Ethnic Discrimination on an Online Marketplace of Vacation Rentals*

Morgane Laouénan†
Roland Rathelot‡

April 29, 2016
Preliminary, Please do not quote

Abstract

We use data from an online marketplace of vacation rentals collected in 19 major cities in North America and Europe to measure discrimination against ethnic-minority hosts. For our purpose, this market has three interesting features: the existence of a detailed reviewing system, the high frequency of transactions and the panel dimension of the data. We take advantage of these features to measure the influence of better signals on prices and provide a credible measure of the extent of statistical discrimination, following a strategy à la Altonji and Pierret (2001). Hosts belonging to an ethnic minority charge 15.5% less than majority hosts in the same cities. Controlling for a rich set of characteristics reduces the ethnic price gap to 3.3%. An additional review increases the price more for minority than for majority hosts. Estimating the parameters of a theoretical pricing model, we find that statistical discrimination can account for the whole ethnic price gap.

Keywords: ethnic discrimination, statistical discrimination, rental market.
JEL: J15, L85.

*We would like to thank David Autor, Manuel Bagues, Leah Boustan, Clément de Chaisemartin, Raj Chetty, Bruno Decreuse, Christian Dustmann, Nathan Hendren, Attila Lindner, Kevin Lang, Victor Lavy, Andrea Moro, David Neumark, Barbara Petrongolo, Fabian Waldinger, Natalia Zinovyeva and participants in seminars at UC Irvine, UCLA (CCPR), UC Santa Barbara, Institut d'Economia de Barcelona (IEB), Umea, Paris Sorbonne, Warwick, DIW Berlin, Sciences Po, Université catholique de Louvain, and at the Workshop on Gender and Ethnic Differentials (Aix-Marseille, 2014) and the Workshop Topics in Labor Economics (Nizza Monferrato, 2015) for fruitful discussions and comments. The views expressed herein are those of the authors and do not necessarily reflect those of any institution. We gratefully acknowledge financial support from the Centre for Competitive Advantage in the Global Economy (Warwick) and from Sciences-Po LIEPP.

†CNRS, Centre d'Economie de la Sorbonne, Sciences-Po LIEPP

‡University of Warwick, CEPR, CAGE

1 Introduction

While ethnic discrimination is a pervasive phenomenon in most markets and most countries, understanding which mechanisms are at work is necessary to design efficient policies. In their recent reviews, Charles and Guryan (2011) and Lang and Lehmann (2012) stress that uncovering discrimination mechanisms is crucial and that empirical attempts are rare and not conclusive. This paper takes advantage of the features of a major online marketplace for short-term rentals to measure the share of statistical discrimination in the ethnic price gap which prevails on that market.

On the online marketplace we study, hosts list their properties, set the daily price and provide information about themselves (at least first name and picture) and their properties (precise location, equipment, local amenities, pictures...). Potential guests propose to book the property at given dates at the price set by the host. In this paper, we study the differential between the prices set by hosts who belong to an ethnic minority and those set by majority hosts. Controlling for a large set of observable characteristics accounts for more than 75% of the raw gap but the "unexplained" gap remains significant.[1] In this paper, we ask whether this unexplained gap is driven by statistical discrimination or another mechanism?

While taste-based discrimination stems from the existence of racial preferences or an aversion towards 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 gaps vary with the amount of information on the market (Farber and Gibbons, 1996; Altonji and Pierret, 2001). Many papers have attempted to apply this method to wage gaps on the labor market. There are however many challenges. First, employment spells are typically long and worker's productivity evolve over time in a way that may depend on the quality of the match. Second, there is typically no good measure of the quantity of information available to employers. Experience (or age) is usually used to proxy this quantity, which is problematic as human capital also varies with age (in a potentially complex and asymmetric way).

We adapt the Altonji and Pierret (2001) approach to our setting, where the quantity and quality of information about a property is well measured. In contrast with the labor market, the short-term rental market is well suited for testing statistical discrimination because (i) transactions happen frequently, compared to changes in the intrinsic quality of the property, (ii) the evolution of the number of reviews and

[1] Edelman and Luca (2014) is the first paper to document a Black-White price gap in New York City in such a setting.

ratings can be observed, (iii) large-sample and longitudinal data are available. New properties start with self-reported information about characteristics. Then, guests who have stayed in a property are allowed to let a quantitative rating and a qualitative assessment of both the accommodation and the host. As time goes, the number of reviews grows, so that more and more information is available to potential guests over time.

We rely on a simple conceptual framework where properties' quality is partly unobservable. When a property has no review, potential guests can only infer unobservable quality using hosts' ethnicity. When a property has reviews, potential guests aggregate the content of reviews and ethnicity to form the best possible guess about the property's observable quality. In case of statistical discrimination, the price gap should decrease with the number of reviews and tend to zero, conditional on observables and on the measure of quality provided by reviews. If the price gap is due to taste-based discrimination or to 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.

Our dataset include daily prices, hosts' and apartments' characteristics, as well as associated reviews. We collected the data relating to 400,000 properties, corresponding to apartments to rent in 19 cities in North America and Europe. 20 waves of data, collected between June 2014 and July 2015 form an unbalanced panel of 3,500,000 observations. The ethnic minority groups we consider are hosts with Arabic/African first names and hosts with pictures coded as African-American (North America only).

We find that the within-city raw ethnic price gap is around 15.5%. The set of observable characteristics about the property (including its precise location) is rich and explains more than 67% of the variance of the prices. Controlling for ethnic differences in these characteristics reduces the ethnic price gap to 3.3%. In cross-section, we document that this unexplained ethnic gap decreases with the number of reviews and is close to zero and insignificant in the subsample of properties with more than than 50 reviews. We then use the longitudinal dimension of our data and document that, as predicted by the theoretical framework, prices increase faster with the number of reviews when the host belongs to an ethnic minority. Finally, we estimate the parameters of the price equation of the model using longitudinal variations in prices and the number of reviews. Our results point out that the ethnic price gap can be entirely accounted for by statistical discrimination.

Our paper contributes to the growing but largely inconclusive literature on the sources of discrimination. On the U.S. labor market, Altonji and Pierret (2001) pioneered the methodology but find little evidence for statistical discrimination in wages on the basis of ethnicity. A related strand of literature uses the fact that the relevant outcome is perfectly observed ex post. Using data from a peer-to-peer lending website, Pope and Sydnor (2011) find that African-American lenders face higher interest rates and lower borrowing probabilities. However, blacks have higher default rates so that net returns on loans made to African-Americans is lower. According to the authors, these results would be consistent with accurate statistical discrimination against African-Americans and taste discrimination against whites. Similarly, Knowles et al. (2001) show that vehicles of African-Americans are more often searched by the police and that statistical discrimination explains more than the observed gap. Using data from television game shows, Anwar (2012) find that non-black contestants erroneously believe that Afro-Americans have lower skill levels while Levitt (2004) and Antonovics et al. (2005) find no evidence of discrimination.

The amount and nature of information available to discriminatory agents can also be manipulated experimentally. In their correspondence studies on the U.S. and Canadian labor markets, Bertrand and Mullainathan (2004) and Oreopoulos (2011) find that adding information or enhancing resumes do not benefit minority applicants. Conversely, on the online rental apartment market, Ewens et al. (2014) find the response to differential quality varies in a way which is consistent with statistical discrimination. A potential limitation of this approach is related to the critique by Heckman (1998): why would someone conceal a favorable piece of information? Even if the amount of information in the resume is randomized, its absence should be interpreted by employers (or customers) as information.

The heterogeneity in agents' prejudice, whether revealed or assumed, is sometimes used to infer which source of discrimination is more prevalent. Bayer et al. (2012) show that the African-American and Hispanic home-buyers pay 3% premiums on the U.S. housing market. The premium is the same when the seller is himself African-American or Hispanic, suggesting statistical discrimination. Zussman (2013) finds that the discrimination towards Arabs on an online market for used cars in Israel is not related to sellers' revealed attitudes towards Arabs. Doleac and Stein (2013) show that online iPod ads featuring dark-skinned hands received fewer and lower offers than light-skinned-hand ones. Outcomes being poorer in thin markets and those with higher racial isolation and crime is consistent with statistical discrimination.

Finally, other approaches have been used to separate sources of discrimination. Charles and Guryan (2008) introduce an indirect test of the Becker 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. Wozniak (2015) shows how some policy that affects a relevant dimension of the unobservables (in her case, drug use) can provide evidence of statistical discrimination against low-skilled African-American men: she uses time variation in drug-testing legislation. Experimental evidence can be complemented by lab games to separate discrimination mechanisms. In the case of the sportscard market, List (2004) finds that the lower offers received by minorities were rather explained by statistical discrimination.[2]

The main contribution of this paper is to apply a test à la Altonji and Pierret (2001) in a context and with data that makes it more credible. The size of the data allows us to provide a precise assessment while the availability of longitudinal high-frequency observations and a reviewing system provide a unique opportunity to test for statistical discrimination. Potential guests typically search for a few hours, stay at the property for a few days and fill up reviews after a couple of extra days. The fact that the market is centralized is also a precious advantage, as the same set of information (prices, characteristics and reviews) are observed by all agents, and by the econometrician. Further, our study is the first one to study ethnic discrimination on the rental market at a large scale, in 19 cities spread in 8 countries in Europe and North America. This online marketplace is also relevant in itself from an economic point of view: launched in 2008, the website proposes more than 800,000 listings in 190 different countries and claims to have served over 10 million guests. We also contribute to the growing literature on the role of information provided by online market intermediaries on markets' outcomes.[3]. Our paper is also 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 and to Agrawal et al. (2014), who find that standardized information about work performed on the platform disproportionately benefits less-developed-country contractors, relative to developed-country ones. Finally, our results are consistent, in negative, with those obtained by Behaghel et al. (2015), who show that setting up an experimental anonymized-resume policy for some vacancies had counter-productive consequences on the hiring rate of ethnic minorities.

The next section presents the context, the data and the first empirical evidence about ethnic price gaps. In the third section, we present our conceptual framework and its predictions. In the fourth section, we provide the empirical results about statistical discrimination. A fifth section provides robustness checks and discusses alternative explanations. Section six concludes.

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

[3]A substantial share of the literature deal with the role of labor-market intermediaries; see e.g. Autor (2001, 2009); Bagues and Labini (2009); Pallais (2014); Horton (2015); Pallais and Sands (2015); Stanton and Thomas (2015); Brown et al. (2016)

2 Context and Data

2.1 Description of the platform

This marketplace gathers hosts looking for opportunities to let their properties and 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. The information about properties and hosts are displayed to potential guests in a standardized way, in order to ease comparison. In practical terms, potential guests usually start by typing the city where and when they want to stay on the search engine. They can filter the results of the search according to the price, or other characteristics (like the number of accommodates, the type of room, the property type, the number of bedrooms...). At that stage, they would typically obtain a list of results, sorted by relevance, with basic information, among which the daily price, a picture of the property, a thumbnail of the host and the rating. When they click on one of the listing, 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. See Appendix A for a screenshot of a listing.

Hosts can revise the price of their properties at any moment. However, the system does not allow negotiation. Once the potential guest has decided which place he preferred among those available during the period selected, he can choose to offer a bid. The bid is then in the hands of the host that can decide, without any justification, to accept or reject the offer, based on the information he has about the potential guest.[4] There is no way for the two parties to communicate (to bargain on the price, for instance) before the deal is done. If the bid is rejected, the potential guest can look for another place. This rejection is not reported on his profile. If the bid is accepted, the deal is done and there is no way to modify its terms.[5] However, the potential guest can decide to cancel his booking. In this case, the terms of the cancellation policy (specified on the listing) apply: depending on the flexibility of the policy, 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 signals on the host's profile that he has cancelled a deal.

Overall, we consider that potential guests are price-takers. Using a simple model of supply and demand, we consider that the existence of discrimination towards hosts,

[4] Rejection happens often; see empirical evidence about search frictions on Airbnb.com in Fradkin (2015).

[5] 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 et al. (2016) for a field study about discrimination against potential tenants.

which triggers a shift in demand, should translate into lower prices. We formalize this idea in the section dedicated to the conceptual framework.

2.2 Data

We collected the information from the publicly-available webpages of the marketplace. Specifically, we store all information visible on the first page of the listing: price that the host is asking, characteristics of the listing, characteristics of the host and all associated reviews and ratings.

We focus on the 19 cities in Europe, Canada and the U.S. with the largest number of listings and a significant share of ethnic minorities: London, Paris, Madrid, Barcelona, Rome, Milan, Florence, Amsterdam, Berlin, Marseille, Vancouver, Toronto, Montreal, Boston, New York City, Miami, Chicago, San Francisco and Los Angeles[6]. We repeated the collection process every 2-3 weeks between June 2014 and June 2015, obtaining 20 waves in total. See the collection date of each wave in Table 11 in Appendix. Our sample include 400,000 distinct properties. The panel is unbalanced: some properties enter the system while others exit.

Table 1 presents the characteristics of the properties and the hosts. The left column display the mean of each characteristics in the full sample, while the right column focuses on the subsamples of active listings, that have gained at least one review over the observation period. Most properties are apartments and the entire place is let in 70% of cases. Properties are rather small, with 1.3 bedrooms on average, and they can host on average three guests. There are no sizable differences between the whole sample and the set of the active listings. Most places propose wireless connection, heating and a washer while some amenities (e.g. cable TV, dryer, or parking space) are less frequent. The presence of a doorman, a gym, a hot tub, or a pool is rare. Some properties add a cleaning fee and charge for additional people. Most do not allow pets or smoking.

Some information about hosts is available (via their profile pages). Aside from the first name, a picture and a free-text description, guests know whether they have other properties and when they joined the platform. They can also guess whether the listing is managed by a couple or not via the presence of two names. Most hosts have only one property, have joined relatively recently (since 2012) and are not considered being in couple.

[6]See Table 10 in Appendix for the number of observations by city.

Table 1: Summary statistics: Property & host characteristics

	Full Sample	Active Listings
Type of property		
Shared Flat	0.332	0.299
Entire Flat	0.668	0.701
Flat	0.836	0.846
House	0.107	0.103
Loft	0.017	0.019
Size		
Person Capacity	3.114	3.214
Nber bedrooms	1.246	1.244
Nber bathrooms	1.163	1.153
Terrace or Balcony	0.231	0.273
Type of bed		
Couch	0.006	0.006
Airbed	0.003	0.003
Sofa	0.030	0.032
Futon	0.011	0.011
Real Bed	0.949	0.948
Amenities		
Cable TV	0.353	0.364
Wireless	0.920	0.943
Heating	0.905	0.926
AC	0.381	0.384
Elevator	0.358	0.347
Wheelchair Accessible	0.100	0.104
Doorman	0.104	0.096
Fireplace	0.080	0.082
Washer	0.722	0.725
Dryer	0.400	0.404
Parking	0.186	0.185
Gym	0.073	0.066
Pool	0.063	0.055
Buzzer	0.384	0.405
Hot Tub	0.074	0.070
Services		
Breakfast served	0.089	0.094
Family/Kids Friendly	0.443	0.469

(Continued on next page)

Table 1: Summary statistics: Property & host characteristics

Suitable for events	0.050	0.054
Rules & Extras		
Additional People	0.572	0.723
Price per Additional People	7.310	8.544
Cleaning price	27.918	30.644
Smoking Allowed	0.289	0.316
Pets Allowed	0.281	0.312
Host Characteristics		
In couple	0.055	0.070
Has multiple properties	0.348	0.369
Member since 2008	0.001	0.001
Member since 2009	0.008	0.009
Member since 2010	0.028	0.033
Member since 2011	0.095	0.108
Member since 2012	0.190	0.212
Member since 2013	0.253	0.269
Member since 2014	0.300	0.295
Member since 2015	0.099	0.062
Number of languages spoken	1.279	1.434
N	409,737	214,530

Notes: Active listings correspond to listings which receive at least one review over the observation period.

The distribution of the number of waves during which we observe each property is in the left panel of Figure 1. It shows that 11% of listings are observed in all waves and half of listings are observed in more than 6 waves. On average, a property is observed 7 times over the period. The number of listings observed per wave is displayed in the right panel of Figure 1.

Figure 1: Number of observations by listing and of listings per wave

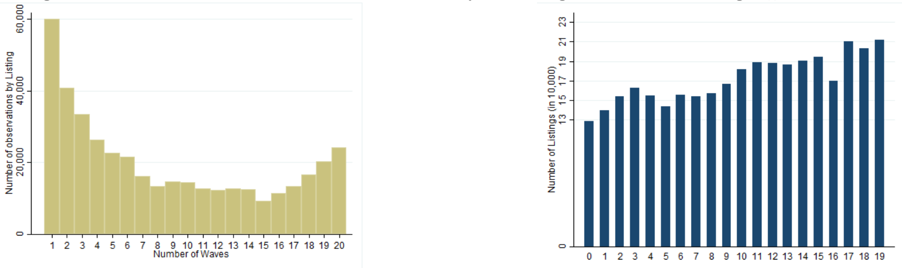

Figure 2 shows the distribution of daily prices. There is a lot of variation in prices across properties. To reduce the influence of outliers, we drop 1% of the observations of the top and the bottom of the price distribution. The first quarter is 51 euros, the median 104 euros and the third quarter 247 euros. The skewness of the distribution implies that the mean price is 139 euros. The daily price varies across cities and according to the amenities of the listing (number of accommodates, bedrooms, bathrooms...). Table 13 (Appendix B) provides details on how amenities affect the price.

Figure 2: Distribution of daily price

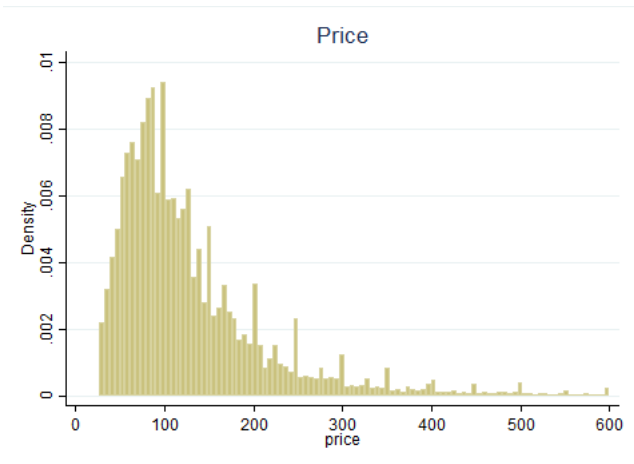

2.3 Ethnic groups and gaps

We consider two groups of ethnic minorities. First, we consider African-Americans, which we identify using the pictures provided on their host profile.[7] This ethnic group is only defined in North America. Second, we consider a group of hosts that have first names associated with Arabic or Sub-Saharan African ethnicity (labeled Arabic/African hereafter). We use the list provided in Jouniaux (2001). This ethnic group is defined both in North America and Europe.[8]

Table 2 displays the share of ethnic groups in the sample and the within-city*wave raw price gap. African-Americans represent roughly 2.0% of the observations in the sample, i.e. 5.3% of the North American observations. Compared to their share in total population, it seems that African-Americans are under-represented on the website or that some African-American hosts do not display a picture of themselves on the website. Hosts with Arabic or African names represent 1.1% of the sample in North America but 2.1% in Europe. There are some differences in the share of minorities across cities but those are not dramatic. NYC has 8% of African-American and 3% of Arabic/African observations, London and Paris both have around 4.5% of Arabic/African observations, while Milan and Rome have less than 1% of Arabic/African observations. Overall, the share of minorities is 5.1%. In the third column of the table, we display the gap between each ethnic group and the majority in daily prices, controlling only for the heterogeneity across cities and waves. The raw price gap is around 7-8% for Arabic/African hosts on both sides of the Atlantic but reaches 32% for African-Americans in North America.

Table 3 displays the ethnic price difference for several specifications. The first column displays within-city raw differential in daily log-prices: location is controlled at the level of city and differences in observables between groups are not taken into account. The raw ethnic gap is quite large (16%) and highly significant. Accounting for ethnic disparities in property observable characteristics reduces the gap to 10% (column 2), which shows that ethnic minorities have on average properties of lower observable quality. Instead of controlling for differences in observable characteristics of the property, we can control for finer heterogeneity of locations within cities. Including

[7]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 on the picture was African-American, (ii) whether nobody on the picture was African-American, (iii) whether it was impossible to say anything about the ethnicity of anyone on the picture or the picture was not showing any human being (pictures of flats, pets, furniture, landscape...). 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% for this dummy variable.

[8]There are other ethnic minorities than those considered in this analysis. Hispanics are difficult to identify in these data, given that first names used among the group are not necessarily distinguishable and picture characterisation is even more difficult. Asians can to some extent be identified through their first names: a preliminary analysis on Indian first names did not show any price differentials.

Table 2: Raw price gaps by ethnic groups

	Sample size	Share	Within-city*wave gap
Majority	3,264,326	94.9%	-
African-American (US/Can)	67,445	2.0%	31.8%
Arabic/African (US/Can)	37,639	1.1%	7.1%
Arabic/African (Eur)	72,276	2.1%	8.5%

Notes: 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.

neighborhood fixed-effects instead of city fixed-effects reduces the ethnic price gap from 16% with no fixed-effects to 7% (column 3). This indicates that ethnic-minority hosts tend to lives in neighborhoods that are less valued by potential guests. Finally, in the fourth column, both neighborhood and property characteristics are included in the regression: the residual ethnic price gap is reduced to 3.3% but is still highly significant. Note that the adjusted R-squared is high in this last specification, equal to .67. Observables are found to explain the largest part of the variance, as the adjusted R-squared is equal to .61 in the second column.

Stratifying this analysis by ethnicity shows that African-Americans start from a higher raw gap than Arabic/African hosts but end up with a lower price when location and characteristics are controlled for: 2.1% vs. 4.7%.

Table 3: Ethnic price gap, by specification

	Log daily rate			
	(1)	(2)	(3)	(4)
Minority	-0.155***	-0.095***	-0.067***	-0.033***
	(0.007)	(0.005)	(0.006)	(0.004)
City*Wave FE	Yes	Yes	Yes	Yes
Neighborhood FE	No	No	Yes	Yes
Property characteristics	No	Yes	No	Yes
Adj R^2	0.15	0.59	0.31	0.67
N obs.	3,441,686	3,324,141	3,441,686	3,324,141

Notes: OLS regression on the daily log-price on the minority dummy, controlling city-wave fixed-effects. See the list of all property characteristics in Table 13. Robust standard errors clustered at the property level. *** $p < 0.01$.

3 Conceptual framework

In this section, we introduce a simple conceptual framework to explain how we expect to separate the different mechanisms behind the ethnic price gap. In particular, we show that, under some assumptions, we can separate statistical discrimination from the other mechanisms. Taste-based discrimination cannot be separated for ethnic gap that comes from differentials in characteristics that are observed by potential guests but not by the econometrician. Our framework also allows us to test that ethnic minorities set lower prices because they have lower outside options.

3.1 Prices and demand as a function of quality

At each period (say, a week), a host shares his working time between two activities: renting his property (looking for guests, communicating with guests, cleaning up) or working on a regular job. L is the amount of labor put in renting and $1-L$ into the regular job. The technology to rent the property is supposed to be with decreasing returns to scale: the number of nights supplied is equal to $N = L^{\tilde{\alpha}}$, with $\tilde{\alpha} \in (0,1)$. The regular job has constant returns to scale. Overall, given the price of a night is P and the wage of the regular job is 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, decrease with P. Taste-based discrimination is embedded in this framework: demand is assumed to be shifted down when $m = 1$, relatively to $m = 0$. To simplify the notations, we write the inverse demand equation as:

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

where β and κ are strictly positive and $\Gamma > 1$ if there is taste-based discrimination, equal to 1 otherwise.

Hosts can set the effort they dedicate and the price to maximize their revenue, under the demand constraint; hence the following program:

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

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

$$p = p_0 + \lambda\alpha w + \lambda\beta q - \lambda\gamma m$$

where $p = \log P$, $w = \log W$, $q = \log Q$, $\gamma = \log \Gamma$, $\alpha = \frac{\tilde{\alpha}}{1-\tilde{\alpha}}$, $\lambda = (\kappa + \alpha)^{-1}$, $p_0 = \lambda\alpha \log(\frac{\kappa}{\tilde{\alpha}(\kappa-1)})$. Combining the log-demand and the log-price equations and eliminating quality, we obtain a relationship involving only the log-demand d, the log-price and the outside log-wage:

$$d = d_0 + (\lambda^{-1} - \kappa)p - \alpha w \tag{1}$$

3.2 Unobserved quality

Quality q is not perfectly observable by potential guests. It can be split in four categories $q = x + \zeta + \nu + u$ where:

- x are the characteristics written in the listing that both potential guests and the econometrician have access to (e.g. precise location);
- ζ are the characteristics in the listing that potential guests have access to but not the econometrician (e.g. interior decoration, on the pictures);
- ν are the unobservable characteristics that become accessible from the reviews (e.g. reliance of the host);
- u are the characteristics that remain unobservable.

Assume that ζ, ν and u have zero mean in the majority group and denote δ_ζ, δ_ν and δ_u the difference between the majority and the minority groups. Given that potential guests observe x and ζ and have no hope to learn about u, reviews are used to learn about ν. When there is no review, the best guess about ν is its expectation conditional on the host's group.[9] Statistical discrimination arise when $\delta_\nu > 0$, so that, everything else equal, the price set up by minority hosts has to be lower to compensate the lower demand induced by a lower average ν. Furthermore, ν is assumed to have a variance σ_ν^2.

A review k is assumed to transmit a signal r_k about ν: $r_k = \nu + \varepsilon_k$, where ε is iid of null expectation and variance σ_ε^2.[10] Using Bayes' rule and by induction, we can show that observing K reviews is equivalent to observe a signal $\bar{r} = \sum_k r_k/K \sim \mathcal{N}(\nu, \sigma_\varepsilon^2/K)$. Denoting $\rho = \sigma_\varepsilon^2/\sigma_\nu^2$ the ratio between the variances of the error term of the reviews, the expectation of the (posterior) belief on ν after observing the reviews is:

$$\mathbb{E}(\nu|\bar{r}, K, m) = \frac{K\bar{r} - \rho\delta_\nu m}{K + \rho}$$

Given that potential guests can observe x, ζ, K, $\bar{r}$ and m, a host with outside option w will set a price:

$$p = p_0 + \lambda\alpha w + \lambda\beta x + \lambda\beta\zeta + \lambda\beta\frac{K\bar{r}}{K+\rho} - \lambda\left(\gamma + \beta\frac{\rho\delta_\nu}{K+\rho} + \beta\delta_u\right)m$$

[9] In what follows, we make the assumption that ν is orthogonal to x and ζ.

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

The econometrician observes p, K, m and a proxy for $\bar{r}$. He also observes a vector of characteristics X from which x has to be inferred. Denote δ_w the difference between the mean of $\log w$ in the majority and the minority groups. The best possible prediction of the log-price based on what is observed by the econometrician is:

$$p = p_0 + \lambda\beta x + \lambda\beta\frac{K\bar{r}}{K+\rho} - \lambda\beta\frac{\rho\delta_\nu m}{K+\rho} - \lambda\left(\gamma + \beta\delta_\zeta + \beta\delta_u + \alpha\delta_w\right)m \qquad (2)$$

The comparison within-listing will help identify the parameter related to statistical discrimination δ_ν but the parameters related to taste-based discrimination γ, to the difference in ζ δ_ζ and to difference in outside options δ_w cannot be distinguished from each other.

3.3 Prices and reviews: Empirical evidence

In order to be able to identify statistical discrimination, we rely on the fact that reviews bring information.[11] First, we need to have some variability in the number of reviews. The left panel of Figure 3 displays the distribution of reviews across the observations of our sample. On the right panel, we have the variation of the number of reviews (between two waves) in the sample. The sample offers a decent amount of heterogeneity in the number of reviews, the empirical distributions being quite similar to that of a Poisson random variable.

Figure 3: Distribution of the number of reviews (left) and of the longitudinal variation in the number of reviews within a property (right)

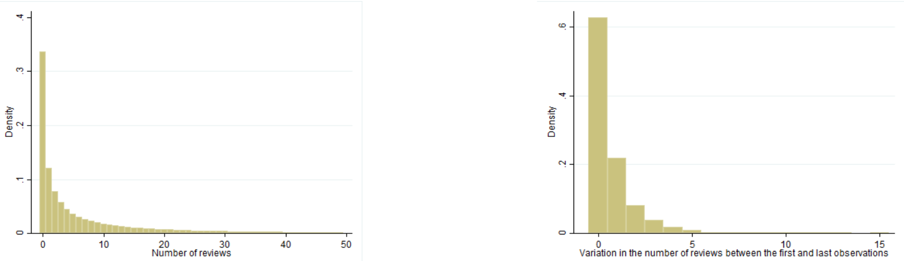

While ratings can vary between 1 and 5 stars (with half-star increments), the distribution of ratings is disproportionately skewed to good ratings, as documented in Fradkin et al. (2016) for Airbnb.com. If we consider the last rating observed for each property of our sample, 49% of observations have 5 stars and 34% 4.5 stars. By contrast, only 4% have 3.5 stars (see Table 12 in Appendix). We also know from our database that 75% of hosts have updated the price of their listings between the first and the last observations.

[11] Fradkin et al. (2016) studies in detail the reviewing system of the marketplace Airbnb.com and shows that reviews are informative.

According to our conceptual framework, hosts should update their prices as new information is available about the quality of their properties, i.e. as the number of reviews increases. An additional review providing less information than a previous one, the model predicts a concave relationship between the price and the number of reviews, converging to some value when the number of reviews tends to infinity. Also, the impact of new information on prices depends on unobservable quality. High-quality properties will benefit from new information while prices of low-quality properties are expected to decrease. Figure 4 provides a qualitative illustration of this Bayesian-updating phenomenon from a simulation of our model.

Figure 4: Simulation in the theoretical model: Prices with the number of reviews, by unobservable quality

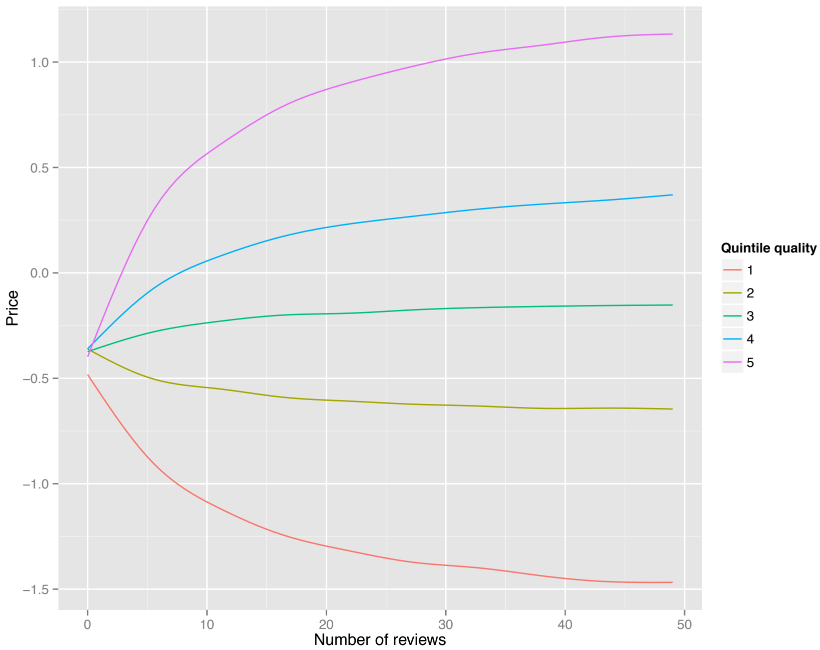

Do we observe this pattern in our data? We use as a proxy for unobservable quality the more recent rating of the properties, which is computed based on the largest number of reviews and is thus the most reliable we have on the listing. This rating can take four values: 5, 4.5, 4, and 3.5 stars and less. We regress the log-price on splines of the number of reviews interacted with the last rating and the full set of characteristics of the properties. We use linear splines with knots at 5, 10, 20, 30 and 50 reviews. The spline specification allows to flexibly allow for the hypothetical concavity of the relationship between prices and number of reviews without forcing it.

$$p_{it} = \sum_{r=3.5}^{5} 1\{\bar{r}_{it} = r\} s_r(K_{it}) + X_{it}\beta_x + \eta_i + \varepsilon_{it} \tag{3}$$

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, $s_r(.)$ are linear splines for each level of the last rating r and η are property fixed-effects. The results of the estimation are displayed in Figure 5. The figure shows that, depending on the last rating, the prices diverge in a way that is close to our conceptual framework. We see this result as additional evidence for the fact that (i) reviews provide information to potential guests, (ii) hosts use reviews and information to update their prices, and (iii) the last rating is a satisfactory proxy for the unobservable quality of the listing.

Figure 5: Estimated prices with the number of reviews, stratified by the most recent rating

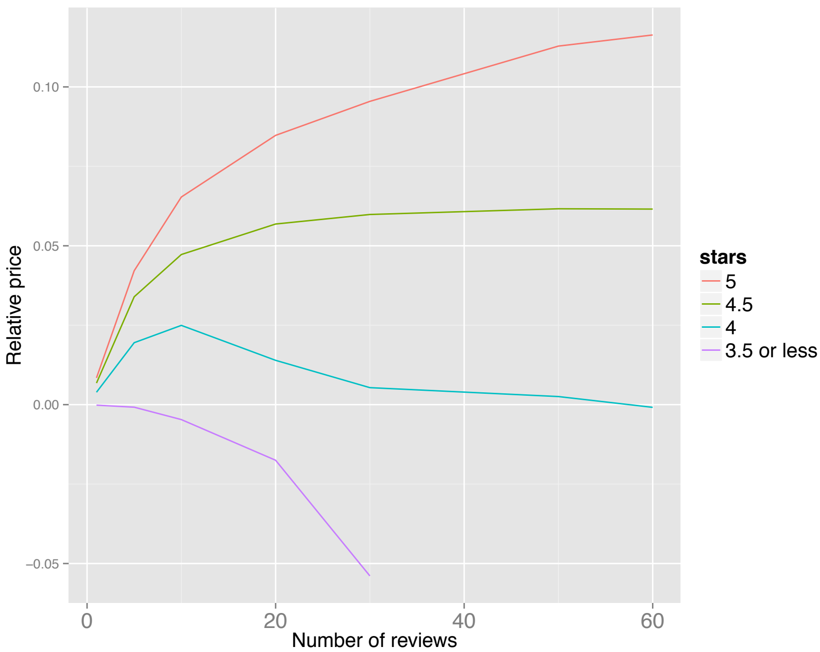

Notes: Equation (3) was estimated by linear regression with property fixed effects. We plot the estimates $\hat{s}_r(.)$ for all values of r, with the normalization $\hat{s}_r(0) = 0$. The number of observations of properties with ratings 3.5 or lower is very small when the number of reviews is higher than 30 and we do not report the corresponding estimates.

4 Ethnic price gaps and statistical discrimination

We first document how the unexplained ethnic price gap changes with the number of reviews. Table 4 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 ethnic gap changes across the samples: from 3.8% for listings with no reviews to 0.4% for listings with more than 50 reviews. While this pattern could be interpreted as suggestive evidence of statistical discrimination, it is subject, as any cross-section analysis to selection issues. For instance, potential guests could accept to be hosted by minorities only if the quality of the property was extremely good, and be less demanding for majority-host listings. In this case, the ethnic gap would be reduced, not because of the existence of statistical discrimination, but simply because the minority-host listings with many reviews are relatively much better than those with less reviews. However, this story would predict a drop in the share of minority listings with the number of reviews. It is not the case, we observe that the share of minority-host listings remains stable around 5.7-5.8% in all columns.

Table 4: Ethnic price gap, for several segments of the number of reviews

	Log daily rate				
	(1)	(2)	(3)	(4)	(5)
Minority	-0.038***	-0.031***	0.020***	-0.015	-0.004
	(0.007)	(0.006)	(0.007)	(0.010)	(0.016)
Nb reviews	0	1-4	5-19	20-49	50+
Minority share	5.7%	5.7%	5.8%	5.8%	5.7%
Adj R^2	0.635	0.708	0.749	0.761	0.759
N obs.	1,031,664	956,442	829,131	352,678	154,226

Notes: OLS regressions of the daily log-price on the minority dummy, controlling for neighborhood FE, property characteristics and ratings (for properties with at least one review). See the list of all property and host characteristics in Table 13. Robust standard errors clustered at the property level. *** $p < 0.01$.

Still, more sophisticated forms of differential selection could accommodate these findings. In order to deal with selection and unobserved heterogeneity, we estimate a within-listing model linking the evolution of prices with the increase in the number of reviews. Following our conceptual framework, we estimate the following model.

$$\Delta p_i = \sum_r 1\{\bar{r}_i = r\}\Delta K_i \beta_r + \Delta K_i m_i \beta_m + X_i \beta_x + \varepsilon_{it}$$

in which Δp is the variation in the log-price between the first and last observation of a property, ΔK is the variation in the number of reviews, X are the characteristics at the first observation and $\bar{r}$ is the rating at the last observation. According to our model, if reviews matter and rating provide some information about unobserved quality, we should have $\beta_r > \beta_{r'}$ if $r > r'$. Besides, in the presence of statistical discrimination, we should have $\beta_m > 0$.

Table 5 presents the results of the estimation of this model for three subsamples. Column 1 reports the estimates on the subsample of listings for which the minimum number of review is lower than 5; column 2 broadens the sample to listings for which the minimum is lower than 20; column 3 presents results on the full sample. The reason behind this stratification is that, because of the concavity of the theoretical relationship between prices and the number of reviews, we expect β_r and β_m to be lower in magnitude when the number of reviews is smaller.

The results in Table 5 are overall consistent with the predictions of the model. Better-quality listings (those with higher final ratings) experience higher increases in prices and the increase is stronger when the increase in the number of reviews is larger, which confirms the results obtained in cross-section. The estimate for β_m, which reflects the relative increase in prices with the number of reviews for minority listings is positive, indicating the presence of statistical discrimination. Interestingly, the coefficient of the minority dummy is close and insignificant, suggesting that, conditional on property characteristics, listings of minority hosts do not experience disproportionate variations compared to majority ones. Finally, the magnitude of the coefficients β_m and β_r are indeed smaller in columns 2 than in 1 and in 3 than in 2, which supports the hypothesis of a concave relationship between prices and the number of reviews.

Instead of working with last-first differences Δp and δK, one could use a fixed-effect specification:

$$p_{it} = \sum_r 1\{\bar{r}_i = r\} K_{it}\beta_r + K_{it} m_i \beta_m + X_{it}\beta_x + \eta_i + \varepsilon_{it}$$

where notations are the same as in equation (3). Compared to the model in last-first differences, this specification relies arguably on smaller differences (in information and prices) and assumes that hosts instantaneously adjust prices to information innovation. Table 6 presents the results of the estimation of this model. Column (1) focuses on a subsample restricted to the first and last observation of each property: the estimate of β_m turns out to be fairly similar to the value found in the first column of Table 5. Column (2) widens the sample to all observations of properties starting with less than 5 reviews. The coefficient decreases notably, which may reflect the fact that hosts' reactions to new information is not fully instantaneous. In columns (3) and (4), we see that, as in Table 5, including observations with more reviews makes the estimate smaller, which is consistent with the assumption of a concave relationship. In column (5), we replace the linear relationships by quadratic one and find

Table 5: Estimation of the model in difference between the first and last observations

	Variation of log-price		
	(1)	(2)	(3)
3.5 stars	-0.021***	-0.023***	-0.027***
	(0.005)	(0.004)	(0.004)
4 stars	-0.017***	-0.019***	-0.021***
	(0.003)	(0.003)	(0.003)
4.5 stars	0.006**	0.003	0.002
	(0.002)	(0.002)	(0.002)
Minority	0.002	0.002	0.003
	(0.002)	(0.002)	(0.002)
3.5 stars $\times \Delta K/100$	-0.206**	-0.191***	-0.179***
	(0.080)	(0.069)	(0.063)
4 stars $\times \Delta K/100$	-0.096***	-0.104***	-0.089***
	(0.036)	(0.028)	(0.022)
4.5 stars $\times \Delta K/100$	0.065***	0.003	-0.035***
	(0.016)	(0.012)	(0.008)
5 stars $\times \Delta K/100$	0.212***	0.130***	0.063***
	(0.015)	(0.011)	(0.007)
Minority $\times \Delta K/100$	0.157***	0.126***	0.076***
	(0.039)	(0.032)	(0.022)
Sample	$\min(K) \leq 5$	$\min(K) \leq 20$	Full
Adj R^2	0.115	0.120	0.123
N obs.	324,589	365,874	385,667

Notes: OLS regressions. Aside from those mentioned in the Table, controls include city*wave FE, neighborhood FE and property characteristics (see Table 13). Robust standard errors clustered at the property level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

that, compared to the majority group, the slope of the price is larger but decreasing with the number of reviews. This is consistent with our theoretical framework.

The previous results provide evidence that statistical discrimination contributes to the ethnic price gap but do not allow us to assess the magnitude of the phenomenon. In order to measure what share of the ethnic gap statistical discrimination explain, we turn back to our conceptual framework and estimate the parameters relating to statistical discrimination $\beta_m = \lambda\beta\delta_\nu$. We use the number of stars s (taking values 3.5, 4, 4.5, or 5) observed in the last observation of each listing as a proxy for $\bar{r}$. We do not observe x and use the vector X of observable characteristics, as well as dummies for the city interacted with the wave in which the listing appeared. As the main outcome, we use the difference in prices, within listing, between the first and

Table 6: Fixed-Effects Estimation

	log-price				
	(1)	(2)	(3)	(4)	(5)
3.5 stars$\times K/100$	0.034	-0.050	-0.109*	-0.120**	-0.184**
	(0.066)	(0.069)	(0.058)	(0.054)	(0.074)
4 stars$\times K/100$	0.082***	0.028	-0.033	-0.055***	-0.021
	(0.030)	(0.029)	(0.024)	(0.020)	(0.027)
4.5 stars$\times K/100$	0.276***	0.188***	0.091***	0.032***	0.094***
	(0.014)	(0.013)	(0.010)	(0.007)	(0.011)
5 stars$\times K/100$	0.383***	0.299***	0.202***	0.117***	0.206***
	(0.014)	(0.013)	(0.009)	(0.007)	(0.014)
Minority $\times K/100$	0.147***	0.089**	0.059	0.025	0.086**
	(0.043)	(0.037)	(0.036)	(0.025)	(0.036)
3.5 stars$\times K/100^2$					0.180
					(0.116)
4 stars$\times K/100^2$					-0.034
					(0.022)
4.5 stars$\times K/100^2$					-0.053***
					(0.008)
5 stars$\times K/100^2$					-0.075***
					(0.011)
Minority $\times (K/100)^2$					-0.059**
					(0.023)
Samples	Min(K)<5	Min(K)<5	Min(K)<21	-	-
	Min Max	Full	Full	Full	Full
N obs.	597,061	2,507,078	3,033,699	3,324,141	3,324,141

Notes: OLS regressions with host fixed effects. Aside from those mentioned in the Table, controls include city*wave FE, neighborhood FE and property characteristics (see Table 13). Robust standard errors clustered at the property level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

the last observation Δp.

We estimate the parameters of the following equation by non-linear least-squares, β_m and ρ being the parameter of main interest.

$$\Delta p = X\beta_x + \sum_{s=3.5}^{5} \beta_s \left[\frac{K_1}{K_1+\rho} - \frac{K_0}{K_0+\rho}\right] - \beta_m m \left[\frac{\rho}{K_1+\rho} - \frac{\rho}{K_0+\rho}\right]$$

where K_0 and K_1 are the number of reviews at the first and last observations. Inference is performed by bootstraping at the property level.

We obtain an estimated value of 8 for ρ. According to the previous equation, 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 $(\underline{p}+\overline{p})/2$ is reached after ρ reviews. In our case, 8 reviews are necessary to reveal half of the information about the unobservables.

β_m is estimated to be equal to .077 (.011), which means that going from 0 to an infinite number of reviews would increase the prices of minority by 7.7%. This figure is of the same order of magnitude as, and even larger than, the ethnic price gap observed in the subset of listings with no reviews (3.8%, see Table 4, column 1). This suggests that the totality of the initial gap would be overcome by the revelation of information. According to this result, the full ethnic gap can be accounted for by statistical discrimination.

5 Additional results and robustness checks

5.1 Ethnic differences in pricing behavior

A potential explanation to explain why minority-host listings have lower prices is that minority hosts have on average lower outside options than majority hosts. This relates, in our conceptual framework, to a lower w. A lower outside wage entails a lower price but it should also lead to a higher demand, conditional on price. We test this prediction using the number of new reviews between two waves as a proxy for demand. This proxy relies on the assumption that the number of new reviews is proportional to the number of nights the property was occupied by a guest. More precisely, we build two outcomes: a dummy for having at least one new review between t and $t+1$, and the log of the number of reviews.

Table 7 presents the results of the regression of these two outcomes on the log-price (at t), controlling for location and observable characteristics: columns 1-3 for the dummy and 4-6 for the log new reviews. In columns 2 and 4, lagged prices are included in a more flexible manner (using splines). In columns 3 and 6, we include the number of reviews (at t) as an additional covariate. In all columns, we find that the coefficient of the minority is close to zero and insignificant. These results suggest minority hosts do not get more demand than majority hosts, despite the lower prices. The ethnic price gap does not seem to reflect differences in pricing behavior induced by differences in outside wages.

Table 7: Variation in the number of reviews between two waves as a function of host ethnicity, controlling for prices

	Dummy for any new review			Number of new reviews		
	(1)	(2)	(3)	(4)	(5)	(6)
Log price	-0.129***		-0.094***	-0.149***		-0.111***
	(0.001)		(0.001)	(0.002)		(0.002)
Minority	-0.001	-0.001	0.002	0.000	0.000	0.004
	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.004)
$K/100$			0.644***			0.633***
			(0.005)			(0.006)
Price functional form	Linear	Spline	Linear	Linear	Spline	Linear
Adj R^2	0.192	0.192	0.249	0.101	0.101	0.174
N	3,324,141	3,324,141	3,324,141	932,478	932,478	932,478

Notes: OLS regressions. Aside from those mentioned in the Table, controls include city*wave FE, neighborhood FE and property characteristics (see Table 13). Robust standard errors clustered at the property level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.2 Do ethnic groups compete on the same market?

In the previous analyses, we have made the implicit assumption that minority and majority hosts compete on the same market. Conversely, it may be that the two markets are segmented: minority hosts receiving almost only guests of their own ethnicities. To investigate this issue, we have to extract information about guests' ethnicity. We have access to the first name of the last ten guests leaving reviews on each listing and we code whether a guest has an Arabic/African name.

For each listing and wave, we compute the share of the new reviews that comes from guests with an Arabic/African first name. We then regress this share on a dummy for the host ethnicity, controlling for the location and the observable characteristics of the listing.

We find some evidence for a very mild ethnic matching: a host with an Arabic/African first name will be 1 percentage point more likely to have a review from a guest with an Arabic/African first name. Overall, despite the mild ethnic matching, our results support the assumption that hosts belonging to different ethnic groups compete on the same market.

Table 8: Ethnic matching between hosts and guests

	Minority share in total guests
Arabic African	0.009***
	(0.000)
Adj R^2	0.012
N obs.	2,272,120

OLS regression. Aside from those mentioned in the Table, controls include city*wave FE, neighborhood FE, property characteristics (see Table 13), number of reviews and ratings. Standard errors are clustered at the property level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.3 Robustness check and results on sub-groups

We check the robustness of our results by running some of the analyses on several sub-samples. First, the marketplace is heterogeneous in terms of sellers. Both professional and non-professional hosts offer properties on the website. Second, we can split the analysis by continent, to make sure that results are comparable in Europe and North America. Finally, we check that our results are not sensitive to the way we control for geographic unobserved heterogeneity. Instead of neighborhood dummies, we build 5000 squared blocks using longitude and latitude of listings.

Table 9 reports the results of these robustness checks. For each sample or specification, Panel A shows the unexplained price gap on the sample of properties with no review. According to our model, the ethnic price gap is maximum at zero review and decreases once information is revealed. Initial gaps appear to be smaller for single-property hosts and in North America but remain in the same ballpark. Controlling for block fixed effects instead of neighborhood fixed effects does not affect the results at all. Panel B shows the result of the estimation of the constrained model. In all cases, the point estimate of b_m is of the same magnitude as (or larger than) the zero-review ethnic price gap.

Table 9: Robustness checks

	Full sample	Single-property	US & Canada	Europe	Block fixed effects
	(1)	(2)	(3)	(4)	(5)
	Panel A. Unexplained ethnic price gap (zero-review sample)				
Minority	-0.038***	-0.027***	-0.024**	-0.056***	-0.037***
	(0.007)	(0.008)	(0.010)	(0.010)	(0.007)
Adj R^2	0.64	0.63	0.66	0.63	0.66
N obs.	1,031,664	687,806	350,499	681,165	1,031,664
Share Minority	.057	.062	.094	.039	.057
	Panel B. Estimation of the constrained model				
b_m	-0.077***	-0.085***	-0.068***	-0.041***	-0.077***
	(0.011)	(0.027)	(0.016)	(0.016)	(0.011)
ρ	8.08	5.69	9.78	9.02	7.76
	(0.714)	(2.15)	(0.847)	(4.15)	(0.687)

Inference is done by bootstraping with 100 replications.

6 Conclusion

This paper shows that, in a popular online platform of short-term rentals, hosts belonging to an ethnic minority experience a 3% price penalty, when differences in locations and observable characteristics are accounted for. Taking advantage of the longitudinal nature of our data, we show that statistical discrimination can be considered to be the only significant driver of the ethnic price gap.

We can draw several conclusions from this finding. First, aside from the issues inherent to any online feedback system, the one proposed by this online platform is effective in supplying useful information to potential guests. Second, in the absence of such a feedback system, the ethnic price gap would be higher than its current value. Third, beside the gains in efficiency that improving the feedback system would have, we can expect that it would also contribute to reduce ethnic price gaps. Our findings suggest that reducing ethnic gaps requires disclosing more abundant and more reliable information about candidates, sellers or hosts. As discussed by Shaw et al. (2011), it remains to understand how platforms can adequately incentivize reviewers to provide informative and helpful reviews.

References

Agrawal, A., N. Lacetera, and E. Lyons (2014): "Does Standardized Information in Online Markets Disproportionately Benefit Job Applicants from Less Developed Countries?" Mimeo.

Aigner, D. J. and G. G. Cain (1977): "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review*, 30, 175–187.

Altonji, J. G. and C. R. Pierret (2001): "Employer Learning And Statistical Discrimination," *The Quarterly Journal of Economics*, 116, 313–350.

Antonovics, K., P. Arcidiacono, and R. Walsh (2005): "Games and Discrimination: Lessons from The Weakest Link'," *Journal of Human Resources*, 40, 918–947.

Anwar, S. (2012): "Testing for discrimination: Evidence from the game show Street Smarts," *Journal of Economic Behavior & Organization*, 81, 268–285.

Arrow, K. J. (1973): "The Theory of Discrimination," in *Discrimination in Labor Markets*, ed. by O. Ashenfelter and A. Rees, Princeton University Press, 3–33.

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

——— (2009): *The Economics of Labor Market Intermediation: An Analytic Framework*, University of Chicago Press.

Autor, D. H. and D. Scarborough (2008): "Does Job Testing Harm Minority Workers? Evidence from Retail Establishments," *The Quarterly Journal of Economics*, 123, 219–277.

Bagues, M. F. and M. S. Labini (2009): "Do Online Labor Market Intermediaries Matter? The Impact of AlmaLaurea on the University-to-Work Transition," in *Studies of Labor Market Intermediation*, National Bureau of Economic Research, Inc, NBER Chapters, 127–154.

Bayer, P., M. D. Casey, F. Ferreira, and R. McMillan (2012): "Price Discrimination in the Housing Market," Working Papers 12-10, Duke University, Department of Economics.

Becker, G. (1957): *The Economics of discrimination, Chicago University Press, Chicago*, Reprint 1971.

Behaghel, L., B. Crpon, and T. L. Barbanchon (2015): "Unintended Effects of Anonymous Rsums," *American Economic Journal: Applied Economics*, 7, 1–27.

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

Brown, M., E. Setren, and G. Topa (2016): "Do Informal Referrals Lead to Better Matches? Evidence from a Firms Employee Referral System," *Journal of Labor Economics*, 34, 161–209.

Castillo, M. and R. Petrie (2010): "Discrimination in the lab: Does information trump appearance?" *Games and Economic Behavior*, 68, 5059.

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

——— (2011): "Studying Discrimination: Fundamental Challenges and Recent Progress," *Annual Review of Economics*, 3, 479–511.

Doleac, J. L. and L. C. Stein (2013): "The Visible Hand: Race and Online Market Outcomes," *Economic Journal*, 123, F469–F492.

Edelman, B. and M. Luca (2014): "Digital Discrimination: The Case of Airbnb.com," Harvard Business School Working Papers 14-054, Harvard Business School.

Edelman, B. G., M. Luca, and D. Svirsky (2016): "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment," Harvard Business School Working Papers 16-069, Harvard Business School.

Ewens, M., B. Tomlin, and L. C. Wang (2014): "Statistical Discrimination or Prejudice? A Large Sample Field Experiment," *The Review of Economics and Statistics*, 96, 119–134.

Farber, H. S. and R. Gibbons (1996): "Learning and Wage Dynamics," *The Quarterly Journal of Economics*, 111, 1007–47.

Fershtman, C. and U. Gneezy (2001): "Discrimination In A Segmented Society: An Experimental Approach," *The Quarterly Journal of Economics*, 116, 351–377.

Fradkin, A. (2015): "Search Frictions and the Design of Online Marketplaces," Mimeo.

Fradkin, A., E. Grewal, D. Holtz, and M. Pearson (2016): "Bias and Reciprocity in Online Reviews: Evidence From Field Experiments on Airbnb," Mimeo.

Heckman, J. J. (1998): "Detecting Discrimination," *Journal of Economic Perspectives*, 12, 101–116.

Horton, J. J. (2015): "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment," Forthcoming in the Journal of Labor Economics.

Jouniaux, L. (2001): *Les 20,000 plus beaux prénoms du monde*, Hachette eds.

Knowles, J., N. Persico, and P. Todd (2001): "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy*, 109, 203–232.

Lang, K. and J.-Y. K. Lehmann (2012): "Racial Discrimination in the Labor Market: Theory and Empirics," *Journal of Economic Literature*, 50, 959–1006.

Levitt, S. D. (2004): "Testing Theories of Discrimination: Evidence From Weakest Link," *Journal of Law and Economics*, 47, 431–452.

List, J. A. (2004): "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," *The Quarterly Journal of Economics*, 119, 49–89.

Oreopoulos, P. (2011): "Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes," *American Economic Journal: Economic Policy*, 3, 148–71.

Pallais, A. (2014): "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review*, 104, 3565–99.

Pallais, A. and E. G. Sands (2015): "Why the Referential Treatment? Evidence from Field Experiments on Referals," Forthcoming in the Journal of Political Economy.

Phelps, E. S. (1972): "The Statistical Theory of Racism and Sexism," *American Economic Review*, 62, 659–661.

Pope, D. G. and J. R. Sydnor (2011): "Whats in a Picture?: Evidence of Discrimination from Prosper.com," *Journal of Human Resources*, 46, 53–92.

Shaw, A. D., J. J. Horton, and D. D. Chen (2011): "Designing Incentives for Inexpert Human Raters," in *Proceedings of the ACM Conference of Computer Supported Cooperative Work (ACM-CSCW)*, ed. by J. Bardham and N. Ducheneaut, ACM, 275–284.

Stanton, C. T. and C. Thomas (2015): "Information Frictions and Observable Experience: The New Employer Price Premium in an Online Market," Mimeo.

Wozniak, A. K. (2015): "Discrimination and the Effects of Drug Testing on Black Employment," *Review of Economics and Statistics*, 97, 548–566.

Zussman, A. (2013): "Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars," *Economic Journal*, 123, F433–F468.

A Online Platform

A.1 Example of listing

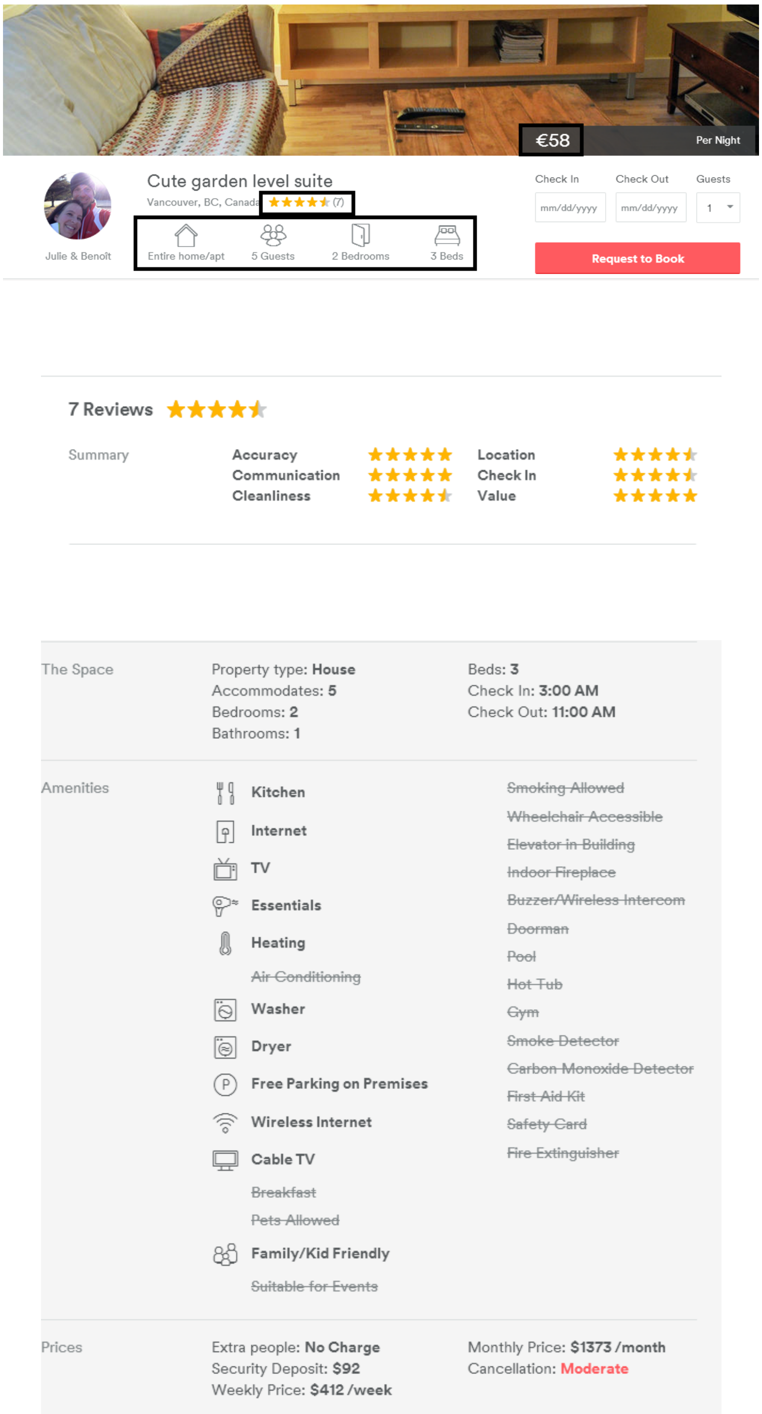

A.2 Peer-reviewing System

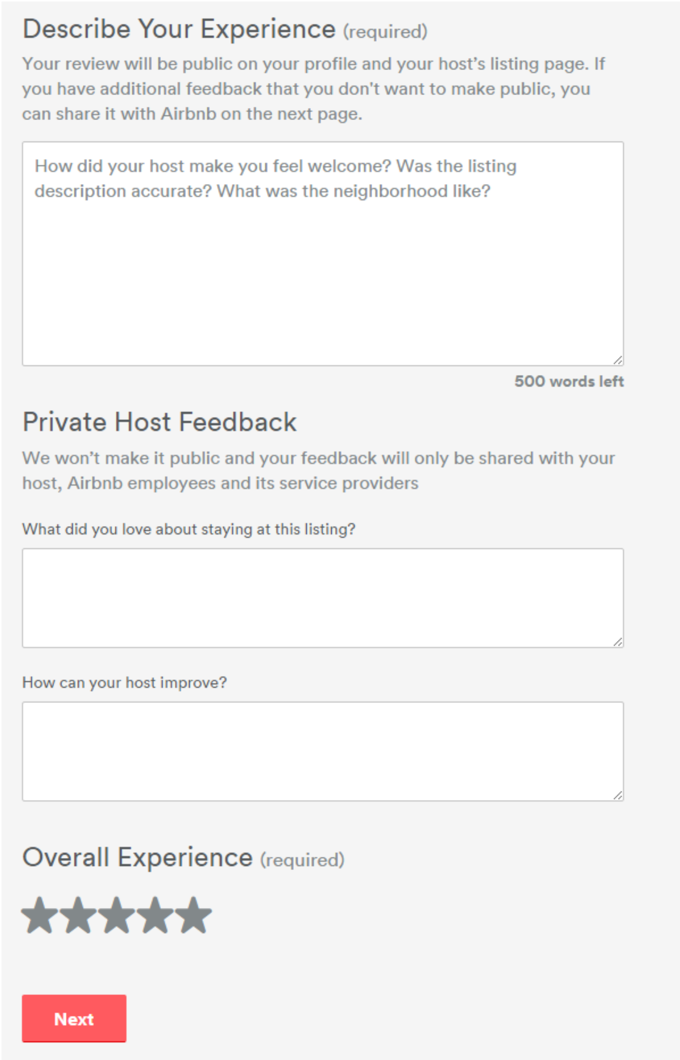
Describe Your Experience (required)

Your review will be public on your profile and your host's listing page. If you have additional feedback that you don't want to make public, you can share it with Airbnb on the next page.

How did your host make you feel welcome? Was the listing description accurate? What was the neighborhood like?

500 words left

Private Host Feedback

We won't make it public and your feedback will only be shared with your host, Airbnb employees and its service providers

What did you love about staying at this listing?

How can your host improve?

Overall Experience (required)

Next

B Data

Table 10: Number of observations by city

City	Obs	Share
Amsterdam	135,261	3.93
Barcelona	230,272	6.69
Berlin	209,830	6.10
Boston	52,209	1.52
Chicago	52,205	1.52
Florence	85,186	2.48
London	369,727	10.74
Los Angeles	211,440	6.14
Madrid	93,003	2.70
Marseille	86,278	2.51
Miami	90,984	2.64
Milan	122,944	3.57
Montreal	104,671	3.04
New-York	477,663	13.88
Paris	640,513	18.61
Rome	204,907	5.95
San-Francisco	133,069	3.87
Toronto	81,613	2.37
Vancouver	59,911	1.74

Table 11: Collection of waves

Wave	Time Period
0	15 June 2014
1	8 July 2014
2	28 July 2014
3	11 August 2014
4	25 August 2014
5	8 September 2014
6	25 September 2014
7	15 October 2014
8	5 November 2014
9	25 November 2014
10	15 December 2014
11	7 January 2015
12	13 January 2015
13	3 February 2015
14	4 March 2015
15	25 March 2015
16	13 April 2015
17	4 May 2015
18	26 May 2015
19	15 June 2015

Table 12: Distribution of the last rating

	Obs	Share
3.5 stars	8,831	4.18%
4 stars	25,606	12.11%
4.5 stars	72,601	34.33%
5 stars	104,418	49.38%

Sample: Listings for which last rating is observed.

Table 13 shows observable characteristics explain a large share of the variance. These covariates are all included in the following regressions. In column (2), neighborhood fixed effects are included in the equation. It shows the adjusted R-squared increase by 11% when including neighborhood fixed-effects.

Table 13: Log daily rate

Shared flat	-0.711***	-0.633***
	(0.004)	(0.003)
Person Capacity (> 2)	0.202***	0.207***
	(0.003)	(0.003)
Nber bedrooms	0.274***	0.298***
	(0.009)	(0.008)
Nber bathrooms	0.090**	0.078**
	(0.034)	(0.030)
Flat	-0.229***	-0.244***
	(0.008)	(0.007)
House or Loft	-0.210***	-0.125***
	(0.009)	(0.008)
Couch	-0.177***	-0.151***
	(0.014)	(0.013)
Airbed	-0.141***	-0.094***
	(0.024)	(0.023)
Sofa	-0.177***	-0.166***
	(0.006)	(0.006)
Futon	-0.153***	-0.113***
	(0.010)	(0.009)
Terrace or Balcony	0.028***	0.037***
	(0.003)	(0.003)
Cable TV	0.140***	0.106***
	(0.003)	(0.002)
Wireless	-0.027***	-0.043***
	(0.005)	(0.004)
Heating	-0.066***	-0.053***
	(0.005)	(0.004)
AC	0.161***	0.137***
	(0.004)	(0.003)
Elevator	0.101***	0.091***
	(0.003)	(0.003)
Wheelchair Accessible	-0.045***	-0.018***

(Continued on next page)

Table 13: Log daily rate

	(0.004)	(0.004)
Doorman	0.113***	0.052***
	(0.005)	(0.004)
Fireplace	0.154***	0.121***
	(0.005)	(0.004)
Washer	-0.062***	-0.026***
	(0.003)	(0.003)
Dryer	0.158***	0.114***
	(0.003)	(0.003)
Parking	-0.125***	0.007*
	(0.003)	(0.003)
Gym	0.066***	0.062***
	(0.006)	(0.006)
Pool	0.115***	0.126***
	(0.007)	(0.006)
Buzzer	0.026***	-0.005*
	(0.002)	(0.002)
Hot Tub	0.015**	0.012*
	(0.005)	(0.005)
Breakfast served	0.012**	0.035***
	(0.004)	(0.004)
Family/Kids Friendly	0.001	0.017***
	(0.003)	(0.002)
Suitable for events	0.078***	0.077***
	(0.006)	(0.006)
Additional People	-0.052***	-0.034***
	(0.002)	(0.001)
Price per Additional People	0.000	-0.001***
	(0.000)	(0.000)
Cleaning price	0.003***	0.002***
	(0.000)	(0.000)
Cancellation Policy	-0.003*	-0.019***
	(0.001)	(0.001)
Smoking Allowed	-0.049***	-0.038***
	(0.002)	(0.001)
Pets Allowed	-0.013***	-0.013***
	(0.002)	(0.002)
Host in couple	-0.041***	-0.024***

(Continued on next page)

Table 13: Log daily rate

	(0.004)	(0.004)
Host has multiple properties	0.059***	0.032***
	(0.003)	(0.002)
Member since 2008-2009	0.085***	0.076***
	(0.012)	(0.011)
Member since 2010-2011	0.057***	0.046***
	(0.005)	(0.004)
Member since 2012-2013	0.020***	0.016***
	(0.003)	(0.002)
City*Wave FE	Yes	Yes
Neighborhood FE	No	Yes
Property characteristics	Yes	Yes
Adj R^2	0.60	0.67
N obs.	3,324,141	3,324,141

Notes: OLS regression on the daily log-price. Robust standard errors clustered at the property level. *** $p < 0.01$.