Ethnic Discrimination on an Online Marketplace of Vacation Rentals*

Morgane Laouénan†
Roland Rathelot‡

February 28, 2016

Preliminary, Please do not quote

Abstract

We use data from an online market of vacation rentals to measure the ethnic price discrimination towards properties' owners and to measure the contribution of statistical discrimination. Following a strategy à la Altonji and Pierret (2001), we take advantage of the existence of a detailed reviewing system to measure the influence of better signals on prices and of the panel dimension of our data. First, controlling for a rich set of characteristics reduces the ethnic price gap from 10% to 3%. Then, using the longitudinal nature of our data, we show that, conditional of the last rating obtained by the listing, an additional review increases the price more for minority than for majority owners. Finally, estimating the parameters of our theoretical 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 Leah Boustan, Clément de Chaisemartin, Raj Chetty, Bruno Decreuse, Christian Dustmann, Sebastien Fontenay, Nathan Hendren, Attila Lindner, Kevin Lang, Victor Lavy, Andrea Moro, David Neumark, Barbara Petrongolo, Fabian Waldinger, Natalia Zinovyeva and participants in seminars at IEB, UC Irvine, UCLA (CCPR), UC Santa Barbara, Umea, 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).

†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 behavioral mechanisms are at work is necessary to design efficient policies. This paper investigates the extent to which ethnic minorities are discriminated on one of the world-leading online marketplace for short-term rentals and separates statistical discrimination from other mechanisms.

On the online marketplace we study, hosts list their properties, set the nightly 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 that belong to an ethnic minority and those set by majority hosts. Controlling for a large set of observable characteristics reduces the price gap by half but the "unexplained" gap remains significant.[1] Is this unexplained gap 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; Charles and Guryan, 2011). In order to isolate statistical discrimination, we adapt the approach by Farber and Gibbons (1996); Altonji and Pierret (2001) to our context. Because statistical discrimination hinges on imperfect information about the quality of the good for sale, additional information should reduce the ethnic price gap. Our setting is well adapted to this method, as the information set about a property that is available to potential guests evolve frequently over time and we can observe a part of the new signals. New properties start with self-assessed information about characteristics. Then, guests that have stayed in a property have the possibility to let a quantitative rating and a qualitative assessment of the accommodation and the host. As time goes, the number of reviews grows and more and more information is available to potential guests.

We build a simple conceptual model in which properties' quality is partly unobservable. When a property has no review, potential guests can only infer unobservable quality using owner's 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

[1] See also Edelman and Luca (2014) for similar results on the website Airbnb.com, about the Black-White price gap in New York City.

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 taste-based discrimination or ethnic gaps 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.

The data we use include nightly prices, hosts' and apartments' characteristics, as well as associated reviews. We collected the data relating to 400,000 properties, corresponding to entire and shared apartments to rent in 19 cities in North America and Europe. 20 waves of data, collected every two or three weeks between June 2014 and July 2015 form an unbalanced panel of 3,500,000 observations. The ethnic minority groups we consider are: (i) owners with Arabic/African/Muslim first names, (ii) owners with Hispanic first names (North America only), (iii) owners with African-American first names (North America only).

We find that the raw within-city ethnic price gap is around 10%. The set of observable characteristics about the property (including its precise location) is rich and explains more than 70% of the variance of the prices. Controlling for ethnic differences in these characteristics reduces the ethnic price gap to 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 40 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 results contribute to an extensive but largely inconclusive literature on the sources of discrimination, namely taste-based and statistical discrimination. On the U.S. labor market, Altonji and Pierret (2001) find that statistical discrimination would play a small role to explain the ethnic wage gaps. Charles and Guryan (2008) provide an empirical test of associations between prejudice and wages implied by the Becker prejudice model and find that around one quarter of the unconditional Black-White wage gap is due to prejudice, while the three other quarters can be due to differences in unobservables or other forms of discrimination. Findings obtained by experimental studies suggest statistical discrimination would play a limited role in explaining ethnic employment gaps in North America. Sending fictitious resumes with randomly assigned African-American- or White-sounding names to 1300 help-want ads, Bertrand and Mullainathan (2004) find that white names receive 50% more callbacks than distinctively black-named applicants. Moreover, they also find that

African-Americans benefit less from resume enhancements (such as honors, more experience), which goes against evidence of statistical discrimination. Oreopoulos (2011) reports the results of a large-scale correspondence study, in which the amount of information (about education, experience...) is allowed to vary across candidates. Adding information does not seem to reduce discrimination, which adds some evidence against statistical discrimination on the hiring process.

Other contexts and sectors have been looked at by the literature on discrimination and their results support evidence of statistical discrimination. Wozniak (2015) uses time variation in drug-testing legislation to provide evidence for substantial statistical discrimination against low-skilled Black men. 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. Anwar (2012) find that non-black contestants of an American game show erroneously believe that Afro-Americans have lower skill levels. In the case of the sportscard market, List (2004) finds that the lower offers received by minorities reflect statistical discrimination. Bayer et al. (2012) show that the black and Hispanic homebuyers pay 3% premiums on the U.S. housing market. Taste-based discrimination is ruled out as the premium is the same when the seller is himself black or Hispanic. Using data from the television game show *The Weakest Link*, Levitt (2004) and Antonovics et al. (2005) disentangle the sources of discrimination in focusing on the behavior of participants deviating from the optimal voting strategy. Their results suggest no evidence of discriminatory voting patterns by males against females or by whites against blacks. Fershtman and Gneezy (2001) run a randomized experiment based on different type of games (trust game, dictator game, ultimatum game) between Ashkenazic and Eastern Jews. They find evidence of differential treatment by ethnic origin mostly driven by statistical discrimination due to prior beliefs.

As we do in this paper, some research make use of online markets to distinguish between the two theories of discrimination. Zussman (2013) finds that the discrimination towards Arabs on an online market for used cars in Israel is rather due to statistical than taste-based discrimination. Through a field experiment on the online US rental apartment market, Ewens et al. (2014) show their results are consistent with evidence of statistical discrimination rather than taste-based models. Using data from a peer-to-peer lending website, Pope and Sydnor (2011) find that blacks lenders face higher interest rates and lower borrowing probabilities. However, blacks have higher default rates so that net returns on loans made to blacks is lower. According to the authors, these results would be consistent with accurate statistical discrimination against blacks and taste discrimination against whites. In a market of iPods, Doleac and Stein (2013) compare offers received to online ads featuring a dark- or light-skinned hand. Black ads receive fewer and lower offers. Outcomes are poorer in thin markets and those with higher racial isolation and crime, which suggest statistical rather than taste discrimination.

We bring several contributions to the literature. First, this study is the first one to investigate the mechanisms behind the ethnic price gap on the market for short-term rentals. Launched in 2008, this growing online marketplace for vacation rentals claims to propose more than 800,000 listings in 190 different countries and to have served over 10 million guests, making it an intrinsically interesting research object. Second, the size of the data allow us to provide a precise assessment while the longitudinal dimension and the availability of high-frequency provides a unique opportunity to test for statistical discrimination. Because of the way labor markets work, the approach proposed by Altonji and Pierret (2001) is difficult to apply. In cross-section, information usually vary very endogenously. In longitudinal analyses, reliable and long-term observation of wages is scarce, especially if one wants to couple it with a measure of productivity that is not available to employers. In our context, 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.

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 Context and Data

2.1 Description of the platform

This marketplace gathers owners 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. Owners also have to provide information about their properties. The information about properties and owners 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 owner and the rating. When they click on one of the listing, they have access to more detailed information, notably the first name of the owner, 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.

At any moment, owners can revise the price of their properties. However, the system does not allow negotiation. Once the potential guest has decided which place he preferred among those available during the period he has chosen, he can choose to place a bid. The bid is then in the hands of the owner that can decide, without any justification, to accept or reject the offer, based on the information he has about the potential guest. 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.[2] However, the potential guest can decide to cancel his booking. In this case, the terms of the cancellation policy (specified on the listing) applies: depending on the flexibility of the policy, different amounts are charged. The owner 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 owner'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 owners, 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 host, characteristics of the listing, 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. We repeated the collection process every 2-3 weeks between June 2014 and July 2015, obtaining 20 waves. See the collection date of each wave in Table 8 in Appendix. Our sample include 400,000 distinct properties. The panel is unbalanced: some properties

[2] 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.

enter the system while others exit.

Tables 1 and 2 present 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, and on which most of the empirical analysis will be based. 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 only half of them can host more than two 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. Aside from the first name, a picture and a free-text description, guests know about the couple status of hosts, whether they have other properties and when they joined the platform. Most hosts do not declare as being in couple, have only one property and have joined relatively recently (since 2012).

The distribution of the number of waves each property is observed 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, for a total of 2,104,366 observations. The number of listings observed per wave is displayed in the left panel of Figure 1. Local maxima are observed for waves 3 and 11, which were respectively collected early August and early December and correspond to peaks in vacation periods.

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

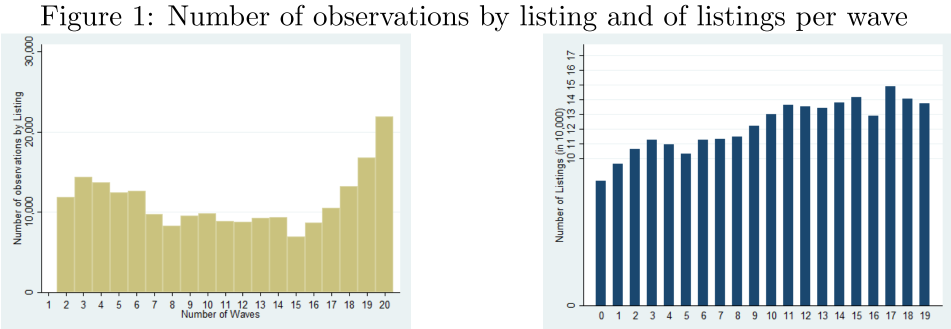

Figure 2 shows the distribution of log price of both entire and shared flats. The nightly price varies a lot across cities and according to the amenities of the listing

Table 1: Summary statistics: Property and host characteristics

	Full Sample	**Active Listings**
Type of property		
Shared flat	0.313	0.289
Apartment	0.830	0.838
House/Loft	0.129	0.127
Size		
Person Capacity (>2 people)	0.501	0.520
Number of bedrooms	1.284	1.283
Number of bathrooms	1.169	1.163
Terrace/Balcony	0.292	0.310
Type of bed		
Couch	0.006	0.006
Airbed	0.003	0.002
Sofa	0.031	0.032
Futon	0.011	0.011
Amenities		
Cable TV	0.376	0.381
Wireless	0.933	0.949
Heating	0.919	0.932
Air Conditioning	0.387	0.389
Fireplace	0.086	0.087
Washer	0.729	0.731
Dryer	0.399	0.405
Elevator	0.356	0.347
Doorman	0.100	0.093
Parking	0.183	0.184
Gym	0.061	0.058
Pool	0.055	0.050
Buzzer	0.399	0.415
Hot Tub	0.068	0.067
Wheelchair Accessible	0.103	0.106
N	3,518,984	2,446,877

Notes: (i) Active listings correspond to listings which receive at least one review over the observation period; (ii)

(number of accommodates/bedrooms/bathrooms...). Table 9 (Appendix B) provides details on how amenities affect the price.

Table 2: Summary statistics: Property and host characteristics (continued)

	Full Sample	**Active Listings**
Services		
Breakfast Served	0.097	0.099
Family/Kid Friendly	0.469	0.488
Suitable for Events	0.059	0.060
Rules & Extras		
Pay for Additional People	0.681	0.794
Price per additional people	8.333	9.180
Cleaning Price	30.722	32.499
Smoking allowed	0.331	0.341
Pets allowed	0.325	0.337
In couple	0.067	0.078
Has multiple properties	0.373	0.389
Member since 2008-09	0.011	0.012
Member since 2010-11	0.153	0.164
Member since 2012-13	0.513	0.531
Verified Email	0.956	0.972
Number of languages spoken	1.409	1.517
N	3,518,984	2,446,877

Notes: (i) Active listings correspond to listings which receive at least one review over the observation period; (ii)

Figure 2: Distribution of nightly price

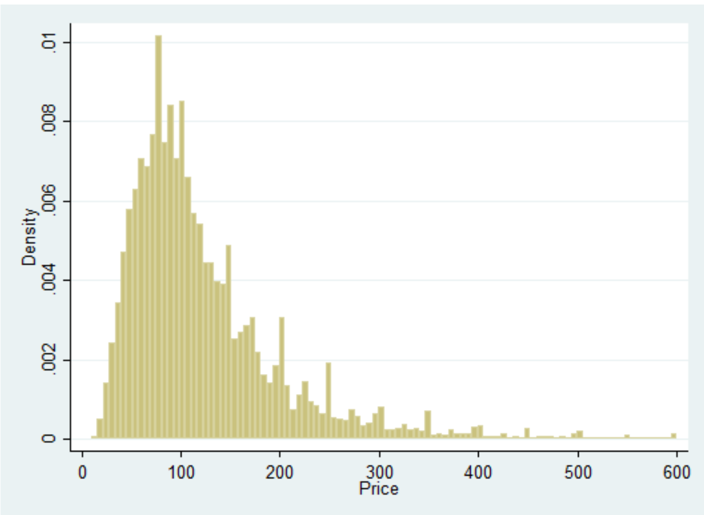

2.3 Ethnic groups and gaps

We consider three groups of ethnic minorities : Arabic/African, African-American and Hispanic. The group of Arabic/African names is determined by comparing hosts' first names to the list provided in Jouniaux (2001). To identify African/American hosts, we use 'distinctively black' first names from Fryer and Levitt (2004); Bertrand and Mullainathan (2004). These 'distinctively black names' are correlated to socioeconomic conditions of parents, so we are actually collecting all pictures from profiles to assess the race of owners to reduce this bias. Finally, the group of Hispanic is determined via popular websites that help parents pick kids' names based on ethnicity and origin.In total, we obtain a list of 7000 Arabic/African names, 15000 Hispanic names, and 1000 African/American names.

Table 3 displays the share of ethnic groups in the sample and the within-city*wave raw price gap. First, it shows the four minority groups are not evenly represented in our sample. While the share of Arabs/Africans in Europe and the share of Hispanic in North America are both relatively high, with 2% and 3% respectively, the proportions of Arabs/Africans and Afro-Americans in North America are low, with 1% and 0.5% respectively. In the third column, we estimate the 'raw gap' between majority and each ethnic group in daily prices, only controlling for heterogeneity across cities. The four raw ethnic gaps are quite large and slightly vary across groups, from 8% for Arabs/Africans in Europe to 11% for hispanics in North America. Overall, the share of minorities is 6.7%.

Table 3: Samples and Raw Gaps by ethnic groups

	Sample size	Share	Within-city*wave gap
Majority	2,340,193	93.3%	-
Arabic/African (Eur)	52,069	2.1 %	8.2%
Arabic/African (US/Can)	28,042	1.1%	8.6%
African-American (US/Can)	14,025	0.6%	9.6%
Hispanic (US/Can)	73,811	2.9%	11.4%

Table 4 displays the ethnic price difference between 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 (9.7%) and highly significant. Accounting for ethnic disparities in property observable characteristics reduces the gap to 5% (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 neighborhood fixed-effects instead of city fixed-effects reduces the ethnic price gap from 9.7% with no fixed-effects to 6.6% (column 3). This indicates that ethnic minorities tend to own properties 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 3.0% but is still highly significant. Note that the adjusted R-squared is high in this last specification, equal to .72. Observables are found to explain the largest part of the variance, as the adjusted R-squared is equal to .54 in the second column.

Table 4: Ethnic price gap, by specification

	Log daily rate			
	(1)	(2)	(3)	(4)
Minority	-0.097***	-0.052***	-0.066***	-0.030***
	(0.002)	(0.001)	(0.002)	(0.001)
City*Wave FE	Yes	Yes	Yes	Yes
Neighborhood FE	No	No	Yes	Yes
Property characteristics	No	Yes	No	Yes
N	2,502,771	2,444,928	2,502,771	2,444,928
Adj R^2	0.17	0.65	0.33	0.72

Notes: (i) See the list of all property and host characteristics in Table 9; (ii) Standard errors in parentheses. *** $p < 0.01$.

3 Conceptual framework

3.1 Prices and demand as a function of quality

At each period (say, a week), an owner 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.

Owners 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, owners 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 owner's group.[3] 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.[4] 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, an owner with outside option w will set a price:

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

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:

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

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

[4] 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.

3.3 Evolution of the price with the number of reviews

In order to be able to identify statistical discrimination, we rely on the fact that reviews bring information. First, we need to have some variability in the number of reviews. Figure 3 displays the distribution of the number of reviews in the sample. The sample offers a decent amount of heterogeneity in the number of reviews, the empirical distribution being quite similar to that of a Poisson random variable. Observations with more than 40 observations represent about 9% of the sample.

Figure 3: Distribution of reviews

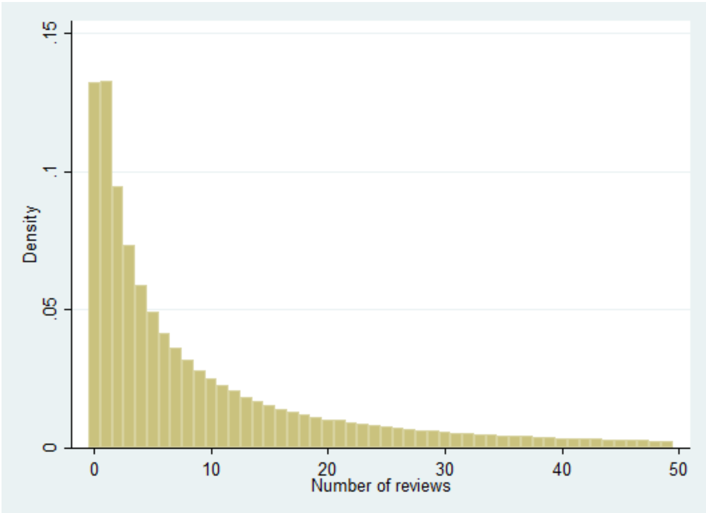

According to our conceptual framework, hosts should update their prices as new information is available about the quality of their properties, i.e. as the number of reviews increases. An additional review providing less information less than a previous one, the model predicts a concave relationship between the price and the number of reviews, converging to some value when the number of reviews tends to infinity. Figure 4 provides an illustration of this Bayesian updating phenomenon from a simulation in our model.

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: 3.5, 4, 4.5 and 5 stars. 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 account for the hypothetical concavity of

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

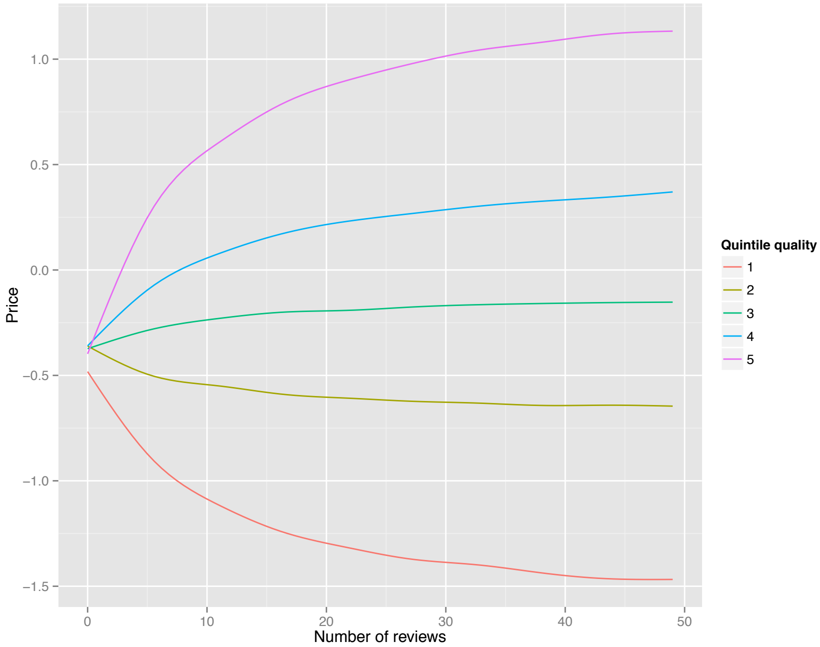

the relationship between prices and number of reviews.

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

The results of the estimation are displayed in Figure 5. The figure shows that, depending on the last rating, the prices diverge in the way predicted by our conceptual framework. This is evidence for the fact that reviews provide information to potential guests, that hosts use reviews and information to update their prices, and the last rating is indeed a proxy for the unobservable quality of the listing.

Figure 5: Estimation: Prices according to the number of reviews, by most recent rating

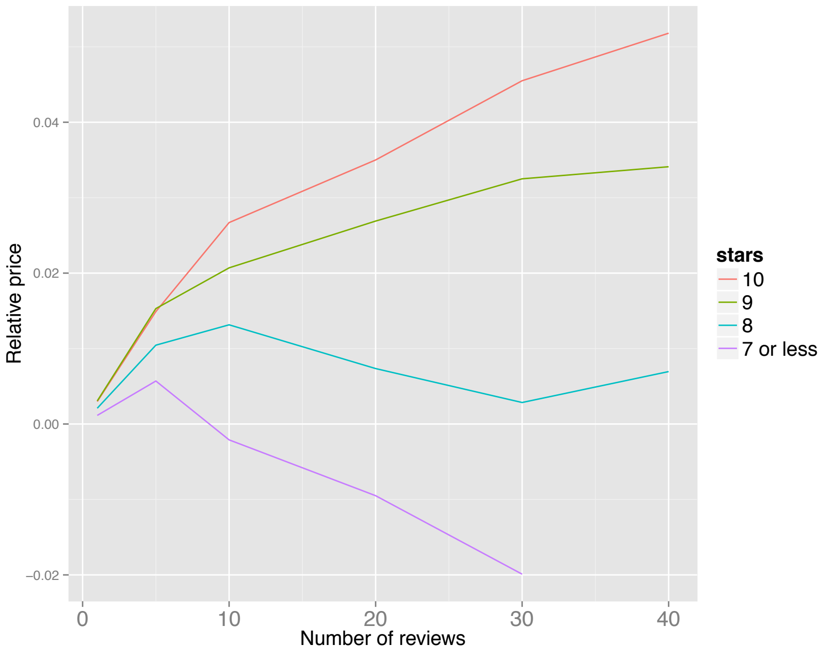

4 Ethnic price gaps and statistical discrimination

We first document how the unexplained ethnic price gap changes with the number of reviews. Table 5 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 4.3% for listings with no reviews to 0% for listings with more than 40 reviews. While this pattern could be interpreted as suggestive evidence of statistical discrimination, it is subject, as any cross-section analysis to selection issues. A possible story is that potential guests only stays at minority if the quality is extremely good, while they are less demanding for majority-owner listings. In this case, the ethnic gap would be reduced, not because of the existence of statistical discrimination, but simply because the minority-owner listings with than 40 reviews are relatively much better than those with a small number of reviews. However, this story would predict a drop in the share of minority listings with the number of reviews. We observe that the share of minority remains stable around 6.7% in all columns.

Table 5: Ethnic price gap, for several segments of the number of reviews

	Log daily rate				
	(1)	(2)	(3)	(4)	(5)
Minority	-0.043***	-0.035***	-0.022***	-0.006**	-0.000
	(0.004)	(0.002)	(0.002)	(0.003)	(0.004)
Nb reviews	0	1-4	5-14	15-39	40+
Neighborhood FE	Yes	Yes	Yes	Yes	Yes
Property characteristics	Yes	Yes	Yes	Yes	Yes
Ratings	-	Yes	Yes	Yes	Yes
Minority share	6.8%	6.7%	6.7%	6.6%	6.7%
N obs.	331,410	805,217	653,035	434,353	220,913
Adj R^2	0.699	0.736	0.760	0.773	0.761

Notes: (i) See the list of all property and host characteristics in Table 9; (ii) Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Still, more sophisticated forms of selection could accommodate both 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 6 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 for when the number of reviews is smaller.

The results in Table 6 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 owners 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. This supports supports the hypothesis of a concave relationship between prices and the number of reviews.

The previous results show 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 the last observation $\dot{p}$.

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

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

Table 6: Estimation of the model in difference between the first and last observations

	Variation of log-price		
	(1)	(2)	(3)
3.5 stars	-0.022***	-0.025***	-0.029***
	(0.004)	(0.004)	(0.004)
4 stars	-0.013***	-0.016***	-0.020***
	(0.003)	(0.003)	(0.002)
4.5 stars	0.011***	0.009***	0.007***
	(0.002)	(0.002)	(0.002)
Minority	-0.001	-0.003	-0.001
	(0.003)	(0.003)	(0.003)
3.5 stars $\times \Delta K/100$	-0.007	-0.026	-0.030
	(0.072)	(0.061)	(0.055)
4 stars $\times \Delta K/100$	-0.036	-0.055**	-0.035*
	(0.028)	(0.022)	(0.018)
4.5 stars $\times \Delta K/100$	0.125***	0.056***	0.001
	(0.014)	(0.010)	(0.008)
5 stars $\times \Delta K/100$	0.317***	0.211***	0.119***
	(0.015)	(0.011)	(0.009)
Minority $\times \Delta K/100$	0.113***	0.100***	0.067***
	(0.033)	(0.026)	(0.020)
Sample	$\min(K) \leq 5$	$\min(K) \leq 20$	Full
Adj R^2	0.119	0.127	0.132
N obs.	204,489	240,724	259,476

Notes: (i) All regressions include city*wave FE, neighborhood FE and property characteristics (See Table 9) ; (ii) Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

where K_0 and K_1 are the number of reviews at the first and last observations.

We obtain an estimated value of 22 for ρ. This can be interpreted as the fact that 22 reviews are necessary to reveal half of the information concerning a listing about the ν (the part of quality that is revealed through reviews and ratings). β_m is estimated to .048, which means that going from 0 to an infinite number of reviews increases the prices of minority by 4.8%. This figure is of the same order of magnitude as the ethnic price gap observed in the subset of listings with no reviews (4.3%, see Table 5, column 1).

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

	Log Demand: Variation of reviews between 2 waves			
	(1)	(2)	(3)	(4)
log price	-0.140***	-0.105***	-0.128***	-0.126***
	(0.001)	(0.001)	(0.002)	(0.001)
Minority	-0.004	0.001	-0.002	-0.004
	(0.002)	(0.002)	(0.003)	(0.003)
$K/100$		0.636***		
		(0.002)		
Samples	Full	Full	reviews<5	reviews<20
obs.	2,181,215	2,181,215	1,455,320	1,926,000

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 hosts belonging to minority groups have on average lower outside options than majority ones. This relates, in our conceptual framework, to a lower w. A lower outside wage entails a lower price but it should also entail 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.

Table 7 presents the result of the regression of the proxy for the log-demand on the log-price and a minority dummy (and the number of reviews at the previous previous for the specification presented in column 2), controlling for location and observable characteristics. In columns 1 and 2, the full sample is used while in column 3 and 4 we focus on observations with less than 5 and 20 reviews. In all columns, we find that the coefficient of the minority is close to zero and insignificant. This is evidence that the ethnic price gap does not reflect differences in pricing behavior induced by differences in outside wages.

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. We investigate this issue, we first code the ethnicity of guests leaving reviews on each listing, using the same procedure as for hosts. In a second step, we regress for each listing and wave, the number of new reviews from each of the ethnic groups (Arabic/African, African-American, Hispanic) on dummies of 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. The magnitudes are similar for African Americans and Hispanic first names. Overall, despite the mild ethnic matching, our results support the assumption that hosts belonging to different ethnic groups compete on the same market.

6 Conclusion

This paper shows that, in a popular online platform of short-term rentals, owners 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. The value of the gap in properties without reviews and the estimate of the statistical-discrimination parameter in our model both point to a value around 4.5% instead of the current 3%, which represents a sizable gain in proportion. 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.

While there is no consensus about the sources of ethnic gaps in employment and wages on the labor market, our findings mirrors those obtained by Agrawal et al. (2014) on the online platform ODesk. They find that standardized information about work performed on the platform disproportionately benefits less-developed-country contractors, relative to developed-country ones.

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.

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.

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.

Charles, 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.

Charles, K. K. and J. Guryan (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.

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.

Fryer, R. G. and S. D. Levitt (2004): "The Causes and Consequences of Distinctively Black Names," *The Quarterly Journal of Economics*, 119, 767–805.

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.

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.

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.

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

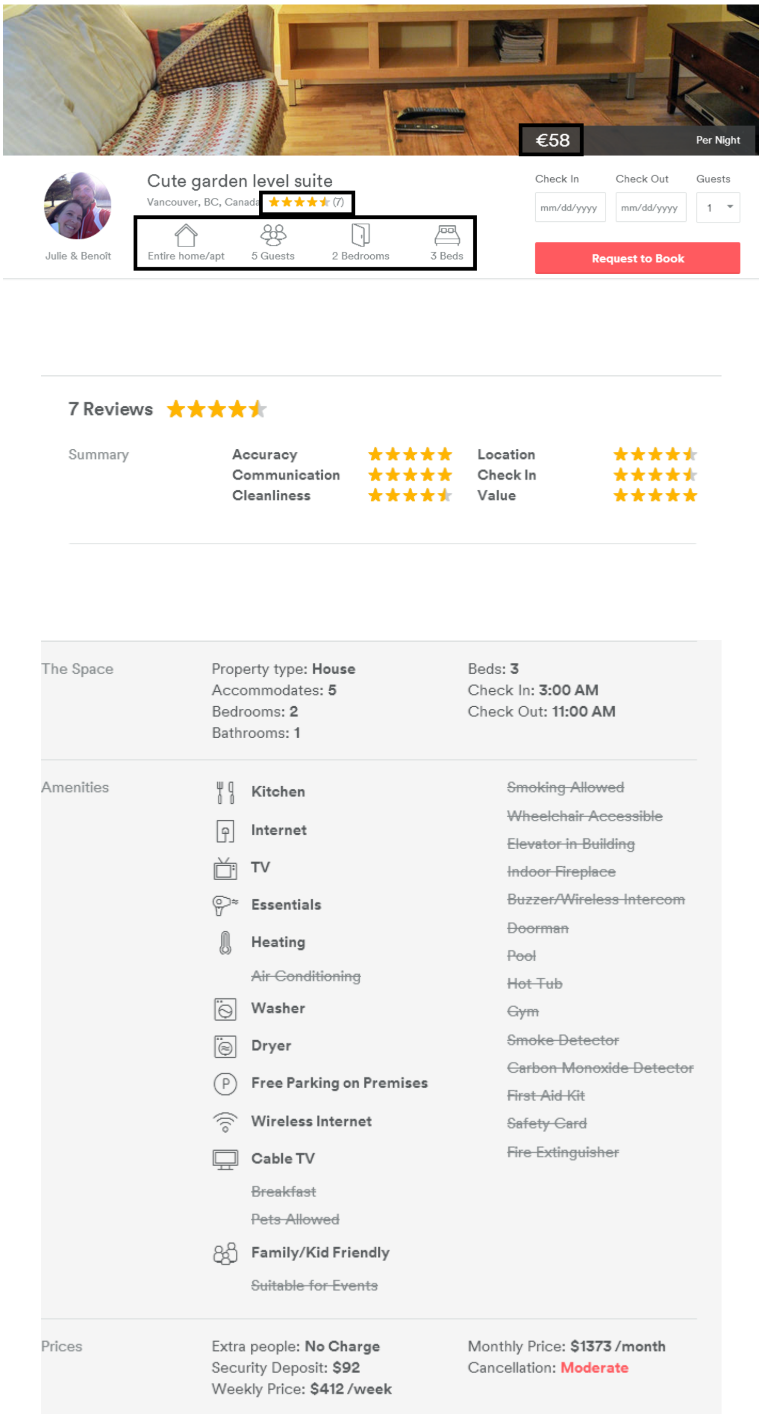

A.2 Peer-reviewing System

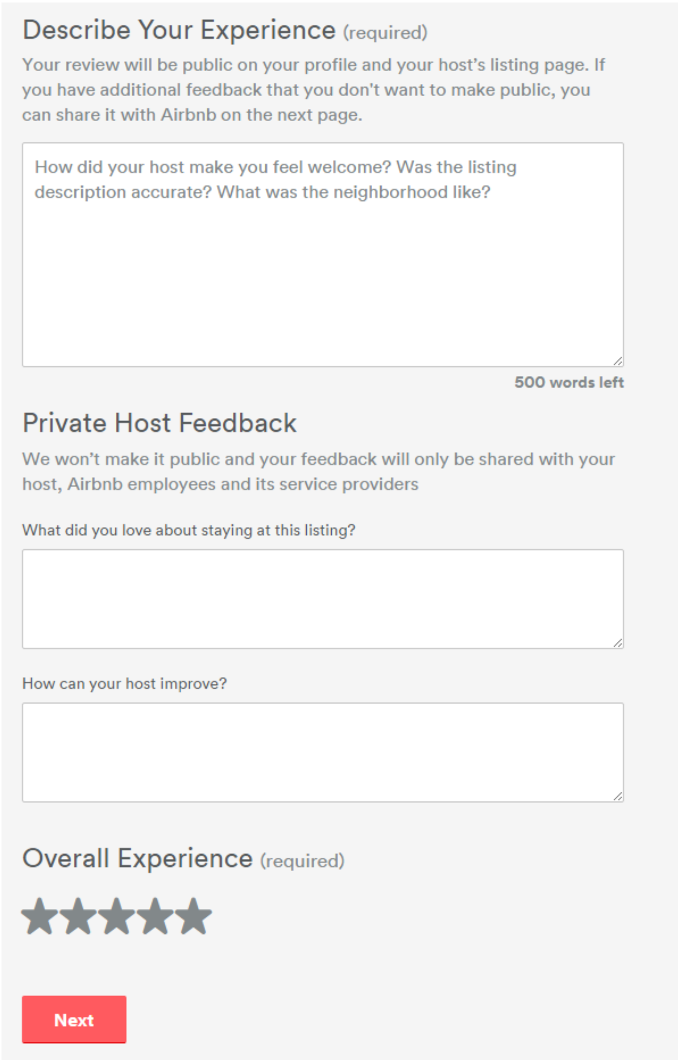
Describe Your Experience (required)
Your review will be public on your profile and your host's listing page. If you have additional feedback that you don't want to make public, you can share it with Airbnb on the next page.
How did your host make you feel welcome? Was the listing description accurate? 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 8: 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 9 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 9: Log daily rate

Shared flat	-0.791***	-0.709***
	(0.001)	(0.001)
Person Capacity (> 2)	0.157***	0.167***
	(0.001)	(0.001)
Nber bedrooms	0.297***	0.319***
	(0.001)	(0.001)
Nber bathrooms	0.097***	0.086***
	(0.001)	(0.001)

Continued on next page

Table 9: Log daily rate

Flat	-0.175***	-0.197***
	(0.002)	(0.002)
House or Loft	-0.160***	-0.073***
	(0.002)	(0.002)
Couch	-0.160***	-0.130***
	(0.005)	(0.004)
Airbed	-0.173***	-0.138***
	(0.008)	(0.007)
Sofa	-0.167***	-0.156***
	(0.002)	(0.002)
Futon	-0.148***	-0.110***
	(0.003)	(0.003)
Terrace or Balcony	0.033***	0.042***
	(0.001)	(0.001)
Cable TV	0.123***	0.090***
	(0.001)	(0.001)
Wireless	0.034***	0.022***
	(0.002)	(0.002)
Heating	-0.018***	0.001
	(0.002)	(0.001)
AC	0.163***	0.139***
	(0.001)	(0.001)
Elevator	0.092***	0.087***
	(0.001)	(0.001)
Wheelchair Accessible	-0.041***	-0.012***
	(0.001)	(0.001)
Doorman	0.103***	0.052***
	(0.001)	(0.001)
Fireplace	0.166***	0.132***
	(0.001)	(0.001)
Washer	-0.039***	-0.000
	(0.001)	(0.001)
Dryer	0.148***	0.101***
	(0.001)	(0.001)
Parking	-0.140***	0.009***
	(0.001)	(0.001)
Gym	0.057***	0.049***
	(0.002)	(0.002)

Continued on next page

Table 9: Log daily rate

Pool	0.099***	0.126***
	(0.002)	(0.002)
Buzzer	0.048***	0.011***
	(0.001)	(0.001)
Hot Tub	0.022***	0.018***
	(0.002)	(0.001)
Breakfast served	0.026***	0.051***
	(0.001)	(0.001)
Family/Kids Friendly	0.003**	0.020***
	(0.001)	(0.001)
Suitable for events	0.099***	0.092***
	(0.002)	(0.001)
Additional People	-0.036***	-0.015***
	(0.000)	(0.000)
Price per Additional People	0.000***	-0.001***
	(0.000)	(0.000)
Cleaning price	0.004***	0.003***
	(0.000)	(0.000)
Cancellation Policy	0.005***	-0.011***
	(0.000)	(0.000)
Smoking Allowed	-0.045***	-0.035***
	(0.001)	(0.001)
Pets Allowed	-0.007***	-0.009***
	(0.001)	(0.001)
Host in couple	-0.023***	-0.007***
	(0.001)	(0.001)
Host has multiple properties	0.028***	0.005***
	(0.001)	(0.001)
Member since 2008-2009	0.065***	0.049***
	(0.003)	(0.003)
Member since 2010-2011	0.044***	0.030***
	(0.001)	(0.001)
Member since 2012-2013	0.020***	0.014***
	(0.001)	(0.001)
City*Wave FE	Yes	Yes
Neighborhood FE	No	Yes
Property characteristics	Yes	Yes
Adj R^2	0.649	0.723

Continued on next page

Table 9: Log daily rate

N obs.	2,444,928	2,444,928

City	Obs	Share
Amsterdam	117,923	4.70
Barcelona	184,672	7.36
Berlin	157,874	6.29
Boston	42,903	1.71
Chicago	43,656	1.74
Florence	61,456	2.45
London	261,441	10.42
Los Angeles	158,128	6.30
Madrid	71,063	2.83
Marseille	54,598	2.18
Miami	68,002	2.71
Milan	71,259	2.84
Montreal	71,395	2.85
New-York	348,466	13.89
Paris	445,742	17.77
Rome	140,873	5.62
San-Francisco	108,532	4.33
Toronto	54,796	2.18
Vancouver	45,361	1.81