



Offline biases in online platforms: a study of diversity and homophily in Airbnb

Victoria Koh¹, Weihua Li^{1,2}, Giacomo Livan^{1,2*}  and Licia Capra¹

*Correspondence: g.livan@ucl.ac.uk

¹Department of Computer Science, University College London, London, UK

²Systemic Risk Centre, London School of Economics and Political Science, London, UK

Abstract

How diverse are sharing economy platforms? Are they fair marketplaces, where all participants operate on a level playing field, or are they large-scale online aggregators of offline human biases? Often portrayed as easy-to-access digital spaces whose participants receive equal opportunities, such platforms have recently come under fire due to reports of discriminatory behaviours among their users, and have been associated with gentrification phenomena that exacerbate preexisting inequalities along racial lines. In this paper, we focus on the Airbnb sharing economy platform, and analyse the diversity of its user base across five large cities. We find it to be predominantly young, female, and white. Notably, we find this to be true even in cities with a diverse racial composition. We then introduce a method based on the statistical analysis of networks to quantify behaviours of homophily, heterophily and avoidance between Airbnb hosts and guests. Depending on cities and property types, we do find signals of such behaviours relating both to race and gender. We use these findings to provide platform design recommendations, aimed at exposing and possibly reducing the biases we detect, in support of a more inclusive growth of sharing economy platforms.

Keywords: Sharing Economy; Social Networks; Homophily; Online User Behavior; Statistical Validation

1 Introduction

Sharing economy platforms are new manifestations of century old phenomena. Resource circulation systems that facilitate the exchange of underutilized goods or services between consumers have long existed in the form of flea markets, garage sales, second-hand shops, just to name a few. However, what used to be small scale and local instances of collaborative consumption, have now become massive online marketplaces, where face-to-face interactions have been replaced by technology-mediated ones [1].

A fundamental question arises about the role that such decentralized, largely unregulated, online platforms play in our societies. Often thought of as level-playing fields, where all participants receive the same opportunities, sharing economy platforms might instead end up acting as *online* aggregators of well-known *offline* human dynamics and biases. Indeed, a number of studies have suggested that some of the big sharing economy players are acting as accelerators of gentrification phenomena that are already underway in large cities. For example, Airbnb has led to the emergence of short-term rent gaps between different areas of New York City [2] and has contributed to exacerbating the affordable

housing crisis in Los Angeles [3]. These phenomena, in turn, typically accelerate preexisting divides along racial lines, fostering inequalities between the Airbnb community and the communities living in the neighbourhoods where Airbnb has a significant presence (see, e.g., [4]).

Moreover, sharing economy platforms have come under multiple allegations over discrimination episodes taking place within the platforms themselves. For example, Uber drivers were found to be twice more likely to cancel trips requested by passengers with African–American sounding names compared to White-sounding names (even though Uber penalizes drivers for cancellations) [5]. Similarly, Airbnb’s hosts were found to be turning down potential guests based upon their racial background [6].

While headway has been made to tackle the unintended consequences brought about by sharing economy platforms, the debate on their socio-economic impact is still in its infancy, and only relies on a handful of studies or on anecdotal evidence. This, in turn, delays the execution of targeted interventions to expose, and possibly reduce, such consequences. The goal of this paper is to contribute to inform such a debate by performing a large scale empirical analysis aimed at detecting systematic statistical evidence of ‘offline’ biases taking place in online sharing economy platforms.

We study the Airbnb hospitality service, and focus first on the composition of its user base, with the aim of assessing its diversity both in general terms and then contextually with respect to the city hosting it. Second, we employ a network methodology to assess the statistical significance of host–guest interactions in Airbnb. In particular, we focus on *homophily*, i.e. the social phenomenon where people gravitate towards those like themselves [7], and on its opposite, heterophily. We also study the tendency to avoid members of a social group with different social traits, which we refer to as *avoidance*. While avoidance is universally deemed as unacceptable, homophily has sometimes been perceived as ‘natural’, and thus judged in a more accepting way. However, several studies have shown that the aggregation of slightly biased individual preferences can lead to unintended and collectively undesirable consequences, as evidenced by Schelling’s work on urban racial segregation [8, 9], and Neal’s work on school children’s development [10].

In performing this study, we make three contributions:

- We gather data about Airbnb hosts, guests, and their interactions for five cities, spanning three different continents (Airbnb Data section). These are Amsterdam (The Netherlands), Dublin (Ireland), Hong Kong (China), Chicago and Nashville (U.S.). We have chosen them so to cover geographically (and culturally) different cities, as well as to cover variances in size, population composition, and cost of living.
- We study the diversity of the Airbnb user base in the above five cities along the dimensions of gender, age, and race. We find the Airbnb community to be predominantly female, and overwhelmingly young and White. In line with the aforementioned literature, we find the majority of hosts to be White even in cities whose racial composition is significantly more diverse (Results section).
- We model Airbnb’s peers and interactions as nodes and edges in a bi-partite graph, and use a statistical method based on network rewiring to systematically identify edges (i.e., guest–host pairings) that cannot be attributed to chance (Method section). We apply such a method to the five cities under study, and, depending on the specific city and property type, find signals of homophily, heterophily and avoidance. We find such signals to be rather strong in the case of gender, rather weak (although still

statistically significant) in the case of race, and mostly absent in the case of age (Results section).

These results echo other findings in the literature (see next section), and provide concrete evidence about how sharing economy platforms are being appropriated in different city contexts, possibly resulting in large divides between the online communities who can enjoy the benefits of the sharing economy and the 'offline' urban communities who are most exposed to its expansion. They also offer an opportunity to inform the design of tailored technology interventions aimed at exposing, and possibly reducing, certain behaviours, while also providing the means to monitor their effects (Discussion section).

2 Related work

Upon its inception, the Internet was expected to create a global level playing field, where the inequalities of the 'offline' world would be overcome thanks to easy access to digital opportunities. Yet, reality has been very different. As it is invariably the case, different social groups are not equally equipped to face technological innovation in its early stages, which typically exacerbates preexisting inequalities [11].

The sharing economy, as a whole, has been no exception. Indeed, a handful of studies have shown that the ability to seize the sharing economy's opportunities is often severely limited by geographical and socio-economic constraints. For example, Airbnb listings are usually more concentrated in wealthier, more attractive areas populated by young and tech-savvy residents [12]. Similarly, TaskRabbit users from areas with low socioeconomic status and/or low population density were found to have a harder time both when selling their services and when seeking to outsource work to potential taskers [13], while individuals living in deprived Chicago suburbs have been found to have a harder time to get an Uber ride [14].

Similarly, the Internet's promise to circumvent physical barriers and improve communication between social groups has not always been upheld. For example, episodes of racial discrimination in online social networks have been extensively documented [15, 16]. Also, a vast amount of scholarly work has been devoted to understanding the formation of online preferential relationships between individuals. This has often been explained either in terms of interest-based homophily, e.g., showing the impact of ideological homophily in determining the opinions and content individuals are exposed to on social media [17], or in terms of homophily driven by *demographics*.

Studies of early social networks, e.g., MySpace [18], have identified race, gender, and age as the main demographic features driving online homophily, and such elements kept recurring in more recent studies. Indeed, evidence of racial [19] and gender [20] homophily has been reported in Facebook and Twitter, respectively, and evidence of both has been documented in the social networks underpinning location sharing applications [21]. Age homophily is somewhat less studied, but still documented in a study of the Facebook social graph [22] and in niche environments such as virtual worlds [23, 24].

Our work follows this stream of literature and investigates whether well known 'offline' biases also take place in sharing economy platforms. A handful of recent studies have started to look at such platforms from this perspective. Indeed, recently published work [5] found evidence of both gender and racial discrimination in Uber and Lyft, as female passengers were disproportionately taken on longer and more expensive routes, while passengers with African American-sounding names were twice as likely to receive trip cancellations from Uber drivers compared to passengers with White-sounding ones (even

though Uber penalizes drivers for cancellations). Similarly, another study [25] found gender and race to have an impact on worker evaluations in online freelance marketplaces.

Evidence of biased behaviour was also found in Airbnb by means of a field experiment [6]. In particular, guests were found to be 16% less likely to have their booking accepted if they had a distinctly African American-sounding name when compared to identical guests with White-sounding names instead. Similarly, in [26] it was found that non-Black hosts charged on average 12% more for an equivalent rental compared to Black hosts, and similar results were replicated in a subsequent study on Airbnb [27], where Asian and Hispanic hosts were found to rent at prices 9.3% and 9.6% lower, respectively, than their White counterparts.

While the above works investigate some specificities of user demographics and interactions in sharing economy platforms, a systematic analysis of these dimensions across the fundamental features of gender, age, and race is still lacking. This work aims at filling this void, by providing (i) an overview of the composition and diversity of Airbnb's community, and (ii) a quantitative method to dissect the anatomy of user-user interactions in sharing economy platforms (and Airbnb in particular), providing statistical evidence of homophily and avoidance between certain user groups.

3 Airbnb data

In order to perform this study, we needed two types of data: demographic characteristics of hosts and guests (i.e., gender, age, race); and their pairing dynamics (i.e., who stayed with whom). Since we hypothesise that peers' behaviours might vary in different geographic (and cultural) contexts, we chose to perform this study on a per city level, rather than treating the whole of Airbnb as a single analytical context.

To begin with, we accessed city snapshots that the website InsideAirbnb^a already makes available. We chose five cities (Amsterdam, Chicago, Dublin, Hong Kong, Nashville) so to have high geographic diversity (these cities span three different continents), as well as high diversity in terms of population composition and cost of living. Records of Airbnb hosts, guests and stays go from 2008 to 2016 for all cities except Nashville, whose Airbnb records start in 2009. For each city, InsideAirbnb makes available a full list of host IDs (from their 'listings' file). We used these IDs to query the Airbnb website and further acquire a host profile picture, the type(s) of property they were renting out (i.e., full property, private room in a shared property, or shared room), and the full list of IDs of all guests that ever left a review to such host (and for what property). We then further queried the Airbnb website with the guest IDs to acquire their profile pictures. Since Airbnb does not explicitly make available a peer's gender, age and race as attribute-value pairs in the peers' profile, we used image processing software on the collected profile pictures to automatically extract this information. In particular, we first used face localisation software to detect whether the profile picture contained a human face, and if so, to identify the portion in the picture containing it. We tested both FaceReact^b and Indico^c on a manually curated sample of 50 Airbnb images, so to contain a mix of pictures with and without human faces, and with and without background clutter. We found Indico to be significantly more accurate, especially for human images taken at an angle rather than straight-facing the camera. We thus continued only with the latter. Having extracted the bounding box containing a human face, we then used face recognition software to extract attributes. We tested Betaface,^d Sightcorp F.A.C.E.,^e and Face++^f on a subset of 250 Airbnb images. We found all three to

Table 1 Number of hosts, guests, and host–guest pairs annotated for each city analysed

City	# Hosts	# Guests	# Host–Guest pairs
Amsterdam	2369	69,923	71,779
Chicago	1706	21,105	22,493
Dublin	1039	2618	2785
Hong Kong	1233	12,103	13,330
Nashville	630	1712	2017

be equally accurate when detecting gender. Sightcorp was found to be significantly more reliable in recognising age groups, and Betaface in extracting race (note that our analyses will focus exclusively on race, not on ethnicity; in particular, we will focus on three main race categories, i.e., White, Black, and Asian). We thus worked with Sightcorp and Betaface in parallel. We manually verified their accuracy on all 250 test images, and found the confidence levels reported by both products to be $0.3 \in [0, 1]$ or higher on images annotated correctly. Hence, we kept such value as a threshold for the ensuing automatic annotation; furthermore, we only retained pictures for which both face recognition software products agreed on both gender and ethnicity. To understand how robust our results are when varying facial annotation accuracy, we repeated all our analyses after (i) increasing the above threshold to 0.5, and (ii) manipulating the data by changing the race annotation on a random sub-sample of the images. The results obtained from such analyses are reported in the Additional file 1.

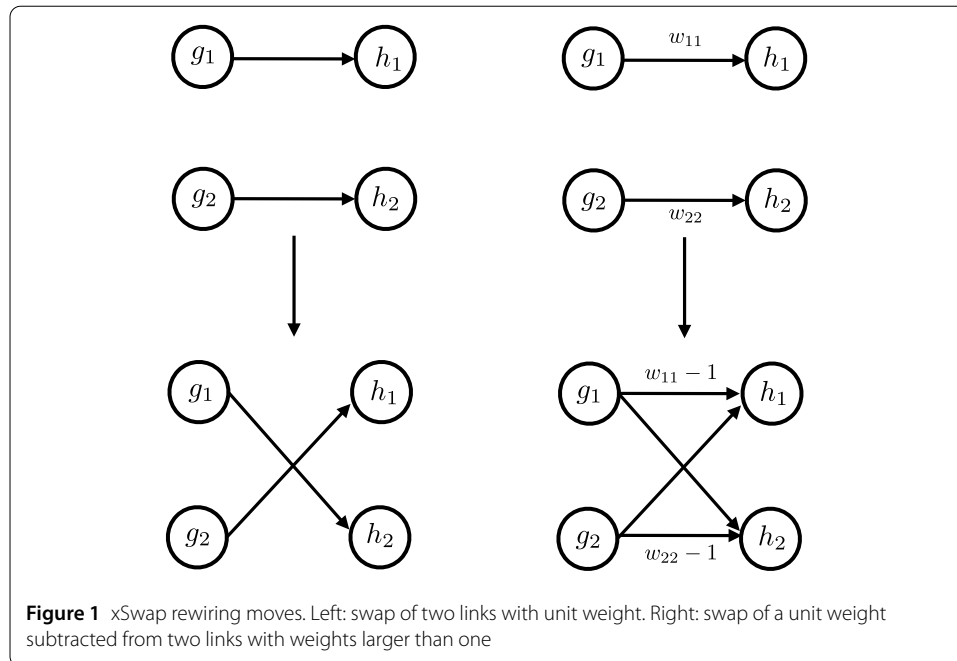
In terms of pairing dynamics, Airbnb does not make visible who stays with whom, nor whether a stay request has been refused or cancelled. However, what it does make visible are reviews that hosts and guests leave to one another. We use these as proxies for the actual pairing dynamics. Studies show that over 65% of stays result in a guest review and 72% result in a host review [28], so most stays are indeed captured by reviews. At present, it is not known whether those who do *not* leave reviews in Airbnb belong to specific users' groups; a past survey study of Tripadvisor reviews [29] did find that certain age and gender groups were more vocal than others, and this might also be the case in this context. Although the method we present next is still applicable, the validity of some of our findings might be impacted, and we will come back to this when we discuss limitations and future work (Conclusion section).

Summary statistics about the number of hosts, guests and pairings that we collected and annotated for each city under study are reported in Table 1.

4 Method

We model Airbnb hosts and guests as nodes in a bi-partite graph, with a directed $g \rightarrow h$ edge with weight w_{gh} representing the number of times guest g stayed at host h . Since we hypothesise that pairing dynamics may vary across cities, as well as across type of rented property (full property vs. shared—the latter comprising both private and shared rooms), we create and analyse a total of 10 (5-cities \times 2-property types) bipartite networks.

Each such network is analysed using a statistical *rewiring* approach designed to assess the significance of pairing patterns in each of the cities studied. More precisely, the method starts from a null hypothesis that a given guest–host pairing occurred randomly. It then proceeds to verify whether this hypothesis holds by creating ensembles of null network models through the rewiring of the original networks' edges, and by comparing the properties of such null network model against those observed in the actual, empirical networks.



Crucially, the procedure is designed to preserve the heterogeneity of the original networks, as it produces null network configurations where the number of stays that each guest and host have had are both kept intact, therefore preserving correlations between demographic features and activity on Airbnb. In the following, we describe the details of this methodology, and the rationale for adopting it.

Starting from an empirical bipartite network, we create a randomised version of it by iteratively performing xSwap operations [30]. These amount to selecting two guest nodes (g_1 and g_2) and one of their corresponding host nodes (h_1 and h_2 , respectively) at random, erasing the existing edges ($g_1 \rightarrow h_1$ and $g_2 \rightarrow h_2$) between both pairs, then reassigning them to each other ($g_1 \rightarrow h_2$ and $g_2 \rightarrow h_1$), as shown in Fig. 1 (left). Should either of the selected edges have a weight larger than one, the strength of the link is reduced by one and a unit weight is redistributed to the new host node (see Fig. 1-right). These operations preserve both the outgoing weight of guest nodes and the incoming weight of host nodes. Therefore, repeated xSwap operations yield configurations which are exactly equivalent to the original networks in terms of their heterogeneity, but are instead fully randomized in terms of the relationships between their nodes [31]. In order to determine how many xSwap operations were needed before the rewired network configurations could be considered distinct enough (i.e., sufficiently uncorrelated) from their original counterparts, we computed the Kendall's Correlation Coefficient as suggested in [32].

For each of the 10 empirical networks under consideration, we generated an ensemble of 1000 randomized network configurations with the above rewiring procedure, thus preserving the original networks' degree sequences. Since such a rewiring procedure generates null configurations that do not deviate substantially from the original networks themselves, we can resort to a fairly parsimonious numerical investigation based on such relatively small number of null configurations. On such configuration ensemble, we computed the frequency of interaction between different *groups* of hosts and guests based on gender, age, and race (e.g., female guests to male hosts, White guests to White hosts, etc.).

Table 2 Airbnb host and guest population by gender (F = female)

City	Full property rental		Shared property rental	
	F Host	F Guest	F Host	F Guest
Amsterdam	59%	55%	59%	58%
Chicago	61%	58%	57%	56%
Dublin	61%	59%	58%	54%
Hong Kong	56%	59%	60%	59%
Nashville	63%	66%	74%	59%

Table 3 Airbnb host population by race (W = White, A = Asian, B = Black)

City	Full property rental			Shared property rental		
	W	A	B	W	A	B
Amsterdam	93%	5%	2%	88%	10%	2%
Chicago	90%	8%	2%	87%	10%	3%
Dublin	96%	4%	1%	92%	7%	1%
Hong Kong	64%	35%	2%	58%	40%	2%
Nashville	92%	3%	5%	92%	5%	2%

We then took the 95% confidence interval to define lower and upper bound probabilities for the occurrence of each pairing combination. This range represents the expected probabilities of pairing combinations taking place by chance. We then computed the pairing probabilities among the same groups in the original empirical network: should the actual values fall below or above the ‘by-chance’ range, we take them to be statistically significant under/over-expressions of certain group interactions with respect to the null hypothesis. Note that the over-expression of interactions between two groups (or within a group) does not necessarily translate into the under-expression of interactions between other groups; we will use the former to detect a tendency of certain groups to *associate*, and the latter to separately detect a tendency to *avoid* certain groups instead.

5 Results

In this section, we report the results obtained when applying the above method to analyse Airbnb users’ pairing dynamics. Before doing so, we provide an overview of the demographic characteristics of Airbnb hosts and guests, whose profile pictures we scraped and annotated using the process described in the Airbnb Data section. We will discuss and elaborate on these results in the Discussion section.

5.1 Demographics

In Table 2 we report summary statistics of Airbnb users in terms of *gender*, broken down by city and by property rental type (full vs shared). As shown, Airbnb has a predominantly female user base: in all cities under study, the net majority of both hosts and guests were found to be female, regardless of the property type, with such proportion getting close to or exceeding 60% in a number of cases.

We next consider *race*, once again as it varies by city and by property rental type (see Table 3). We focus our attention on hosts only in this case, so to compare our results with available census information on the demographics of the cities under study. In all cities, the majority of hosts were found to be racially White, even in those with a markedly diverse racial makeup (e.g., Hong Kong, Chicago, Nashville). We shall comment extensively on this finding in the next section.

Table 4 Quintiles of Airbnb's host age distributions for full property rentals. Q_1 denotes the bottom 20% of the age distribution, Q_2 denotes users falling between the bottom 20% and 40% of the age distribution, and so on

City	Hosts				
	Q_1	Q_2	Q_3	Q_4	Q_5
Amsterdam	<26	[26, 31)	[31, 34)	[34, 39)	>39
Chicago	<26	[26, 31)	[31, 35)	[35, 39)	>39
Dublin	<24	[24, 29)	[29, 33)	[33, 38)	>38
Hong Kong	<27	[27, 31)	[31, 35)	[35, 39)	>39
Nashville	<27	[27, 31)	[31, 34)	[34, 39)	>39

The last demographic feature we consider is *age*. Table 4 reports the breakdown of each considered city's host population based on quintiles (we found the summary statistics to be very similar in the case of hosts, as well as when separating hosts based on full or shared property rentals). As shown in Table 4, Airbnb is a rather 'young' community, with the vast majority of users found to be between their mid-twenties and late thirties.

5.2 Rewiring analysis

We now present results of Airbnb host/guest pairing dynamics and compare them with the results of the rewiring analytical method illustrated above. We break down results by focusing on one demographic characteristic at a time for both hosts and guests. In this work, we did not investigate interactions between different features (e.g., between age of a host and gender of a guest); the same methodology could be used in the future to study these interactions.

Gender-related pairings. As shown in Table 5, we found very different, strongly city-dependent patterns. Indeed, we found same gender interactions to be prevalent in some cities (e.g., Amsterdam) while in others we found interactions between hosts and guests of different genders to be prevalent (e.g, Nashville). Yet, these results *per se* are not necessarily informative, as they partially echo each city's Airbnb population composition, and become statistically relevant only when compared with the *expected* rate of interaction measured under the null hypothesis of random interaction encoded in the above link rewiring procedure. As reported in Table 5, depending on the city and property type we detect very different over/under-expression patterns. Namely, when looking at full property rentals, we find same-gender interactions to be over-expressed (homophily), and interactions between different genders to be under-expressed (avoidance), in all cities. Things change considerably when looking at shared properties, where in the case of Dublin we find interactions between different genders to be over-expressed (heterophily) and, symmetrically, same gender interactions to be under-expressed, while in Nashville we find interactions to be compatible with the null hypothesis.

Race-related pairings. Results capturing pairing dynamics in terms of race are shown in Tables 6 and 7 for full and shared properties, respectively, once again broken down per city and per rental type. As a general observation, we can note that the signals we detect are not as strong as the ones detected for gender, with most over/under-expressions being of the order of a few fraction percentage points. Yet, strictly speaking, in most cases we detect an over-expression of White-to-White, Asian-to-Asian, and Black-to-Black guest–host pairings, regardless of whether the property was a shared rental or not (although with notable exceptions in Dublin and Nashville). Of note, the presence of racial homophily was

Table 5 Pairings between guest–host genders (F = female; M = male) in Airbnb. Values in brackets represent 95% confidence level intervals obtained from the rewiring analysis, while the values to their right denote the corresponding empirically observed frequencies. Upward green (downward red) arrows highlight over-expressed (under-expressed) values

City	FF	FM	FM	FM	MF	MF	MM	MM
Full property rental								
AMS	[29.39; 29.44]%	30.23% ↑	[23.98; 24.03]%	23.19% ↓	[25.61; 25.65]%	24.82% ↓	[20.93; 20.97]%	21.76% ↑
CHI	[33.57; 33.66]%	34.81% ↑	[25.04; 25.13]%	23.89% ↓	[23.60; 23.69]%	22.45% ↓	[17.61; 17.70]%	18.84% ↑
DUB	[35.17; 35.40]%	37.80% ↑	[25.29; 25.52]%	22.88% ↓	[22.85; 23.08]%	20.44% ↓	[16.23; 16.46]%	18.87% ↑
HK	[33.24; 33.33]%	33.79% ↑	[24.51; 24.61]%	24.06% ↓	[24.16; 24.27]%	23.72% ↓	[17.89; 17.99]%	18.44% ↑
NAS	[41.29; 41.49]%	41.78% ↑	[22.60; 22.80]%	22.32% ↓	[23.14; 23.34]%	22.86% ↓	[12.56; 12.76]%	13.05% ↑
Shared property rental								
AMS	[28.55; 28.62]%	29.27% ↑	[20.65; 20.72]%	20.00% ↓	[29.38; 29.44]%	28.72% ↓	[21.28; 21.34]%	22.01% ↑
CHI	[31.50; 31.58]%	32.91% ↑	[24.64; 24.72]%	23.32% ↓	[24.53; 24.61]%	23.21% ↓	[19.16; 19.24]%	20.57% ↑
DUB	[30.88; 31.47]%	29.59% ↓	[26.59; 26.85]%	28.14% ↑	[22.68; 22.94]%	24.24% ↑	[19.32; 19.59]%	18.03% ↓
HK	[33.25; 33.37]%	33.89% ↑	[23.96; 24.07]%	23.44% ↓	[24.79; 24.91]%	24.28% ↓	[17.77; 17.88]%	18.40% ↑
NAS	[42.33; 42.60]%	42.55% ↑	[30.91; 31.18]%	30.96%	[15.27; 15.54]%	15.32%	[10.95; 11.22]%	11.17%

Table 6 Pairings between racial backgrounds of Airbnb guests and hosts (W = White, A = Asian, B = Black) in full property rentals. Values in brackets represent 95% confidence level intervals obtained from the rewiring analysis, while values below them denote the corresponding empirically observed frequencies. Upward green (downward red) arrows highlight over-expressed (under-expressed) values

AMS	Full property rental		
	W	A	B
W	[76.73; 76.75]% 76.81% ↑	[7.62; 7.64]% 7.59% ↓	[1.14; 1.15]% 1.14%
A	[4.07; 4.08]% 4.07% ↓	[0.40; 0.41]% 0.43% ↑	[0.06; 0.06]% 0.06%
B	[1.88; 1.89]% 1.87% ↓	[0.18; 0.19]% 0.20% ↑	[0.03; 0.03]% 0.03% ↓
CHI	W	A	B
W	[71.66; 71.71]% 72.07% ↑	[9.41; 9.45]% 9.32% ↓	[2.06; 2.08]% 2.01% ↓
A	[6.73; 6.77]% 6.64% ↓	[0.88; 0.91]% 0.98% ↑	[0.02; 0.02]% 0.02%
B	[1.71; 1.73]% 1.50% ↓	[0.22; 0.23]% 0.24% ↑	[0.04; 0.05]% 0.15% ↑
DUB	W	A	B
W	[82.34; 82.42]% 82.28% ↓	[6.98; 7.05]% 7.10% ↑	[1.52; 1.55]% 1.43% ↓
A	[3.32; 3.38]% 3.73% ↑	[0.28; 0.33]% 0.07% ↓	[0.09; 0.12]% 0.00% ↓
B	[0.87; 0.90]% 0.57% ↓	[0.10; 0.13]% 0.22% ↑	[0.07; 0.09]% 0.22% ↑
HK	W	A	B
W	[35.72; 35.82]% 37.41% ↑	[22.39; 22.49]% 20.74% ↓	[0.74; 0.76]% 0.72% ↓
A	[20.05; 20.15]% 18.32% ↓	[12.47; 12.57]% 14.34% ↑	[0.40; 0.42]% 0.38% ↓
B	[1.03; 1.06]% 1.12% ↑	[0.62; 0.65]% 0.57% ↓	[0.03; 0.03]% 0.05% ↑
NAS	W	A	B
W	[78.95; 79.04]% 78.92% ↓	[6.66; 6.73]% 6.52% ↓	[2.14; 2.18]% 2.21% ↑
A	[2.83; 2.89]% 2.96% ↑	[0.22; 0.26]% 0.22%	[0.08; 0.10]% 0.00% ↓
B	[3.36; 3.42]% 3.40%	[0.24; 0.29]% 0.38% ↑	[0.10; 0.12]% 0.11%

found to be strongest in Hong Kong, both in terms of White-to-White and Asian-to-Asian guest–host pairings.

If we now change focus to pairing dynamics between different racial groups, we found results to be almost symmetric: that is, the majority of pairings between hosts and guests from different racial groups were under-expressed, although with a few exceptions (see, e.g., Dublin’s full property case), and these results appear to be largely independent from the property rental type.

We verified the robustness of the above findings against possible confounding factors associated with the users’ economic status. Namely, we checked whether the homophily

Table 7 Pairings between racial backgrounds of Airbnb guests and hosts (W = White, A = Asian, B = Black) in shared property rentals. Values in brackets represent 95% confidence level intervals obtained from the rewiring analysis, while values below them denote the corresponding empirically observed frequencies. Upward green (downward red) arrows highlight over-expressed (under-expressed) values

AMS	Shared property rental		
	W	A	B
W	[72.34; 72.37]% 72.39% ↑	[10.37; 10.40]% 10.31% ↓	[1.24; 1.25]% 1.24%
A	[8.30; 8.33]% 8.27% ↓	[1.18; 1.20]% 1.24% ↑	[0.14; 0.15]% 0.13% ↓
B	[1.46; 1.47]% 1.44% ↓	[0.20; 0.21]% 0.24% ↑	[0.02; 0.03]% 0.03%
CHI	W	A	B
W	[66.48; 66.53]% 66.92% ↑	[12.87; 12.92]% 12.57% ↓	[2.11; 2.13]% 2.03% ↓
A	[7.73; 7.74]% 7.32% ↓	[1.49; 1.53]% 1.89% ↑	[0.24; 0.26]% 0.28% ↑
B	[2.04; 2.07]% 1.96% ↓	[0.40; 0.42]% 0.39% ↓	[0.06; 0.07]% 0.16% ↑
DUB	W	A	B
W	[80.20; 80.29]% 80.20%	[6.60; 6.67]% 6.69% ↑	[0.79; 0.82]% 0.75% ↓
A	[5.11; 5.18]% 5.19% ↑	[0.39; 0.44]% 0.27% ↓	[0.08; 0.10]% 0.16% ↑
B	[1.07; 1.11]% 1.07%	[0.11; 0.13]% 0.11%	[0.05; 0.05]% 0.00% ↓
HK	W	A	B
W	[29.04; 29.16]% 30.72% ↑	[25.74; 25.86]% 23.90% ↓	[0.81; 0.83]% 0.84% ↑
A	[19.64; 19.76]% 18.04% ↓	[17.52; 17.64]% 19.51% ↑	[0.54; 0.57]% 0.52% ↓
B	[0.90; 0.92]% 0.81% ↓	[0.77; 0.80]% 0.87% ↑	[0.03; 0.04]% 0.06% ↑
NAS	W	A	B
W	[76.54; 76.69]% 77.13% ↑	[6.26; 6.37]% 6.60% ↑	[2.05; 2.11]% 1.49% ↓
A	[5.11; 5.22]% 4.79% ↓	[0.39; 0.47]% 0.21% ↓	[0.17; 0.22]% 0.43% ↑
B	[1.88; 1.95]% 1.70% ↓	[0.16; 0.21]% 0.21%	[0.11; 0.15]% 0.32% ↑

and avoidance patterns shown in Tables 6 and 7 might arise as a byproduct of correlations between racial background and wealth/income. We used the number of Airbnb properties owned and the price charged for a week-long stay as proxies for a host's income. We then performed two types of analysis: (i) a matched pair analysis [33], to measure the rate of interaction of White guests across groups of White and non-White hosts with similar levels of income; and (ii) a rewiring analysis on the sub-networks obtained after removing all hosts belonging to the top and bottom thirds of the income distribution. In both cases, by focusing our attention on hosts of similar economic status, we try to control for hosts' wealth. The results obtained from the matched pair analysis were statistically signif-

icant only in Hong Kong and Chicago (full properties); in these settings, they do confirm those reported in Table 6; results obtained from the rewiring analysis were all significant and fully in line with the ones reported in Tables 6 and 7. Full results are reported in the Additional file 1.

Age-related pairings. When comparing the empirical frequency of interaction between the age groups reported in Table 4, we again detect a number of slight deviations from the null model, both in terms of over and under-expressions. Yet, such relationships do not follow any discernible pattern neither in terms of city nor in terms of property type. We therefore chose not to report them here.

All in all, our findings show that, once online, interactions between sharing economy users display a tendency towards homophily and avoidance. The results we obtained should not be interpreted as weak statistical signals. Indeed, the null networks we employed in our analysis are constrained to preserve the original networks' degree sequences intact, which deliberately limits the extent to which the original networks can be randomised, and leads to the proximity we observe between the empirically measured interaction rates between genders and races and their counterparts in the null network ensembles. Yet, such a tendency is not universal, and our findings reveal patterns that vary across cities.

6 Discussion

In this section, we discuss the results presented above, trying to offer an interpretation based on their societal context, and proposing recommendations concerning the design of sharing economy platforms. We break down the discussion into two parts: we first consider the Airbnb's user base in each city under study, and reflect upon it relative to the city's demographic, economic, and historical context; we then move our discussion to our findings on pairing dynamics.

6.1 City demographics vs. Airbnb community

Echoing the findings from other studies of the sharing economy, our investigation into the user base of Airbnb revealed a disparity between its communities and the city-level demographics surrounding them, both in terms of age, gender and race. As far as *age* is concerned, we found Airbnb hosts and guests to be overwhelmingly young (mid-twenties to mid-thirties). This can be interpreted as a reflection of the broader age-related digital divide phenomenon [34].

In terms of *gender*, we have found the Airbnb community to be predominantly female. In 2015, Airbnb reported that 54% [35] of their guests were female. Based on the data we have collected up to 2017 for the cities under study, such percentage seems to be substantially higher (60% and above), both for hosts and for guests, both for private and for shared rental properties. We do not know whether this is a signal of an evolutionary mechanism, whereby more female than male users join the platform attracted by homophily, as they already see more female users already being engaged with it.

Perhaps the most notable results were found in terms of *race*. In the following, we focus our attention only on hosts, so to compare our results with available census information on the demographics of the cities under study. In all cities, the majority of hosts were found to be racially White. In particular, Dublin was found to have the highest proportion of White hosts (at 96% for full property rental); this is expected for a city whose resident population was reported to be 90% racially White in the latest Census⁸ (2011).

Things are considerably different in the more diverse cities we analyzed, where we found well-known inequalities along racial lines to be largely replicated in Airbnb's interactions, with systematic evidence of a divide between the Airbnb community and the local demographics. Indeed, we found up to 64% of Airbnb hosts in Hong Kong to be White (for full property rentals), even when 92% of the population is reported to be of Chinese (i.e., Asian) racial background as per the 2016 Census.^h Hong Kong is well-known to be plagued by rising wealth inequality [36] and exorbitant property prices, that are known to be among the most unaffordable in major cities [37]. This is coupled with high income inequality, with Hong Kong's racially White population amassing relatively high household income, while the same does not hold for the majority of the local Chinese population [38]. Owning a spare property to rent out might thus be a privilege mostly in the hands of the White population.

Similarly of interest is Amsterdam's large majority of racially White hosts, resting at roughly 90% of the total Airbnb user pool. This too is substantially higher than expected, given that a third of Amsterdam's population is composed of migrants recognized to be of non-western racial origins.ⁱ This finding is in line with research conducted by the Netherlands' Central Commission for Statistics (CCS), which has highlighted the existence of societal integration issues [39] amongst the Netherlands' non-western population, which we speculate to be reflected into a decreased ability to secure properties to rent out.

Chicago and Nashville were found to have the highest proportion of Black hosts and guests recorded among the cities under study, averaging around 3–4% across the two property rental types. However, this result too is notable, given that the most recent census reports Chicago's and Nashville's total populations to be 32%^j and 28%^k racially Black, respectively—an order of magnitude more than what found on Airbnb. This resonates with recent research suggesting that Airbnb is a conduit for racial gentrification where the old, local community members of a neighborhood lose out in housing and in wealth [4].

6.2 Homophily and avoidance

Our statistical investigation on pairing dynamics detected evidence of homophily both in terms of gender and race. Gender homophily is well documented to be 'built in' even in young children [40, 41], so it is not surprising we could detect it in our results too. Conversely, it was interesting to see it supplanted by heterophilous behavior (i.e., a statistical over-expression of interactions between hosts and guests of different genders) in the case of Nashville, regardless of the property type, and both in Dublin and Hong Kong when switching from full property to shared property rentals. This is even more interesting when considering that full properties obviously do not imply any shared space between hosts and guests, and often allow to avoid any live interaction through automated checkin procedures. We speculate that a possible explanation behind this might lie in the different communities that naturally self-select based on property type, with those selecting shared accommodation most likely being more open-minded and prone to meeting different people (see [42] for similar findings in couchsurfing platforms).

In addition, we also detected less strong statistical signals of homophily when analysing pairing dynamics based on racial background. Once again, this is somewhat to be expected, as racial homophily has been detected in a broad variety of social environments [7], ranging from labour markets [43] to online social networks [19]. Symmetrical to this,

we also detected phenomena of racial avoidance (still with quite weak statistical signals in most cases), i.e., under-expressions of relationships between guests and hosts belonging to different racial groups. This, again, resonates with pre-existing literature. For example, racial avoidance has been found to partially explain relocation patterns within countries [44]. These results were accompanied by a few exceptions where we detected under-expressed homophily (White-to-White relationships in Dublin's and Nashville's full properties) and heterophily (e.g., Asian-to-White relationships in Dublin).

Since our results are of a purely statistical nature, we can only highlight what relationships are over/under-represented, without making any claims on causality. In particular, we are in no position to distinguish between avoidance and outright discrimination. Yet, some of the trends we observe are worrying and raise questions about potential countermeasures that platforms might adopt in order to monitor their progress and possibly control them. For example, following research on unconscious biases [45], platforms could design interventions aimed at providing users with detailed information about the peers they chose to interact with (or not) in the past, possibly highlighting systematic preferences or deviations from the outcomes that would be obtained under an unbiased selection process. Interventions could also go a step beyond raising awareness of individual behaviours. For example, they could encourage behaviours to enhance heterophily (which we already detected in some cities) by means of incentive systems similar to those that are already in place to promote service excellence (e.g., rewarding outstanding Airbnb hosts with a 'superhost' status). In this fashion, users with a history of interactions with peers from different racial backgrounds could be rewarded with badges or statuses highlighting their role as diversity champions. Last but not least, platforms could incentivise users to give up potentially unnecessary steps in the interaction process where additional, and potentially biasing, information about other peers is usually acquired; for example, Airbnb's 'instant booking' option, where a guest's request is automatically accepted by the platform, without an explicit consent action from the host, has been an exemplary step in this direction.

6.3 Limitations

We ought to acknowledge three main limitations of this work: first, in the generation of the null network models, we have not enforced the preservation of temporal constraints (i.e., it is possible for a stay that occurred between guest g_1 and host h_1 in year y_1 to be swapped with a stay between g_1 and h_2 , despite host h_2 joining Airbnb only in year $y_2 > y_1$). We chose to adopt this simplified approach under the assumption that Airbnb demographics composition has not changed significantly between 2008 and 2016 (e.g., women have consistently made up the majority of the Airbnb host community [46]). In the future, we will consider generating null network models that preserve timing constraints (see, e.g., [47]).

Second, our findings rely on the accuracy of several image processing tools, to automatically annotate profile pictures in terms of gender, age and race. If the accuracy of these annotations is low, then the findings are void. In this paper, we have tried to reduce this risk by cross validating annotations across several image processing tools, and by verifying the robustness of our findings with respect to variations in their accuracy. Even so, we had to disregard any user whose profile picture did not present a (recognisable) human face, or where the estimated confidence of the annotation was low. Platform owners are most likely in possession of more accurate demographic information, explicitly provided at the time of user registration; they could thus skip the image annotation step (Airbnb Data

section) and directly use this information to annotate nodes in the bipartite graph, then proceeding with the application of the statistical network analysis method we proposed (Method section) to extract more robust results.

Finally, we ought to acknowledge that our analysis of Airbnb pairing dynamics was limited to what the platform makes externally visible (i.e., reviews that hosts and guests leave to one another after a stay); results might differ if one had the opportunity to apply our method to the whole history of interactions (including stays that resulted in no reviews, and reservation requests that were cancelled/refused). Once again, platform owners do possess the whole interaction history, and might thus want to repeat this study so to validate our findings on a complete network of host/guest stays. Yet, we have reason to believe our results would hold regardless. Indeed, the large samples our analysis relies on are such that only major differences in the tendency to leave reviews between groups would affect the significance of the findings reported in this paper.

7 Conclusion

In this paper, we have gathered and analyzed data to assess the diversity of the Airbnb community, especially with respect to the cities where it is embedded, and we have presented a method based on the statistical analysis of networks to detect homophily, heterophily and avoidance between different groups in the Airbnb community. To the best of our knowledge, network rewiring techniques, and, more generally, null network ensembles have never been employed as a tool to detect bias, and this application represents an element of novelty of our work.

Our findings suggest that, in all cities under study, certain user groups (e.g., young, White, female) are substantially over-represented compared to the local population; furthermore, statistically significant signals of gender and racial homophily were detected, across all cities and regardless of the property rental type.

Taken together, our findings echo sentiment that perhaps, contrary to all branding, the sharing economy community might not be that diverse. Rather, platforms such as Airbnb might be acting as accelerators for gentrification processes that are already well underway in major cities. While policy and legislation interventions are needed to regulate who benefits from sharing economy platforms [12], technological considerations also deserve attention: for example, Airbnb, like most sharing economy platforms, requires hosts to have a traditional bank account at the time of registration, to which money will then be deposited when guests visit. This might hinder the ability for many hosts from socio-economic deprived backgrounds to join the platform in the first place (which might be reflected in the very low representation from certain racial backgrounds in our studies); alternative solutions, such as on-demand payment services provided by platforms like Bit-Pesa,¹ could lower the barrier to entry.

Additional material

[Additional file 1](#): Supplementary information

Funding

WL and GL were supported by an EPSRC Early Career Fellowship in Digital Economy (Grant No. EP/N006062/1).

Abbreviations

F, female; M, male; A, asian; B, black; W, white; Q, quintile; AMS, Amsterdam; DUB, Dublin; CHI, Chicago; HK, Hong Kong; NAS, Nashville.

Availability of data and materials

The Airbnb data analysed during the current study is available at <http://insideairbnb.com/get-the-data.html>. The additional face image datasets collected during the course of this study are not publicly available since they contain personally identifiable information (profile pictures) and the Ethics committee that approved this study prohibits public sharing. However, the complete list of URLs from where these images have been collected is available, and thus the data could be fully reconstructed.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

VK, WL, GL and LC designed and performed research and wrote the paper. VK and WL analyzed data. All authors read and approved the final manuscript.

Endnotes

- ^a <http://insideairbnb.com/get-the-data.html>
- ^b <https://github.com/cooyzee/face-react>
- ^c <https://indico.io/>
- ^d <https://www.betaface.com/wpa/>
- ^e <https://face-api.sightcorp.com/>
- ^f <https://www.faceplusplus.com/>
- ^g <http://www.cso.ie/en/census/census2011reports/census2011thisisirelandpart1/>
- ^h <http://www.bycensus2016.gov.hk/en/bc-mt.html>
- ⁱ <https://www.cbs.nl/nl-nl/achtergrond/2016/47/bevolking-naar-migratieachtergrond>
- ^j <https://www.infoplease.com/node/70107>
- ^k <https://suburbanstats.org/population/tennessee/how-many-people-live-in-nashville>
- ^l <http://www.bitpesa.co/>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 15 November 2018 Accepted: 20 March 2019 Published online: 29 March 2019

References

1. Zloteanu M, Harvey N, Tuckett D, Livan G (2018) Digital identity: the effect of trust and reputation information on user judgement in the sharing economy. *PLoS ONE* 13(12):0209071
2. Wachsmuth D, Weisler A (2018) Airbnb and the rent gap: gentrification through the sharing economy. *Environ Plan A: Econ Space* 50(6):1147–1170
3. Lee D (2016) How Airbnb short-term rentals exacerbate Los Angeles's affordable housing crisis: analysis and policy recommendations. *Harv L & Pol'y Rev* 10:229
4. Griffith MW (2017) Airbnb as a racial gentrification tool? *Brookly Deep*
5. Ge Y, Knittel CR, MacKenzie D, Zoepf S (2016) Racial and gender discrimination in transportation network companies. Technical report, National Bureau of Economic Research
6. Edelman B, Luca M, Svirsky D (2017) Racial discrimination in the sharing economy: evidence from a field experiment. *Am Econ J Appl Econ* 9(2):1–22
7. McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: homophily in social networks. *Annu Rev Sociol* 27(1):415–444
8. Schelling TC (1969) Models of segregation. *Am Econ Rev* 59(2):488–493
9. Schelling TC (1971) Dynamic models of segregation. *J Math Sociol* 1(2):143–186
10. Neal JW (2010) Hanging out: features of urban childrens peer social networks. *J Soc Pers Relatsh* 27(7):982–1000
11. Aghion P, Akcigit U, Bergeaud A, Blundell R, Hémous D (2015) Innovation and top income inequality. *Rev Econ Stud*
12. Quattrone G, Proserpio D, Quercia D, Capra L, Musolesi M (2016) Who benefits from the sharing economy of Airbnb? In: Proceedings of the 25th international conference on World Wide Web, pp 1385–1394. International World Wide Web Conferences Steering Committee
13. Thebault-Spieker J, Terveen LG, Hecht B (2015) Avoiding the South side and the suburbs: the geography of mobile crowdsourcing markets. In: Proceedings of the 18th ACM conference on computer supported cooperative work & social computing. ACM, New York, pp 265–275
14. Thebault-Spieker J, Terveen L, Hecht B (2017) Towards a geographic understanding of the sharing economy: systemic biases in uberx and taskrabbit. *ACM ToCHI*
15. Tynes BM, Garcia EL, Giang MT, Coleman NE (2011) The racial landscape of social networking sites: forging identity, community, and civic engagement. *ISJLP* 7:71
16. Daniels J (2009) *Cyber Racism: White Supremacy Online and the New Attack on Civil Rights*. Rowman & Littlefield Publishers

17. Bakshy E, Messing S, Adamic LA (2015) Exposure to ideologically diverse news and opinion on Facebook. *Science* 348(6239):1130–1132
18. Thelwall M (2009) Homophily in myspace. *J Assoc Inf Sci Technol* 60(2):219–231
19. Wimmer A, Lewis K (2010) Beyond and below racial homophily: erg models of a friendship network documented on Facebook. *Am J Sociol* 116(2):583–642
20. Bamman D, Eisenstein J, Schnoebelen T (2014) Gender identity and lexical variation in social media. *J Sociolinguistics* 18(2):135–160
21. Guha S, Wicker SB (2015) Do birds of a feather watch each other?: homophily and social surveillance in location based social networks. In: Proceedings of the 18th ACM conference on computer supported cooperative work & social computing. ACM, New York, pp 1010–1020
22. Ugander J, Karrer B, Backstrom L, Marlow C (2011) The anatomy of the Facebook social graph. [arXiv:1111.4503](https://arxiv.org/abs/1111.4503)
23. Huang Y, Shen C, Williams D, Contractor N (2009) Virtually there: exploring proximity and homophily in a virtual world. In: International conference on computational science and engineering, 2009. CSE'09 vol 4. IEEE Press, New York, pp 354–359
24. Foucault Welles B, Van Devender A, Contractor N (2010) Is a friend a friend?: investigating the structure of friendship networks in virtual worlds. In: CHI'10 extended abstracts on human factors in computing systems. ACM, New York, pp 4027–4032
25. Hannák A, Wagner C, Garcia D, Mislove A, Strohmaier M, Wilson C (2017) Bias in online freelance marketplaces: evidence from taskrabbit and fiverr. In: CSCW, pp 1914–1933
26. Edelman BG, Luca M (2014) Digital discrimination: the case of airbnb.com
27. Kakar V, Franco J, Voelz J, Wu J (2016) Effects of host race information on airbnb listing prices in San Francisco
28. Fradkin A, Grewal E, Holtz D (2017) The determinants of online review informativeness: evidence from field experiments on Airbnb. Technical report, Working paper
29. Gretzel U, Yoo K-H (2008) Use and impact of online travel reviews, pp. 35–46
30. Hanhijärvi S, Garriga GC, Puolamäki K (2009) Randomization techniques for graphs. In: Proceedings of the 2009 SIAM international conference on data mining. SIAM, Philadelphia, pp 780–791
31. Livan G, Caccioli F, Aste T (2017) Excess reciprocity distorts reputation in online social networks. *Scientific Reports* 7
32. Bardoscia M, De Luca G, Livan G, Marsili M, Tessone CJ (2013) The social climbing game. *J Stat Phys* 151(3–4):440–457
33. Ho DE, Imai K, King G, Stuart EA et al (2011) Matchit: nonparametric preprocessing for parametric causal inference. *J Stat Softw* 42(8):1–28
34. White P, Selwyn N (2013) Moving on-line? An analysis of patterns of adult Internet use in the UK, 2002–2010. *Inf Commun Soc* 16(1):1–27
35. Airbnb (2015) Airbnb summer travel report. <http://blog.airbnb.com/wp-content/uploads/2015/09/Airbnb-Summer-Travel-Report-1.pdf>
36. Hu F, Yun M (2013) Hong Kong Poverty Line Shows Wealth Gap With One in Five Poor. Bloomberg
37. Noble J (2014) Economic inequality underpins Hong Kong's great political divide. *Financial Times*
38. Government of Honk Kong (2014) Honk kong poverty situation report 2014. Technical report, Government of Honk Kong. https://www.povertyrelief.gov.hk/pdf/poverty_report_2014_e.pdf
39. voor de Statistiek CB (2014) Integration process is uphill struggle. Technical report, Central Bureau voor de Statistiek. <https://www.cbs.nl/en-gb/news/2014/47/integration-process-is-uphill-struggle>
40. Shrum W, Cheek NH Jr, MacD S (1988) Friendship in school: gender and racial homophily. *Sociol Educ*: 227–239
41. Stehlé J, Charbonnier F, Picard T, Cattuto C, Barrat A (2013) Gender homophily from spatial behavior in a primary school: a sociometric study. *Soc Netw* 35(4):604–613
42. Klein M, Zhao J, Ni J, Johnson I, Hill BM, Zhu H (2017) Quality standards, service orientation, and power in airbnb and couchsurfing
43. Jacquemet N, Yannelis C (2012) Indiscriminate discrimination: a correspondence test for ethnic homophily in the Chicago labor market. *Labour Econ* 19(6):824–832
44. Rathelot R, Safi M (2014) Local ethnic composition and natives' and immigrants' geographic mobility in France, 1982–1999. *Am Sociol Rev* 79(1):43–64
45. Lai CK, Marini M, Lehr SA, Cerruti C, Shin J-EL, Joy-Gaba JA, Ho AK, Teachman BA, Wojcik SP, Koleva SP et al (2014) Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. *J Exp Psychol Gen* 143(4):1765
46. Airbnb (2017) Women hosts and Airbnb: building a global community. <https://tinyurl.com/yd5rhp2a>
47. Li W, Aste T, Caccioli F, Livan G (2018) Reciprocity and success in academic careers. [arXiv:1808.03781](https://arxiv.org/abs/1808.03781)

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
