethnic listings should decrease with the number of reviews relative to nonminority listings. We show that this is not the case. As we do not directly observe the number of transactions, we use the number of new reviews between two waves as a proxy. We estimate our nonlinear model using this difference (or the difference in logs) as the outcome:

$$(9) \qquad \Delta K_{it} = \sum_{\bar{r}=3.5}^{5} b_r \mathbf{1}\{\bar{r}_i = \bar{r}\} \frac{K_{it}}{K_{it} + \rho} - b_m m_i \frac{K_{it}}{K_{it} + \rho} + \mu_i + X_{it} b_x + \varepsilon_{it},$$

where $\Delta K_{it} = K_{t+1} - K_{it}$. We want to measure whether minorities accumulate more or fewer reviews (proxy for the number of transactions) as the number of reviews increases. If minorities are overpessimistic and learn about their type gradually, we

Table 10—Ethnic Differentials in the Accumulation of Reviews over Time

	ΔK	$\Delta \log K$
5 stars $\times f(K)$	0.086	−2.27
	(0.002)	(0.026)
4.5 stars $\times f(K)$	0.071	−2.27
	(0.002)	(0.023)
4 stars $\times f(K)$	0.030	−2.52
	(0.002)	(0.036)
$\leq$3.5 stars $\times f(K)$	−0.005	−3.20
	(0.003)	(0.073)
Minority $\times f(K)$	−0.018	−0.103
	(0.005)	(0.068)
Adjusted R^2	0.409	0.199
Observations	2,253,612	1,901,981

Notes: OLS regressions of equation (9), using the estimated $\hat{\rho} = 13.6$ and $f(K) = K/(K+\hat{\rho})$, where K is the number of reviews and ρ is the number of reviews that are necessary to make up for half of the gap due to statistical discrimination. Stars represent the last-known average ratings. Minority is an indicator that identifies the minority host; i.e., $m_i = 1$. Values in row Minority $\times f(K)$ are estimates of the coefficients on the term $m_i\left(K_{it}/(K_{it}+\rho)\right)$. Under our assumptions, the interaction $m \times f(K)$ is an estimate of $-b_m$. In column 1, the outcome is the difference between two dates in the number of reviews. In column 2, the outcome is the difference in the log number of reviews. On top of the covariates included in the table, we include property/host characteristics. Robust standard errors clustered at the listing level.

should observe that minorities have a more decreasing pattern of the number of transactions per period compared to majority hosts. In this case, b_m should be positive.

In Table 10, we find that the coefficient associated to minority hosts is negative (which entails a positive b_m), significant, and small in the first column and insignificant in the second column. Taken at face value, the magnitude of the first coefficient suggests that minority hosts would initially get 0.02 more reviews than majority hosts. While the sign of the coefficient is consistent with minorities being pessimistic about their perspectives on the website, the magnitude of the coefficient suggests that it should be a minor contributor to the overall story.

V. Conclusion

This paper documents that Airbnb hosts who belong to an ethnic minority experience a 3 percent price penalty when differences in locations and observable characteristics are accounted for. Taking advantage of the longitudinal nature of our data, we show that the ethnic gap can be fully explained by statistical discrimination. About one-quarter of the gap comes from differentials in average unobservable quality across groups. Three-quarters can be attributed to the fact that potential guests hold inaccurate beliefs about the average quality of properties held by minority hosts compared to majority hosts.

We can draw several conclusions from these findings. First, aside from the issues inherent to any online feedback system, the one featured by this online platform is effective in supplying useful information to potential guests. In the absence of such a feedback system, the ethnic price gap would be higher than its current value.

Second, aside from gains in efficiency, improving the feedback system would also contribute to reducing ethnic price gaps. Third, minority hosts are still largely penalized by the existence of inaccurate beliefs that potential guests hold against them even though the review system mitigates their influence.

We believe that the evidence provided in this paper is relevant to the current debate about discrimination in online platforms. While there is no reason to make ethnicity particularly salient on these platforms, policies consisting of concealing more information about actors' identity may backfire if ethnic gaps are due to statistical discrimination. We see our results as advocating another way to reduce ethnic gaps: disclosing more abundant and reliable information about candidates, sellers, or hosts. As discussed by Shaw, Horton, and Chen (2011), it remains to be understood how platforms can adequately incentivize reviewers to provide informative, unbiased, and relevant reviews. Further research is required to understand how interventions on information disclosure affects ethnic gaps.

On Airbnb, like on many other online marketplaces, interactions between agents are limited. While we have no evidence about how our results can generalize to other platforms, online or not, they are consistent with those obtained by Pallais (2014) and Agrawal, Lacetera, and Lyons (2016) on the online platform oDesk (now Upwork). Pallais (2014) finds that providing public information about workers' abilities has, on average, a positive effect on workers' probability to be hired. Agrawal, Lacetera, and Lyons (2016) find that standardized information about work performed on the platform disproportionately benefits less developed country contractors relative to developed country ones. The approach we follow in this paper may be adapted to study ethnic discrimination on several other widely used online platforms, including labor markets.

While our identification strategy allows us to pin down statistical discrimination (and the share of it that is due to inaccurate beliefs), we cannot disentangle other factors like taste-based discrimination, ethnic differentials in characteristics that are observable to potential guests but not to econometricians (e.g., pictures contents), or hosts' opportunity cost of time. While statistical discrimination (and inaccurate beliefs) appears to explain most of the gap, taste-based discrimination could be offset by differentials in characteristics, for instance. Therefore, we cannot rule out the existence of taste-based discrimination on Airbnb. Another caveat is that the analysis is made conditional on location. Because ethnic minorities tend to live in neighborhoods that are less valued by potential guests,[24] minority hosts suffer in reality from larger price gaps than those computed conditional on location.

REFERENCES

Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons. 2016. "Does Standardized Information in Online Markets Disproportionately Benefit Job Applicants from Less Developed Countries?" *Journal of International Economics* 103: 1–12.

Aigner, Dennis J., and Glen G. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *Industrial and Labor Relations Review* 30 (2): 175–87.

[24] We find that locations where the share of ethnic minority among Airbnb hosts is 1 standard deviation (i.e., 9.6 percentage points) higher are valued 5.5 percent less in prices.

Airbnb. 2014–2017. Airbnb, Inc. https://www.airbnb.com/ (accessed between June 2014 and November 2017).

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113 (1): 151–84.

Altonji, Joseph G., and Charles R. Pierret. 2001. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics* 116 (1): 313–50.

Antonovics, Kate, Peter Arcidiacono, and Randall Walsh. 2005. "Games and Discrimination: Lessons from 'The Weakest Link'." *Journal of Human Resources* 40 (4): 918–47.

Anwar, Shamena. 2012. "Testing for Discrimination: Evidence from the Game Show *Street Smarts*." *Journal of Economic Behavior and Organization* 81 (1): 268–85.

Arnold, David, Will Dobbie, and Crystal S. Yang. 2018. "Racial Bias in Bail Decisions." *Quarterly Journal of Economics* 133 (4): 1885–1932.

Arrow, Kenneth J. 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets,* edited by Orley Ashenfelter and Albert Rees, 3–33. Princeton, NJ: Princeton University Press.

Autor, David H. 2001. "Wiring the Labor Market." *Journal of Economic Perspectives* 15 (1): 25–40.

Autor, David H. 2009. "Introduction." In *Studies of Labor Market Intermediation*, edited by David H. Autor, 1–23. Chicago: University of Chicago Press.

Autor, David H., and David Scarborough. 2008. "Does Job Testing Harm Minority Workers? Evidence from Retail Establishments." *Quarterly Journal of Economics* 123 (1): 219–77.

Bagues, Manuel F., and Mauro Sylos Labini. 2009. "Do Online Labor Market Intermediaries Matter? The Impact of AlmaLaurea on the University-to-Work Transition." In *Studies of Labor Market Intermediation*, edited by David H. Autor, 127–154. Chicago: University of Chicago Press.

Bayer, Patrick, Marcus Casey, Fernando Ferreira, and Robert McMillan. 2017. "Racial and Ethnic Price Differentials in the Housing Market." *Journal of Urban Economics* 102: 91–105.

Becker, Gary. 1957. *The Economics of Discrimination.* 2nd ed. Chicago: University of Chicago Press, 1971.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94 (4): 991–1013.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. 2019. "The Dynamics of Discrimination: Theory and Evidence." *American Economic Review* 109 (10): 3395–3436.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2016. "Stereotypes." *Quarterly Journal of Economics* 131 (4): 1753–94.

Brown, Meta, Elizabeth Setren, and Giorgio Topa. 2016. "Do Informal Referrals Lead to Better Matches? Evidence from a Firm's Employee Referral System." *Journal of Labor Economics* 34 (1): 161–209.

Castillo, Marco, and Ragan Petrie. 2010. "Discrimination in the Lab: Does Information Trump Appearance?" *Games and Economic Behavior* 68 (1): 50–59.

Charles, Kerwin Kofi, and Jonathan Guryan. 2008. "Prejudice and Wages: An Empirical Assessment of Becker's *The Economics of Discrimination*." *Journal of Political Economy* 116 (5): 773–809.

Charles, Kerwin Kofi, and Jonathan Guryan. 2011. "Studying Discrimination: Fundamental Challenges and Recent Progress." *Annual Review of Economics* 3 (1): 479–511.

Cui, Ruomeng, Jun Li, and Dennis J. Zhang. 2020. "Reducing Discrimination with Reviews in the Sharing Economy: Evidence from Field Experiments on Airbnb." *Management Science* 66 (3): 1071–94.

Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2018. "Measuring Bias in Consumer Lending." Princeton University Industrial Relations Section Working Paper 623.

Doleac, Jennifer L., and Luke C.D. Stein. 2013. "The Visible Hand: Race and Online Market Outcomes." *Economic Journal* 123 (572): F469–92.

Edelman, Benjamin, and Michael Luca. 2014. "Digital Discrimination: The Case of Airbnb.com." Harvard Business School Working Paper 14-054.

Ewens, Michael, Bryan Tomlin, and Liang Choon Wang. 2014. "Statistical Discrimination or Prejudice? A Large Sample Field Experiment." *Review of Economics and Statistics* 96 (1): 119–34.

Farber, Henry S., and Robert Gibbons. 1996. "Learning and Wage Dynamics." *Quarterly Journal of Economics* 111 (4): 1007–47.

Farronato, Chiara, and Andrey Fradkin. 2018. "The Welfare Effects of Peer Entry in the Accommodation Market: The Case of Airbnb." NBER Working Paper 24361.

Fershtman, Chaim, and Uri Gneezy. 2001. "Discrimination in a Segmented Society: An Experimental Approach." *Quarterly Journal of Economics* 116 (1): 351–77.
Fradkin, Andrey. 2017. "Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2939084.
Fradkin, Andrey, Elena Grewal, and David Holtz. 2018. "Reciprocity and Unveiling in Two-sided Reputation Systems: Evidence from an Experiment on Airbnb." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2939064.
Hawramani, Ikram. 2015. *The Arabic Baby Name Book: More Than 5000 Names for Boys and Girls.* Independent publisher.
Heckman, James J. 1998. "Detecting Discrimination." *Journal of Economic Perspectives* 12 (2): 101–16.
Horton, John J. 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment." *Journal of Labor Economics* 35 (2): 345–85.
Jouniaux, Léo. 2001. *Les 20,000 Plus Beaux Prénoms du Monde*. Paris: Hachette.
Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 109 (1): 203–32.
Lang, Kevin, and Jee-Yeon K. Lehmann. 2012. "Racial Discrimination in the Labor Market: Theory and Empirics." *Journal of Economic Literature* 50 (4): 959–1006.
Laouénan, Morgane, and Roland Rathelot. 2022. "Replication data for: Can Information Reduce Ethnic Discrimination? Evidence from Airbnb." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E120078V1.
Levitt, Steven D. 2004. "Testing Theories of Discrimination: Evidence from *Weakest Link*." *Journal of Law and Economics* 47 (2): 431–52.
List, John A. 2004. "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field." *Quarterly Journal of Economics* 119 (1): 49–89.
Neumark, David. 2018. "Experimental Research on Labor Market Discrimination." *Journal of Economic Literature* 56 (3): 799–866.
Oreopoulos, Philip. 2011. "Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes." *American Economic Journal: Economic Policy* 3 (4): 148–71.
Pallais, Amanda. 2014. "Inefficient Hiring in Entry-Level Labor Markets." *American Economic Review* 104 (11): 3565–99.
Pallais, Amanda, and Emily Glassberg Sands. 2016. "Why the Referential Treatment? Evidence from Field Experiments on Referrals." *Journal of Political Economy* 124 (6): 1793–1828.
Phelps, Edmund S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (4): 659–61.
Pope, Devin G., and Justin R. Sydnor. 2011. "What's in a Picture? Evidence of Discrimination from Prosper.com." *Journal of Human Resources* 46 (1): 53–92.
Rubinstein, Yona, and Dror Brenner. 2014. "Pride and Prejudice: Using Ethnic-Sounding Names and Inter-Ethnic Marriages to Identify Labour Market Discrimination." *Review of Economic Studies* 81 (1): 389–425.
Shaw, Aaron D., John J. Horton, and Daniel L. Chen. 2011. "Designing Incentives for Inexpert Human Raters." In *CSCW '11: Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work,* 275–84. New York: Association for Computing Machinery.
Stanton, Christopher T., and Catherine Thomas. 2018. "Experience Markets: An Application to Outsourcing and Hiring." Harvard Business School Working Paper 18-096.
Wozniak, Abigail K. 2015. "Discrimination and the Effects of Drug Testing on Black Employment." *Review of Economics and Statistics* 97 (3): 548–66.
Zussman, Asaf. 2013. "Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars." *Economic Journal* 123 (572): F433–68.