Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona*

Miquel-Àngel Garcia-López[1 2], Jordi Jofre-Monseny[2], Rodrigo Martínez Mazza[2] and Mariona Segú[4]

June 2019

Abstract

In this paper, we assess the impact of the arrival and expansion of Airbnb on housing rents and prices in the city of Barcelona. Examining highly detailed data on rents and both transaction and posted prices, we use several econometric approaches that exploit the exact timing and geography of Airbnb activity in the city. These include i) panel fixed-effects models with neighborhood-specific time trends, ii) an instrumental variable shift-share approach in which tourist amenities predict where Airbnb listings will locate and Google searches predict when listings appear, and iii) event-study designs. For the average neighborhood in terms of Airbnb activity, our preferred results imply that rents have increased by 1.9%, while transaction (posted) prices have increased by 5.3% (3.7%). The estimated impact in neighborhoods with high Airbnb activity is substantial. For neighborhoods in the top decile of Airbnb activity distribution, rents are estimated to have increased by 7%, while increases in transaction (posted) prices are estimated at 19% (14%).

Keywords: Housing markets, short-term rentals, Airbnb

JEL Classification: R10, R20, R31, Z30

[1]Universitat Autònoma de Barcelona and IEB
[2]Universitat de Barcelona and IEB
[3]RITM, Université Paris Sud, Paris Saclay

*We are grateful to Gabrielle Fack, Apostolos Filippas, Laurent Gobillon, Hans Koster, Miren Lafourcade, Florian Oswald, Henry Overman and Elisabet Viladecans-Marsal for their valuable comments. We also thank participants of UEA Dusseldorf 2018, UEA Columbia 2018 and RUES Paris 2019.

1 Introduction

Tourism has grown enormously in recent decades. Between 1990 and 2017, the worldwide number of international tourist arrivals increased from about 400 million to 1300 million (WTO, 2018). This pattern is particularly true for urban tourism; in fact, the number of visitors to the 132 most popular world cities increased by 45% between 2009 and 2015.[1] Peer-to-peer platforms such as Airbnb have recently entered the market through partly accommodating the increased demand for tourism in cities. As a consequence, they have contributed to increasing the overlap between tourism and housing markets by allowing owners of residential properties to enter the hospitality sector.

Proponents of peer-to-peer platforms argue that short-term rentals provide residents with an additional source of income while allowing for tourism decentralization within cities. From an economic point of view, home-sharing platforms can be seen as an efficiency improvement in markets where goods are not fully used (Barron et al., 2018). One could argue that empty apartments during holiday periods are efficiency losses that can be reduced by means of short-term rentals. However, if home-sharing platforms are used by owners to permanently shift from long- to short-term rentals for tourists, the supply of units in the long-term market will be reduced, increasing housing prices and rents. Those that oppose peer-to-peer platforms also emphasize that short-term rental units in residential areas might constitute a negative externality for residents in terms of noise or uncivil behavior and that they cause gentrification and displacement of long-term residents.[2] Complaints about the gentrification effects and nuisances of short-term rentals have resulted in the implementation of local policies that limit the expansion of platforms such as Airbnb. Examples of such policy responses include the requirement to have a specific permit (Barcelona, Berlin, Paris, San Francisco, and Los Angeles), limiting the rental period (Amsterdam, New York, Paris, and San Francisco), paying a rental tax (Amsterdam and San Francisco), or outlawing short-term rentals in some conditions (Berlin and New York).

Despite all these local policy responses, we still have limited evidence on the effect of peer-to-peer platforms on housing markets. In this paper, we analyze the effects of the arrival and expansion of Airbnb in Barcelona. Barcelona constitutes an ideal city to study the effects of Airbnb on local housing markets for several reasons. First, Barcelona has experienced a tourism boom, with the number of passengers in the city's airport increasing from 20 to 47 million between 2000 and 2017. In fact, Barcelona is currently the 7th most visited destination in Europe (and the 17th worldwide). Second, Airbnb accounts for the lion's share of peer-to-peer platforms activity in the city, far ahead of its competitors[3]. Third, Airbnb penetration in Barcelona is high, with Barcelona being Airbnb's 6th top destination worldwide.[4]

Table 1 compares the size of the stock of Airbnb listings in Barcelona, New York, Los Angeles and Paris in 2015. Despite (substantial) legal uncertainties regarding the use of peer-to-peer platforms such as Airbnb in Barcelona, about 2.06% of all units are listed

[1] Mastercard Global Destination Cities Index.

[2] Similarly, the hotel industry opposes home-sharing platforms, as they are viewed as a threat to fair competition. Zervas et al. (2017) empirically studies the effect of Airbnb on hotels revenues.

[3] For Barcelona, the market share of Airbnb is around 70% according to the datahippo project (https://datahippo.org/es/) which collects data for several peer-to-peer platforms since 2017. We do not use these data as they do not cover the period we study.

[4] 'You'll never guess which city has the most Airbnb listings'. Forbes. J Bishop 2017.

on Airbnb. [5] This figure is higher than in New York (1.31%) and Los Angeles (0.86%), while it is very similar to Paris (2.09%). However, if we measure Airbnb listings relative to the number of rented units, the percentage for Barcelona rises to 6.84%, a figure that is significantly higher than in the other cities in the table.[6] This high penetration of Airbnb in Barcelona is likely to be explained by the (large) difference between the returns of short- and long-term rentals. At the bottom of Table 1, we provide estimates of the returns of Airbnb relative to those of long-term rentals. In 2015, the average long-term rental price in Barcelona was €11 per night (€735 per month), while the average Airbnb price (short-term rental) was €71 per night. Hence, a unit on Airbnb for 10 days a month would yield the same monthly income as a long-term rental.

Table 1: Airbnb activity in 2015 in selected cities

	Barcelona	New York	Los Angeles	Paris
Airbnb Listings	16,951	45,260	30,000	35,000
as % of total units	2.06	1.31	0.86	2.56
as % of rented units	6.84	1.92	1.56	4.97
Average Airbnb price/day (€)	71	131	114	81
Long-term rent/day (€)	11	59	75	37
Days/month for same revenue	10	14	20	14

Notes: Data on Barcelona come from Cadastral Records and INCASOL, data on New York and Los Angeles come from US Census Bureau, Zillow Rent Index and *airdna*, and data for Paris come from INSEE and OLAP. All Airbnb data have been obtained through InsideAirbnb.

To study the effect of Airbnb listings on housing rents and prices, we combine publicly available web-scraped data on Airbnb listings in Barcelona with high-quality data on housing markets. Specifically, we have access to i) individual-level data on the universe of transactions of second-hand apartments sold in the 2009-2016 period and ii) all posted ads for rentals and sales from a major real estate website (Idealista) that were active each December in the 2007-2017 period. We aggregate the information at the geographical level of small neighborhoods, which leaves us with a panel dataset of 220 small geographical areas that have an average population of about 7,000 inhabitants.

Throughout the empirical analysis, our dependent variable is the average residual resulting from a hedonic regression in which log rents or housing prices are regressed on time dummies and unit characteristics. In all regressions, we control for neighborhood and time fixed effects. Since Airbnb has grown the most in central parts of the city, our main identification concern is that neighborhoods that experienced higher Airbnb penetration might also be simultaneously experiencing processes of gentrification.[7] We adopt several strategies to control for the potential confounding effects of gentrification.

[5] Barcelona's regulation of short-term rental platforms has not substantially changed during recent years. A city law passed in 2007 (Housing Rights Act 18/2007) states that tourist apartments that are neither primary nor secondary residences are required to have a business activity permit. Hence, when Airbnb first arrived in Barcelona around 2009, the short-term rentals of entire apartments without a permit were illegal. Nevertheless, enforcement of the law was very low until 2016, when the number of inspections substantially increased.

[6] Compared to traditional tourist accommodation, the number of active listings during 2017 was equal to 32% of the total number of beds in hotels in the city in that same year.

[7] For the US, processes of urban revival have been described and studied by Baum-Snow and Hartley (2016); Couture and Handbury (2016), while Behrens et al. (2018) focus on the changes in local businesses associated with gentrification processes. González-Pampillón et al. (2019) provide some evidence of gentrification in the city center of Barcelona.

First, we control for the time-varying neighborhood demographic characteristics that are associated with gentrification processes. Second, we allow for each small neighborhood to have a specific linear time trend. Third, we apply an IV strategy, where the instrument is the interaction between i) a measure of proximity to the city's tourist amenities at the neighborhood level and ii) a Google Trends measure that tracks changes in Airbnb activity over time. The proximity to tourist amenities predicts where Airbnb will locate, while Airbnb searches in Google predict when listings appear. We indirectly verify the exclusion restriction by showing that proximity to tourist amenities does not predict rent and price growth in the pre-Airbnb period (i.e., before 2013). Finally, we also estimate event-study regressions. We dummify Airbnb activity (beyond the 90th percentile) and estimate interaction terms between year dummies and the Airbnb intensity indicator that identifies those areas with more Airbnb listings. This approach allows us to check if prior to the expansion of Airbnb, housing markets were evolving similarly in neighborhoods that after 2012 would experience high Airbnb penetration compared to others.

Our findings indicate that Airbnb has increased both rents and prices. For rents, our preferred specification result suggests that 54 more active listings in a small neighborhood (about the average level in 2016) increase rents by 1.9%, while transaction and posted prices increase by 5.3% and 3.67%, respectively. However, our estimates imply that local impacts can be substantial in the most tourist parts of the city. In particular, our results imply that 200 listings (the average number of listings in the top decile of the Airbnb activity distribution in 2016) increase rents by 7% and transaction and posted prices by 19% and 14%, respectively. We develop a stylized model of the housing market, where owners can decide to rent long term to residents or short term to tourists. Consistent with our findings, the model predicts that Airbnb will increase housing prices and rents, with the former effect being larger than the latter. In terms of mechanisms, the model predicts that Airbnb reduces the supply of residential housing units. In line with this prediction, our analysis shows that Airbnb listings reduce the number of resident households in the neighborhood.

Despite Airbnb being a recent phenomenon, there are already some research papers that estimate the effect of Airbnb on housing markets in the US. Barron et al. (2018) perform an analysis similar to ours. They look at the impact of Airbnb on rents and house prices for all cities in the US.[8] Their main strategy consists of using a 'shift-share' instrument, where the time variation comes from Google Trends of 'Airbnb' searches, while the cross-sectional variation is a neighborhood 'touristiness' index based on the location of restaurants. They find that a 1% increase in Airbnb listings increases rents by 0.018% and housing prices by 0.026%. Finally, Koster et al. (2018) study the effects of Airbnb bans implemented by several, but not all, governments in the Los Angeles area. Exploiting changes in prices at the administrative border, they find that banning Airbnb decreases prices by about 5%.

We contribute to the literature in several ways. First, as explained above, in Barcelona, the difference in returns between Airbnb and long-term rentals is large, resulting in high levels of Airbnb activity. Second, we have access to multiple high-quality datasets to track granular changes in both housing rents and prices. These individual-level data allow us to measure changes in rents and housing prices while controlling for composition changes

[8] An earlier contribution is Sheppard et al. (2016) that focuses on New York City. Their results suggest that doubling Airbnb in a 300-meter circle around a property translates to an increase in its value by 6% to 9%

in rented or sold units. Third, this is the first study to estimate the effects of Airbnb in the context of a large European city. This is relevant given the underlying differences between European and US cities. For instance, there might be less excess capacity in European cities, where guest houses or basement apartments (below a main house) are virtually nonexistent. For the case of Barcelona, our analysis below shows that only a small proportion of housing units active on Airbnb are primary residences. Despite these differences, the results that we find are remarkably close to those found in Barron et al. (2018). Compared to Koster et al. (2018), our paper focuses on a different channel through which Airbnb affects housing markets. The spatial RD design, which constitutes the main analysis in Koster et al. (2018), compares changes in prices across municipality borders following Airbnb bans. This neatly identifies the price increase of a property due to the possibility of using Airbnb. However, properties located across a border might be part of the same housing market, and thus, their spatial RD estimates do not capture changes in rents and prices that are caused by supply reductions. Our last contribution is precisely to provide direct evidence on the supply mechanism by showing that Airbnb reduces the number of households living in the neighborhood.

The paper is organized as follows. In Section 2, we develop a stylized model to understand the channels through which short-term rentals might affect the residential markets of neighborhoods. Section 3 describes the Airbnb, rents and housing prices data and describes the most relevant variables. A description of the empirical strategies we follow is provided in Section 4. The main results are presented and discussed in Section 5, while Section 6 contains the instrumental variables and event-study results. Finally, some concluding remarks are provided in Section 7.

2 Theoretical framework

In this section, we develop a theoretical framework to study the effects of short-term rentals to tourists on the market for long-term rentals. The model is a partial equilibrium model and focuses on the rental market of a central neighborhood n. In this neighborhood, there is a housing stock of size C that is owned by absentee landlords.

Each resident in the city consumes one unit of housing. Owners can rent their apartments through a long-term rental to a resident and obtain an annual market rent Q or, alternatively, rent short-term to tourists and obtain an annual rent of T, where $T > Q$. T reflects the (exogenous) tourists' willingness to pay to stay in the neighborhood. We assume that there is an individual-specific cost b_j to rent short-term to tourists, which reflects the legal uncertainties or the costs of running an Airbnb business. b_j is heterogeneous across owners since they can differ in their risk aversion towards legal uncertainties or their access to legal services, as well as their valuation of privacy. If $T - b_j > Q$, the owner rents short term to tourists, while if $T - b_j \leq Q$, the owner rents long term to a resident. This allows T and Q to differ in equilibrium, which is a salient feature of the data for the case of Barcelona. In equilibrium, there is a marginal owner who is indifferent between renting to residents or to tourists, $T - b_j^* = Q$, which implies that owners with $b_j < b_j^*$ rent short term, while those with $b_j \geq b_j^*$ rent long term.

The utility that a resident obtains in neighborhood n is $U_i^n = Y_i - Q - \alpha F_b(b_j^*) + a_i$, where Y_i is income and $\alpha F_b(b_j^*)$ is a term reflecting the negative externality that tourism can impose on residents due to noise or uncivil behavior.[9] Finally, a_i is an idiosyncratic term reflecting the preference of resident i to live in neighborhood n. If we normalize the utility level that resident i can obtain elsewhere in the city to Y_i, the willingness to pay to live in neighborhood n of the marginal resident is $Q(a_i^*) = -\alpha F_b(b_j^*) + a_i^*$, with everyone with $a_i > a_i^*$ living in n and everyone with $a_i \leq a_i^*$ living somewhere else.

In equilibrium, the long-term rental market in neighborhood n must clear, which implies $C(1 - F_b(b_j^*)) = 1 - F_a(a_i^*)$. Without loss of generality, we assume that b_j and a_i are $U \sim (0, 1)$ which simplifies the market clearing condition to $C(1 - b_j^*) = 1 - a_i^*$. Combining the market clearing condition, the willingness to pay of the marginal resident, $Q = a_i^* - \alpha b_j^*$, and the definition of the marginal owner, $T - b_j^* = Q$, we obtain the share of owners that rent short term to tourists:

$$b_j^* = \frac{C - 1 + T}{1 + C - \alpha} \tag{1}$$

Equation 1 indicates that the penetration of Airbnb is determined by T, which reflects tourists' willingness to pay to stay in the neighborhood. In Appendix A, we show that the number of Airbnb listings in a neighborhood is closely related to the average price of its listings, as equation 1 predicts. T reflects how attractive a neighborhood is to tourists. In Section 4.2, we also document that the proximity to relevant tourist attractions is a strong predictor of Airbnb activity at the neighborhood level. Hence, the predictions of the model in terms of where Airbnb activity should be higher are supported by the data.

When the neighborhood is of little interest to tourists, i.e., $T \leq 1 - C$, residents outbid tourists, and there are no short-term apartments in the neighborhood. For $T \geq 1 - C$, $b_j^* > 0$, and the equilibrium price of long-term rentals is obtained by inserting the market clearing condition in the residents' willingness to pay function:[10]

$$Q = 1 - C + (C - \alpha) b_j^* \tag{2}$$

Equation 2 indicates that more units in the short-term rental market affect long-term rents through two different mechanisms that work in opposite directions. First, one additional unit in the short-term market reduces the number of long-term residents, which mechanically increases a_i^*, as the market clearing condition reveals. To put it differently, reducing the supply of long-term units increases their price. Second, a marginal increase in b_j^* means higher negative externalities, which contribute to lower long-term rents.

We follow the approach of Barron et al. (2018) to relate rents and housing prices. The market is assumed to be in a steady state, and the price of a housing unit (P) is given by the present value of discounted cash flows to the landlord:

$$P = \sum_{t=1}^{\infty} \delta^t \left[(1 - b_j^*) Q + \int_0^{b_j^*} (T - b) db \right] = \frac{1}{1 - \delta} \left[Q + (T - Q) b_j^* - \frac{(b_j^*)^2}{2} \right] \tag{3}$$

[9] Tourism as a negative externality is in line with the local population's perception of tourism as a negative phenomena in Barcelona. This fact has been documented by the biannual opinion poll made by the local authorities since 2011, which surveys citizens' perceptions of Barcelona's largest problems since 2011. In this poll, tourism was mentioned, on average, as the fourth major problem of the city during the entire period, reaching the top of the ranking in 2017.

[10] If $2 - \alpha \geq T$, b_j^* becomes zero, as the negative externalities drive all residents out of the neighborhood.

Assuming δ as the discount factor, the cash flow in each period reflects the fact that $1 - b_j^*$ units are rented long-term at price Q, and b_j^* units are rented in the short-term market at rate T paying the cost b_j. Equation 3 indicates that Airbnb increases housing prices (P) more than rents (Q) as part of the stock available for rent obtains a return of $T - b_j$ that is higher than Q.

Equations 2 and 3 guide our empirical analysis consisting of regressing changes in housing rents or prices on changes in Airbnb activity at the neighborhood level. In terms of mechanisms, one immediate prediction of the model is that Airbnb displaces residents. Equations 1 to 3 suggest the main identification threat that we face. The effect of Airbnb activity on residential housing markets will be biased if neighborhoods where Airbnb penetration is high (because T is high) experience changes in the residents' willingness to pay to live there at the time when Airbnb expands in the city.

The model also illustrates that Airbnb has strong redistributive impacts. We defer this analysis to Appendix B. There, we show that owners benefit from Airbnb either because they obtain the short-term rental rate (which is higher than the long-term rate) or because Airbnb increases the long-term rental rates. In contrast, long-term residents are worse off since they face higher rents and experience the negative externalities of tourism. This uneven distribution of costs and benefits of Airbnb is likely to explain the opposition that Airbnb has met in cities where the residential and tourism markets overlap the most.

3 Data and variables

3.1 Neighborhood definition

Our geographical unit of analysis is the basic statistical area (BSA). BSAs are built and used by the Barcelona City Hall for statistical purposes. There are 233 BSAs with an average of 7,122 inhabitants. We believe that BSAs are the appropriate neighborhood definition, as they are designed to contain similar populations in terms of socioeconomic characteristics. On a more practical note, their size is sufficient to generate meaningful measures of housing rents and prices for neighborhoods over time.

3.2 Airbnb

To measure Airbnb activity, we use information extracted directly from the website of Airbnb. These data have been collected at different points in time by an independent Internet user who has made the data publicly available. The dataset is called InsideAirbnb, and for Barcelona, it contains 21 data points and covers the time period April 2015 to February 2018. For every listing, there is information on the host ID, geographical coordinates, room characteristics, date the host registered, and date of each review received. Even though Airbnb is not the only peer-to-peer platform active in the city, we consider that its listings are a good proxy for the short-term rental market. On the one hand, as explained above, its market share is by far the highest among its competitors. On the other hand, most short-term rentals are advertised through more than one platform simultaneously, implying that adding listings from a second platform would entail significant problems of double counting.

For our purposes, it is crucial to identify the period that a listing has been active, with its entry and exit dates, and any possible break period in between these two points.

The main complication we face is that the scraped information starts only in 2015, and hence, a listing count does not provide information for the period prior to April 2015. To build a measure of Airbnb listings for the period of analysis, we use a method similar to that proposed by Zervas et al. (2017). We consider that a listing is active in a given quarter if it has received at least one review during that quarter. According to Airbnb, 72% of guests leave a review, which supports proxying stays with platform reviews.[11]

The potential consequences of Airbnb might be very different if Airbnb is used to rent out excess capacity (home-sharing), or if units are rented through Airbnb as their primary use. We label as 'commercial' the listings that correspond to this second category. Listings in Airbnb are *entire apartments*, *private rooms* or *shared rooms*. To obtain an idea of the magnitude of home-sharing versus commercial use of Airbnb, we consider as commercial listings those that are i) multi-hosted properties (those properties whose host has more than one listing) and ii) single-hosted *entire apartments* that have a minimum of 5 reviews per quarter. This definition is clearly conservative, as many entire properties are rented as separate *private rooms*.[12] Despite this, more than 75% of all listings in every single year in our sample correspond to this (conservative) commercial category. Hence, although there is some genuine home-sharing, Airbnb in Barcelona is mostly a commercial activity.

3.3 Rents and Prices

We use two sources of data to obtain information on rents and prices at a fine spatial level. In particular, we have two measures for prices (transaction prices and posted prices) and one measure for rents (posted rents). For transaction prices, we use data from the Catalan Tax Authority from transaction tax records, which contain information on the price, exact location, date of transaction, size of the housing unit, year of construction, and a variable reflecting the quality of the dwelling. We have the universe of transactions that occurred in Barcelona during the period 2009-2016.[13] We label this dataset ITP (*Impuesto sobre Transmisiones Patrimoniales*) or transaction prices.

For posted rents and prices, we use information from the online real estate portal Idealista. With more than one million ads and an average of 17 million weekly views, Idealista is by far the most important Spanish real estate portal. Idealista has provided us with all ads that were active for the city of Barcelona in December of every year for the period 2007-2017.[14] The data include the exact location, the posted rent or price and the size of the unit, among other characteristics.[15]

Having two measures of prices is useful because both transaction and posted prices

[11] An alternative approach would be to use the entry date and assume that listings never exit, which is the preferred method in Barron et al. (2018). In the case of Barcelona, we consider that this approach is problematic. First, approximately 25% of all listings do not have any reviews at all. Second, the entry date indicates the time when the host registered. If the host has multiple listings (which is the case for the majority of listings in Barcelona), it is not possible to know the entry date of each listing.

[12] This practice increased after July 2016. Short-term rentals of *entire apartments* without a permit is illegal, but enforcement was very low before July 2016. It is less clear if renting a *private room* is also against the law, and in practice, enforcement with respect to *private rooms* has been low throughout the period we study.

[13] We keep only those sales transactions with a declared value of less than 10,000,000 euros.

[14] We have dropped the following data: for sales, we drop those sales ads with posted prices below 10,000 euros and those of less than 20 square meters, and for rental ads, we drop all ads with monthly rents below 100 euros or above 30,000 euros.

[15] Other characteristics that are available and that we use are number of floors, number of rooms, presence of air-conditioning, lift and boxroom, and whether it is a studio, penthouse, or duplex.

have advantages and disadvantages. Posted prices might differ from agreed upon prices since there is often some bargaining in the process. Official transaction prices should, in principle, measure prices more precisely. However, in practice, the transaction (ITP) data have two limitations. First, there might be a nonnegligible time lapse between the date at which parties agree on a price and the date when the ITP tax is paid. Second, there is some fraud in the ITP tax that consists of underreporting the ITP price and thus the tax base. For rents, we cannot compare posted to actual rents. However, Chapelle and Eymeoud (2018) show that bargaining is less of an issue for rents and that online posted prices are a good measure of actual rents.

3.4 Descriptive Statistics

In Figure 1, we plot the evolution over time of Airbnb activity, together with that of rents and prices. Airbnb experienced a very rapid increase from its first entrance in 2009 up to 2017, when the growth stopped because of City Hall's increased efforts to reduce tourist apartments operating without a license. In 2016, the average BSA had 54 listings, while *High Airbnb Areas* (those BSAs in the top decile) had an average of 200 active listings. In these areas of the city, approximately 5% of all housing units are listed on Airbnb. The (substantial) variation in Airbnb activity across neighborhoods is further explored in Figure 2, which shows the distribution of Airbnb listings across BSAs for the last quarter of 2016. Airbnb activity is higher around the city center and, to some extent, along portions of the beach line. Note that Airbnb activity is low in many parts of the city.

Figure 1: Airbnb listings, rents and prices: 2007-2017

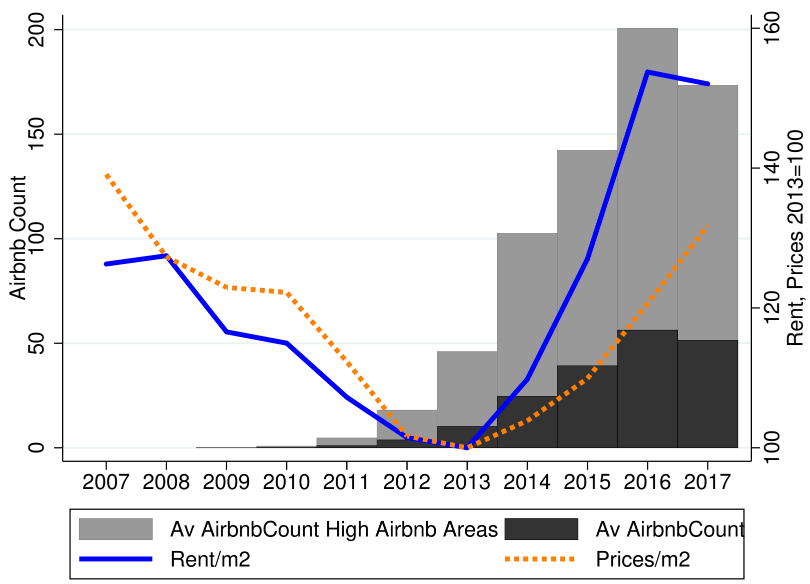

Notes: This graph plots the evolution over time of the BSA averages in Airbnb listings, rents and posted prices (per square meter) for the period 2007-2017. Rents and prices are normalized to their 2013 value. The dark gray bars represent the average Airbnb listings for all BSA, while the light gray bars depict the average listings for *High Airbnb Areas* (BSAs in the top decile of the Airbnb listings distribution in 2016).

As for the evolution of housing rents and prices, it can be observed that the period that we study is a turbulent one. Following the financial crises and the burst of the Spanish

Figure 2: Airbnb activity across BSAs in 2016

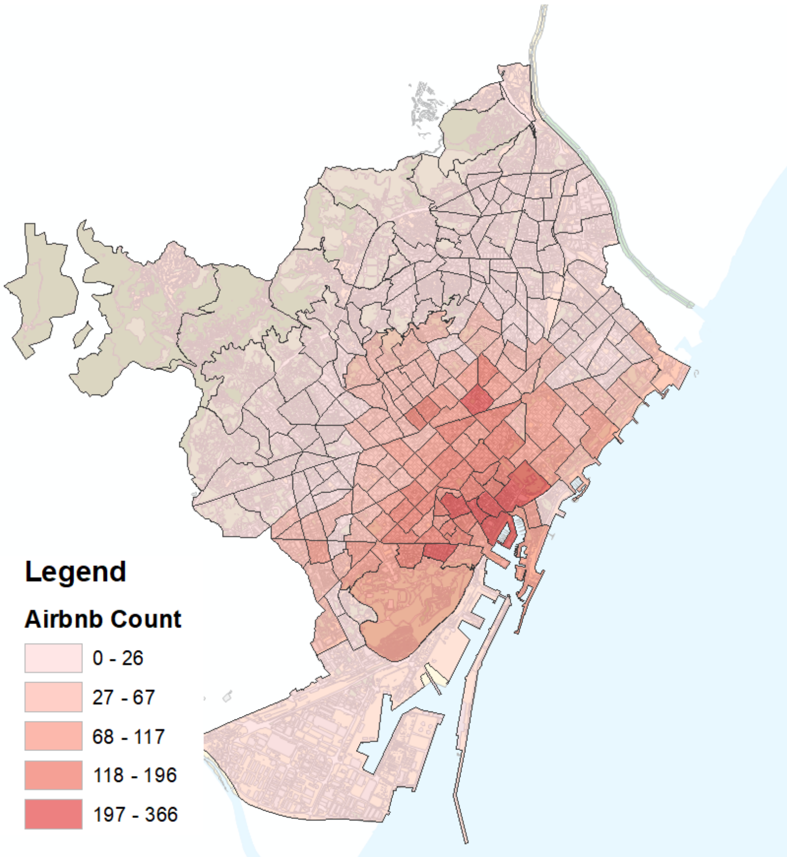

Notes: This graph plots the number of active Airbnb listings in the fourth quarter of 2016 at the BSA level.

housing bubble, rents and prices fell until 2013, when they started recovering. While prices reached the precrisis levels towards the end of the period, rents surpassed their precrisis levels around 2015 and kept rising. In fact, housing affordability is one of the main concerns among residents.[16] The recovery of rents and prices coincides in time with the expansion of Airbnb. In this paper, we attempt to determine how much of this overlap is causal.

For completeness, in Table 2, we report descriptive statistics of Airbnb listings, housing rents and prices, and sociodemographic characteristics that are used as controls in the regression analyses. More specifically, we report the BSA means for the years 2012 and 2016 for two different samples: all BSAs and *High Airbnb Areas* (BSAs in the top decile of the Airbnb listings distribution in 2016).

[16] See *'El acceso a la vivienda, el principal problema de los barceloneses'* Macedo, G., October 2010.

Table 2: Descriptive statistics: Variables' means across BSA for 2012 and 2016.

	2012		2016	
	All BSAs	High Airbnb Areas	All BSAs	High Airbnb Areas
Airbnb Count	2.79	13.35	47.84	178.58
Rent (€/m^2)	11.83	12.93	16.39	20.19
Posted Price (€/m^2)	3250	3338	3753	4282
Transaction Price (€/m^2)	2269	2356	2619	3027
Population	6978	7750	6973	7514
Population Density	0.03	0.04	0.03	0.04
Mean Age	43.36	42.10	43.69	42.08
% of Foreign Population	0.18	0.32	0.17	0.33
House Average Occupancy	2.47	2.41	2.48	2.41
Unemployment Rate	10.48	10.81	7.80	7.83
Income Index	98.37	96.48	102.78	104.58

Notes: Columns 1 and 3 report the mean for all BSAs in 2012 and 2016. Columns 2 and 4 report the means of *High Airbnb Areas* (BSAs in the top decile of the Airbnb activity distribution in 2016).

4 Empirical Strategies

4.1 Baseline Specification

Our main analysis consists of estimating the variants of the following fixed-effects specification:

$$log(Y_{n,t}) = \beta Airbnb\ count_{n,t} + \gamma X_{n,t} + \mu_n + \tau_t + \varepsilon_{n,t} \tag{4}$$

where $Y_{n,t}$ is our measure of housing rents or prices at the BSA level, $Airbnb\ count_{n,t}$ is the number of active listings in quarter t, τ_t are time fixed effects, and μ_n are BSA fixed effects that account for time-invariant neighborhood characteristics. Our dependent variable $log(Y_{n,t})$ is the average residual at the BSA-time period level of a (micro-level) regression in which log rents (or log housing prices) are regressed on time dummies and unit characteristics.[17] This controls for price changes across neighborhoods that could be explained by changes in the composition of units rented or sold across BSAs and over time. For example, it allows us to control for the fact that some BSAs might have a growing proportion of high-end apartments being sold or rented over time. Throughout the regression analyses, we weight BSA-year cells by the relevant number of ads or sales. Standard errors are clustered at the BSA level to account for serial correlation within panel units (Bertrand et al., 2004).

Our main concern regarding identification is that neighborhoods in which Airbnb activity grew the most during our period of study might be experiencing processes of sociodemographic change, which might have a direct impact on housing rents and prices. Specifically, Airbnb has grown the most in central parts of the city that have also been experiencing processes of urban revival in the last two decades. We adopt several strategies to control for the potential confounding effects of gentrification.

First, we introduce in equation 4 a set of time-varying demographic controls at the

[17] We construct a panel on the BSA-year (data from Idealista) and BSA-quarter (transaction prices)

BSA level ($X_{n,t}$), namely, average age, log of population density, average household occupancy rate, unemployment rate, relative income, and percentage of foreign residents. This allows us to control for yearly changes in variables associated with processes of gentrification. In some specifications, we allow for neighborhoods with different characteristics to have different linear time trends by introducing interactions between a linear time trend and the control variables measured in 2012, i.e., $X_{n,2012} \times t$.

Second, a more demanding approach in terms of data is presented in equation 5:

$$log(Y_{n,t}) = \beta Airbnb\ count_{n,t} + \gamma X_{n,t} + \mu_n + \tau_t + \rho_n \times t + \varepsilon_{n,t} \tag{5}$$

which includes the term $\rho_n \times t$ that fits a BSA-specific linear time trend by estimating interaction terms between a BSA-specific coefficient (ρ_n) and a time variable t. This is a very flexible specification since it allows for each BSA to have its own linear time trajectory in housing rents and prices. Here, the variations that we exploit are deviations from each BSA's own specific linear time trend. However, if Airbnb affects not only (rent or price) levels but also trends in these variables, including BSA-specific linear time trends, it will confound the effect of Airbnb with the BSA-specific time trend (Wolfers, 2006). Thus, we will resort to a procedure previously applied in the taxation (Kleven et al., 2014) and minimum wage (Monras, 2015) literature and estimate linear time trends using data prior to 2013 only (i.e., the pre-Airbnb period). Specifically, the method involves the estimation of the following two equations at a neighborhood-time level:

$$log(Y_{n,t}) = \mu_n + \tau_t + \rho_n \times t + \epsilon_{n,t}, \text{ for } t \leq 2012 \tag{6}$$

$$\widetilde{log(Y_{n,t})} = \beta Airbnb\ count_{n,t} + \gamma X_{n,t} + \tau_t + \varepsilon_{n,t}, \text{ for all } t \tag{7}$$

The first equation predicts the outcome based on BSA dummies, time dummies, and BSA-specific linear time trends. Based on these coefficient estimates, we predict $log(Y_{n,t})$ for the entire sample years and compute the residuals, $\widetilde{log(Y_{n,t})}$. In the second stage (equation 7), we regress these residuals against Airbnb listings, time dummies and the time-varying demographic controls ($X_{n,t}$).

4.2 Instrumental Variables Fixed-Effects Models

We also estimate equation 4 through two-stage least squares regression using as an instrument a shift-share variable that combines i) cross-sectional variation across BSAs in tourist amenities and ii) aggregate time variation in Airbnb activity.

Starting with the cross-sectional 'share' component of the instrument, we build an index that measures proximity to tourist amenities. To that end, we use TripAdvisor to produce a complete list of the city's tourist amenities.[18] We then proceed by geolocating the tourist amenities and by obtaining the number of reviews in Google to weight their relative importance.[19] Our measure of tourist amenities is built as follows:

[18] TripAdvisor is a website that offers tourism-related content. According to the site, it currently has over 390 million monthly unique visitors. We exclude the more endogenous and less historical amenities such as areas known for restaurants, bars or clubs.

[19] Although TripAdvisor also provides reviews, these are fewer in number than on Google.

$$TouristAmenities_n = \sum_k \frac{1}{dist_{n,k}} * Reviews_k \tag{8}$$

where k indicates the amenity, $dist_{n,k}$ is the distance in meters between the centroid of each BSA n and amenity k, and $Reviews_k$ is the number of reviews in Google. Figure 3 shows the location of these amenities, where the size of each circle is proportional to the number of reviews received.

Figure 3: **Location of tourist amenities**

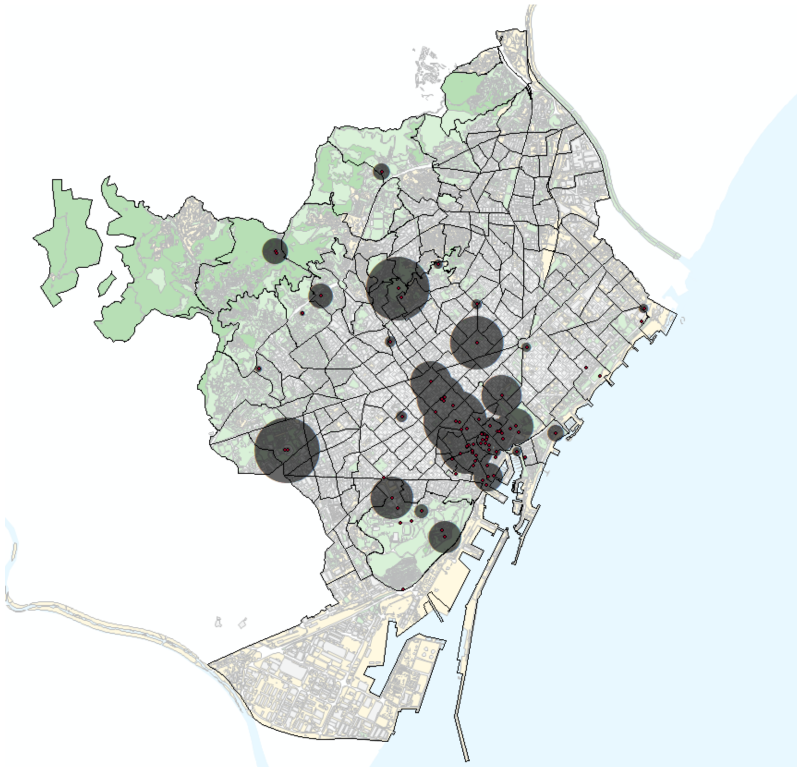

Notes: The size of the circles is proportional to the number of reviews received.

As can be see in Figure 3, tourist amenities are concentrated in the city center, although three of the most important ones have a less central location. These are the *Sagrada Familia* (easternmost large circle), *Parc Güell* (northernmost large circle) and *Camp Nou* (westernmost large circle). Being the 'share' component of the instrument, the tourist amenities index should predict where Airbnb listings will appear. Panel a) in Figure 4 plots this relationship by binning the data for deciles of the tourist index distribution. The graph clearly shows that BSAs that are closer to tourist amenities tend to show the highest number of Airbnb listings. This relationship can be rationalized by the model of Section 2. Tourists' willingness to pay (T in the model) is higher close to tourist amenities, where tourists outbid residents, which results in high levels of Airbnb activity. Consistent with this prediction, in Appendix A, we show that, indeed, Airbnb activity is higher in BSAs where the price of Airbnb is also higher.

Turning to the 'shift' component of the instrument, we follow Barron et al. (2018) and use worldwide searches in Google for the term 'Airbnb Barcelona'. This variable is measured at a monthly level and is normalized to 100 for the month with the highest number of searches. Panel b) of Figure 4 shows that the number of Google Trends searches for 'Airbnb Barcelona' tracks the time variation in Airbnb activity very well.

Figure 4: Airbnb activity, tourist amenities and Google Trends searches.

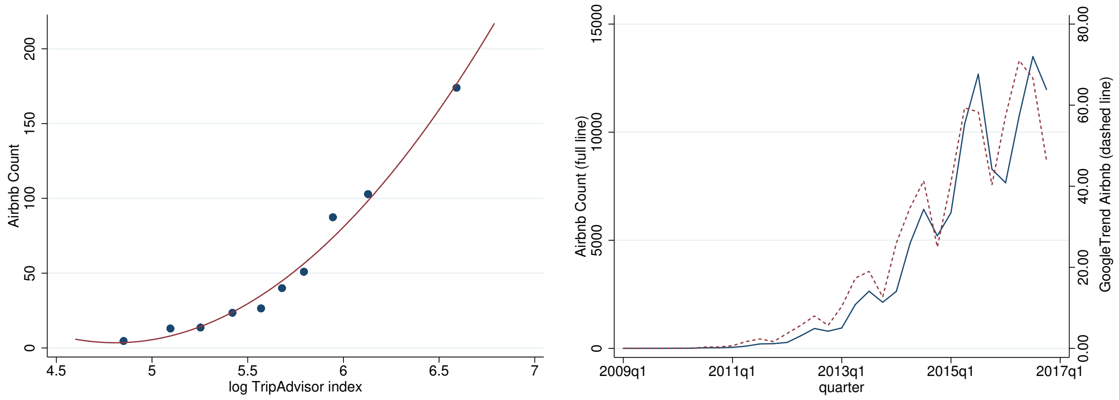

(a) Airbnb activity and tourist amenities **(b)** Airbnb activity and Google Trends searches

Notes: Graph (a) shows BSA Airbnb listings as a function of the tourist amenities index (bins are deciles of the tourist amenities distribution). Graph (b) shows the time profile of Airbnb listings (solid line, left axis) and Google Trends searches for 'Airbnb Barcelona' (dashed line, right axis).

The rationale behind the instrument works as follows. The proximity to tourist amenities predicts where Airbnb will locate, while searches in Google Trends for the term 'Airbnb Barcelona' predict when listings will appear. The relevance of the instrument is testable, and we return to it below. As for instrument validity, recent research on shift-share instruments indicates that the main identification threats are related to the 'share' component of the instrument (Goldsmith-Pinkham et al., 2018). Since our specifications contain a BSA fixed effect, instrument validity hinges on the assumption that the cross-sectional 'share' component, proximity to tourist amenities, is only correlated with changes in housing rents and prices through Airbnb listings. For example, our instrument would be invalid if residents' valuation of proximity to tourist amenities (or any other BSA characteristic that correlates with it) has changed over the period we study. If the instrument is valid, proximity to tourist amenities should not explain changes in housing rents and prices prior to the arrival of Airbnb. We address this issue at length below.

4.3 Event study plots

Finally, we also conduct event study exercises by focusing on the *High Airbnb Areas* (BSAs in the top decile of the Airbnb listings distribution in the last quarter of 2016).[20] Then, we perform regressions of the following type:

$$(9) \qquad log(Y_{n,t}) = \sum_{t \neq 2012} \delta_t \times 1\{HighAirbnbArea_{2016}\} + \gamma X_{n,t} + \mu_n + \tau_t + \varepsilon_{n,t}$$

where $HighAirbnbArea_{2016}$ indicates if the BSA belongs to the top decile in the 2016 Airbnb listings distribution. As in previous regressions, we include time and BSA fixed effects and time-varying demographic characteristics ($X_{n,t}$). We estimate $HighAirbnbArea_{2016}$

[20] In robustness tests, we alternatively use the top 5% and the top quartile.

× year interactions, leaving 2012 as the base year. Each of these coefficients quantifies the price difference between the *High Airbnb Areas* and the control group relative to the 2012 difference. Again, we choose 2012 as the last pre-Airbnb year, as starting in 2013, Airbnb's activity became more significant. This exercise allows us to check if, prior to the arrival of Airbnb, *High Airbnb Areas* showed parallel trends in housing rents and prices compared to other neighborhoods.

5 Main Results

5.1 Graphical evidence

Before proceeding to the regression results, in Figure 5, we show graphical evidence of the effect of Airbnb on housing markets. Specifically, we plot raw average (log) prices and rents series over time for *High Airbnb Areas* (i.e., those in the top decile of the distribution of Airbnb listings in 2016) versus the rest. In Panel (a) we graph rents, while in Panels (b) and (c), we show the corresponding graphs for transaction prices (ITP) and posted prices (Idealista), respectively. For completeness, in Panel (d), we plot Airbnb listings.

Figure 5: Evolution of rents and prices for *High Airbnb Area* vs. the rest

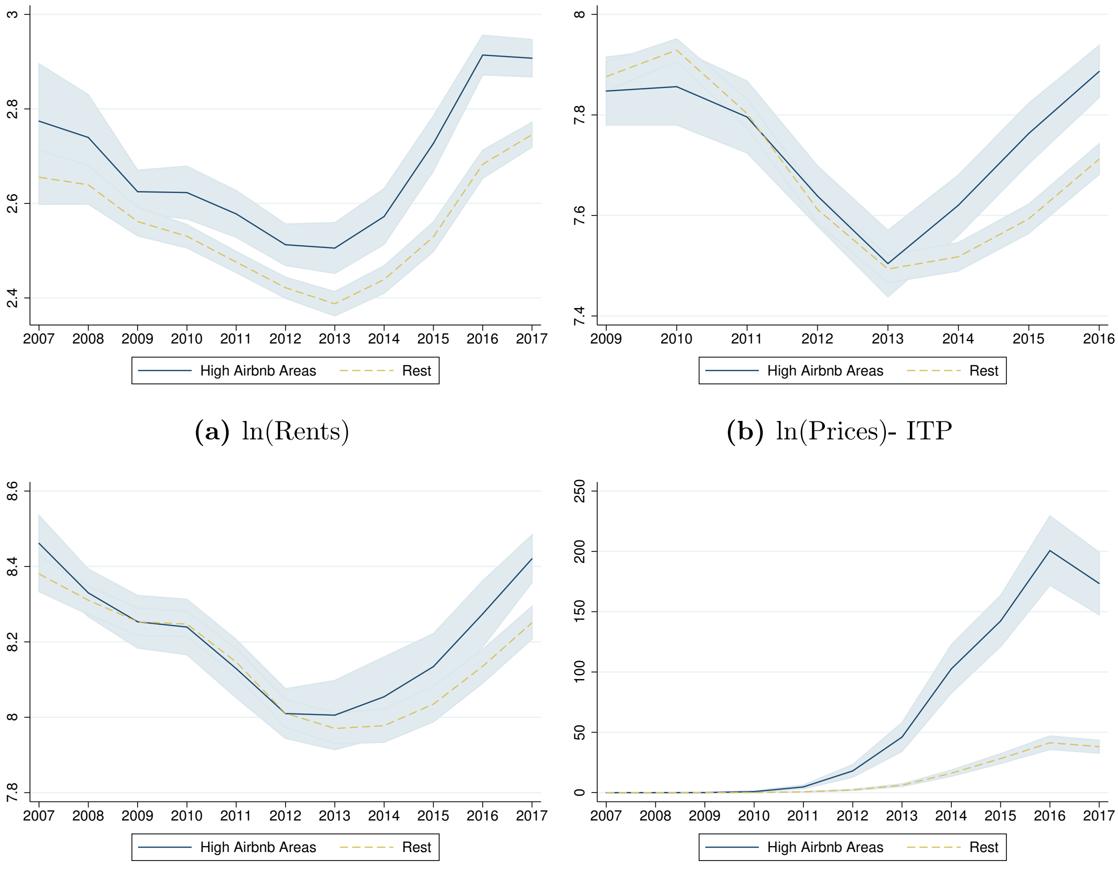

(a) ln(Rents)

(b) ln(Prices)- ITP

(c) ln(Prices)- Posted

(d) Airbnb Count

Notes: Graph plots raw averages and the appropriate confidence intervals. *High Airbnb Area* are those in the top decile of the Airbnb listings distribution in 2016.

The levels of both rents and prices tend to be higher in BSAs with more Airbnb activity. More interestingly for our analysis, while the series for the period before 2012 appear fairly parallel, the gaps in rents and prices seem to widen, coinciding with the expansion of Airbnb in 2013 and onwards, especially for rents and transaction prices, where the divergence is more noticeable. In the first three figures, the difference between the two groups is statistically significant at the end of the period, while this is not the case for the first years. Finally, in Panel (d), we report the evolution of the count of Airbnb listings by group. It can be seen that while the number of listings increased drastically for the *High Airbnb Areas*, the increase was very modest for the other BSAs, reflecting the fact that Airbnb is highly concentrated in particular areas of the city.

These graphs are suggestive evidence that neighborhoods experiencing higher Airbnb penetration experienced higher rents and prices growth with the arrival and expansion of Airbnb. Nevertheless, since these series might be affected by other confounding factors that could be biasing the results, we move to our main empirical strategies described in Section 4.

5.2 Results for baseline specifications

In Table 3, we report our baseline results for the impact of Airbnb on rents (Panel A) and prices (Panels B and C). As explained above, throughout the table, the dependent variable is the average BSA-time period residual of a micro regression in which log rents (or log prices) are regressed on housing characteristics and time dummies.

In column 1, we regress the outcome of interest against the number of Airbnb listings while controlling only for time and BSA-fixed effects. Then, in column 2, we add BSA time-varying demographic controls. Coefficients are positive and significant for both rents and prices, which implies that an increase in the number of listings translates into an increase in rents and prices. The effects on prices are larger than on rents, especially for transaction prices (ITP). It is interesting to observe that the presence of contemporaneous demographic controls has no large impact on the estimates for rents, while it slightly decreases coefficients for prices, although not in a statistically significant way. Nevertheless, we keep the demographic controls in subsequent specifications for the sake of completeness.

In column 3, we include differential demographic linear time trends by introducing interactions between a linear time trend and the control variables measured in 2012, i.e., $X_{n,2012} \times t$. The coefficients for prices are somewhat reduced, while they remain fairly constant for rents. In column 4, we report the results of specifications that fit BSA-specific time trends as described in equation 5. These allow for both observable and unobservable characteristics to impact neighborhood trends. Doing so increases the coefficient for rents (though not significantly) and reduces the coefficients for prices, especially for posted prices where the coefficient becomes nonsignificant. However, as mentioned before, one caveat of this approach is that if Airbnb impacts rent and price trends rather than levels, the BSA fixed effects will absorb part of the Airbnb effect on the outcomes.

In column 5, we repeat the analysis after detrending the data following the procedure described in equations 6 and 7. In a first step, the pre-Airbnb data are used to estimate BSA-specific time trends, which are then used to detrend all data points. Here, the coefficient for rents slightly decreases and gets closer to the specifications reported in columns 1 and 2. For prices, they both increase with respect to columns 3 and 4, and their magnitude becomes more similar to each other.

Across the different specifications, the results indicate that higher Airbnb penetration leads to increases in both rents and prices, with the effects on prices being larger than for

Table 3: Impact of Airbnb on rents and prices - FE Model

	Panel A: Rents				
	(1)	(2)	(3)	(4)	(5)
Airbnb Count (x100)	0.036***	0.035***	0.041***	0.058***	0.034*
	(0.008)	(0.009)	(0.010)	(0.020)	(0.018)
	Panel B: Transaction Prices (ITP)				
	(1)	(2)	(3)	(4)	(5)
Airbnb Count (x100)	0.118***	0.097***	0.080***	0.065***	0.084***
	(0.021)	(0.019)	(0.021)	(0.024)	(0.026)
	Panel C: Posted Prices (Idealista)				
	(1)	(2)	(3)	(4)	(5)
Airbnb Count (x100)	0.081***	0.068***	0.045***	0.022	0.085***
	(0.010)	(0.009)	(0.011)	(0.019)	(0.021)
Time FE	X	X	X	X	X
BSA FE	X	X	X	X	X
Dem Controls	-	X	X	X	X
Time Trends	-	-	Dem	BSA	-
Detrended	-	-	-	-	X

Notes: Significance is indicated by * p<0.1, ** p<0.05, and *** p<0.01. Standard errors, in parentheses, are clustered at the BSA level. Each column represents a different specification. Panel A reports the results for rents, while Panels B and C report the corresponding estimates for transaction and posted prices. Outcomes are average residuals at the BSA-time period level, as explained in the main text. Regressions weighted with the total number of ads (for rents and posted prices) and of transactions (for prices). The analysis takes place at the BSA-year level for rents and posted prices and BSE-quarter for transaction prices. There are 2,123 observations in Panel A, 7,005 in Panel B and 2,229 in Panel C. Demographic controls are average age, log of population density, average household occupation rate, unemployment rate, average relative income, and percentage of foreign residents.

rents. This result is in line with the theoretical prediction of equation 3. Since housing units that are on Airbnb yield, on average, a higher return than those units that are rented to residents, the price increase due to Airbnb exceeds that of long-term rentals.

Among the specifications that allow for linear time trends to vary across neighborhoods (columns 3 to 5), our preferred specification is that in column 5 for the reasons detailed in Section 4. The results of this specification turn out to be close to those reported in column 2, which corresponds to a more parsimonious specification with BSA and time fixed effects and time-varying control variables. Overall, we consider the estimates in column 2 as our baseline results for two reasons. First, the length of the time period before the expansion of Airbnb (i.e., <= 2012) might be too short to obtain robust estimates of BSA-specific time trends, i.e., the ρ_n's in equation 6. Second, the event-study exercises shown below indicate that the parallel trends assumption holds before 2013, suggesting that specifications that fit BSA-specific linear time trends are unnecessary.

We focus on the results in column 2 to gain insight into the economic size of the estimated effects. Taking these estimates at face value, our estimates imply that an increase in 100 Airbnb listings translates into a 3.5% increase in rents, 9.7% in transaction prices and 6.8% in posted prices. Given that the average increase in Airbnb activity in the period 2012-2016 is of 54 listings, our estimates imply an average increase in rents of 1.89% and increases in transaction and posted prices of about 5.24% and 3.67%, respectively.

The large degree of heterogeneity in Airbnb activity across BSAs implies that Airbnb has not affected all neighborhoods equally. In Figure 6, we illustrate these heterogeneous impacts by plotting the result of multiplying the coefficients obtained in column 2 by the Airbnb activity experienced by each BSA in 2016. While the implied effects are very close to zero for the less central BSAs, our estimates imply some local impacts that are substantial. For the *High Airbnb Areas*, Airbnb has increased rents, transaction prices and posted prices by an average of 7%, 20% and 14%, respectively.

Figure 6: Implied effects of Airbnb across BSAs

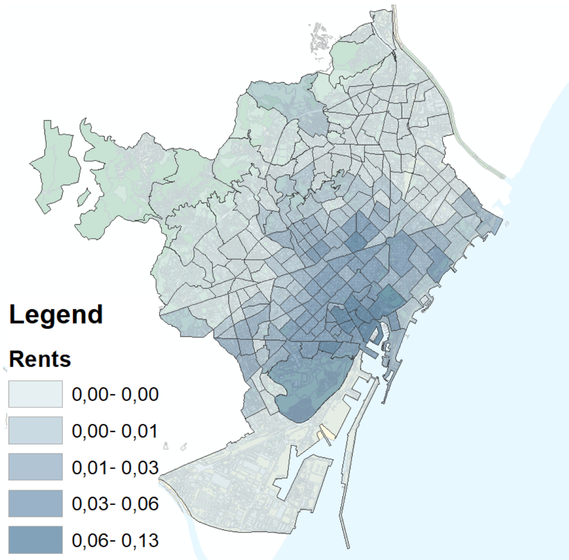

(a) Impact on Rents

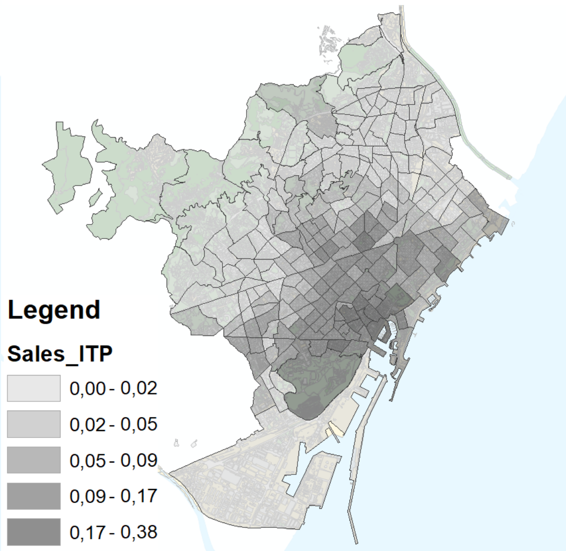

(b) Impact on ITP Prices

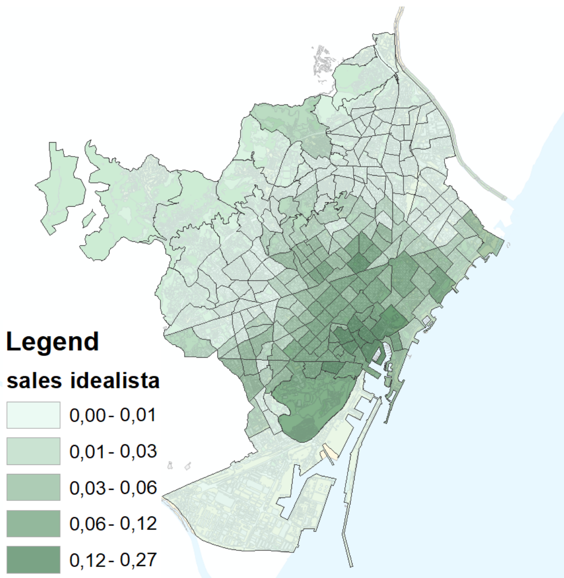

(c) Impact on Posted Prices

Notes: These maps plot the implied impacts of Airbnb on rents and on transaction (posted) prices of the results reported in column 2 of Table 3.

5.3 Robustness checks

In this section, we show that the results are robust to alternative measures of Airbnb activity. In the last sensitivity test, we also show that our findings are not entirely driven by neighborhoods in the historical city center (*Ciutat Vella*). The results are reported in Table 4, and for the sake of comparability, in column 1, we reproduce the baseline results (those in column 2 of Table 3).

Our measure of Airbnb activity reflects the contemporaneous activity of Airbnb. That is, each ad or transaction is matched to the level of Airbnb activity of the quarter in which the ad was active or the transaction took place. In column 2, we consider a specification in which Airbnb activity is measured over a longer time window. Specifically, each ad or transaction is matched to a moving average measure of Airbnb activity that averages contemporaneous activity with that of the previous three quarters (Airbnb count MA). The purpose of this measure is to take into account the seasonality of Airbnb. Then, although the BSAs are relatively similar in terms of size, we compute a measure of Airbnb density by dividing the number of listings over the total number of housing units (column 3). Then, in column 4, we take the log of the number of Airbnb listings to reproduce the log-log specification used by Barron et al. (2018). The last row of Table 4 provides the average of each of the alternative measures of Airbnb activity to ease comparability across estimates.

Table 4: Impact of Airbnb on rents and prices: Robustness checks

	Baseline (1)	AbnbCount MA (2)	AbnbDens (3)	log AbnbCount (4)	No CV (5)
		Panel A: Rents			
Airbnb	0.035***	0.029***	0.0068	0.0098***	0.053***
	(0.009)	(0.008)	(0.005)	(0.003)	(0.010)
		Panel B: Transaction Prices (ITP)			
Airbnb	0.097***	0.107***	0.039***	0.034***	0.093***
	(0.019)	(0.023)	(0.005)	(0.006)	(0.025)
		Panel C: Posted Prices			
Airbnb	0.068***	0.070***	0.019***	0.017***	0.096***
	(0.009)	(0.010)	(0.004)	(0.004)	(0.025)
Time FE	X	X	X	X	X
BSA FE	X	X	X	X	X
Dem Controls	X	X	X	X	X
Mean 4Q2016	56	49	1.57%	1.76	43

Notes: Significance is indicated by * p<0.1, ** p<0.05, and *** p<0.01. Standard errors, in parentheses, are clustered at the BSA level. Panel A reports the results for rents, while Panels B and C report the corresponding estimates for transaction and posted prices. Outcomes are average residuals at the BSA-time period level. Regressions are weighted with the total number of ads (for rents and posted prices) or transactions (for prices). The analysis takes place at the BSA-year level for rents and posted prices and BSE-quarter level for prices. There are 2,123 observations in Panel A, 7,005 in Panel B and 2,229 in Panel C. In the last column, there are 1,935 observations in Panel A, 6,421 in Panel B and 2,247 in Panel C. Demographic controls are average age, the log of population density, average household occupation rate, unemployment rate, average income and percentage of foreign residents.

Overall, our findings are robust to using alternative measures of Airbnb activity. Interestingly, and despite the underlying differences between the two studies, our results (reported in column 4) are similar in magnitude to those found by Barron et al. (2018) for the US. They find that a 1% increase in Airbnb listings increases housing rents and prices by 0.018% and 0.026%, respectively. Our estimates are a bit lower for rents (0.0098), while Barron et al. (2018)'s estimate for housing prices is in between our estimates for posted prices (0.017) and transaction prices (0.034). Finally, in column 5, we drop all BSAs from *Ciutat Vella*. The results indicate that our findings are not driven by some specific BSAs in the city center with very extreme levels of Airbnb activity.

5.4 Mechanisms

As we have shown above, Airbnb in Barcelona is mostly a commercial activity. Hence, as the model in Section 2 clarifies, the main mechanism behind the increase in rents and housing prices is most likely a reduction in the supply of long-term rentals. To provide direct evidence of this mechanism, one would ideally look at the number of units rented to residents. Since these data are not available, we examine instead the number of households, which includes both owner and renter households.[21]

In Table 5, we report the results of running specifications 1 to 4 of Table 3. The outcome is the log of the number of households. The results indicate that Airbnb listings have a negative and strongly significant effect across all four specifications. If we focus on column 2, the estimates imply that 100 Airbnb listings decrease the number of households by 2.8%. Given that the average number of households is 2850, 100 Airbnb listings would reduce the number of households by 80. Thus, these results support the supply reduction of long-term rentals as the mechanism at work. The results of Table 5 also lend credibility to the hypothesis that the increases in housing rents and prices that we estimate are caused by Airbnb activity and not by ongoing processes of gentrification. If the effects that we estimate were driven by processes of gentrification or urban revival, we should observe an increase (and not a reduction) in the number of households in the neighborhood.

Table 5: Impact of Airbnb on the number of households

	Outcome: log(Households)			
	(1)	(2)	(3)	(4)
Airbnb Count (x100)	-0.018***	-0.028***	-0.016***	-0.010***
	(0.005)	(0.006)	(0.005)	(0.004)
Time FE	X	X	X	X
BSA FE	X	X	X	X
Dem Controls	-	X	X	X
Time Trends	-	-	Dem	BSA

Notes: Significance is indicated by * p<0.1, ** p<0.05, and *** p<0.01. Standard errors, in parentheses, are clustered at the BSA level. Each column represents a different regression with the log the number of households. The analysis takes place at the BSA-year level for the period 2009-2016. There are 1,827 observations. Demographic controls are average age, average household occupation rate, unemployment rate, average relative income and percentage of foreign residents.

[21] Alternatively, one could look at the number of signed rental agreements from official records. However, this information is not provided at the BSA level, and it only starts in 2013.

6 Results for alternative empirical strategies

6.1 Instrumental Variable results

In Table 6, we report both the first- and the second-stage results of the instrumental variable approach described in Section 4.2. Columns 2, 4 and 6 report the second-stage results for rents, transaction prices and posted prices, respectively. The specification corresponds to that of equation 4, where Airbnb activity is instrumented with the interaction between the cross-sectional tourist amenities index (equation 8) and the Google Trend searches for 'Airbnb Barcelona'. In terms of control variables, the specification corresponds to column 2 in Table 3. Columns 1, 3 and 5 report the first-stage coefficients. The F-test of excluded instruments is much larger than 10, which is the standard rule of thumb accepted by practitioners (Angrist and Pischke, 2008). Hence, the instrument is not weak and predicts well when and where Airbnb listings appear. Moving to the second-stage results, the coefficients remain positive and statistically significant at the 1% significance level. In terms of magnitude, they are remarkably similar to their OLS counterparts of column 2 in Table 3, although admittedly the estimated coefficient for transaction prices is larger (although not in a statistically significant sense).

Table 6: Impact of Airbnb on rents and prices: Instrumental variable estimates

	Rents		Sales ITP		Posted Prices	
	(1)	(2)	(3)	(4)	(5)	(6)
	Airbnb	Ln(Rent)	Airbnb	Ln(Prices)	Airbnb	Ln(Prices)
Airbnb		0.022***		0.158***		0.074***
Count (x100)		(0.011)		(0.024)		(0.014)
TouristAmenities	0.005***		0.004***		0.005***	
*GoogleTrends	(0.000)		(0.000)		(0.000)	
N	2.138	2.138	7.018	7.018	2.247	2.247
Time FE	X	X	X	X	X	X
BSA FE	X	X	X	X	X	X
Dem Controls	X	X	X	X	X	X
F-stat	191.80		158.70		158.61	

Notes: Significance is indicated by * p<0.1, ** p<0.05, and *** p<0.01. Standard errors, in parentheses, are clustered at the BSA level. Columns 1, 3 and 5 are the first-stage coefficients of the IV regressions, while columns 2, 4 and 6 are the second-stage coefficients. Regressions are weighted by the number of rent/sales ads. All regressions include time and BSA fixed effects and demographic controls (average age, log of population density, average household occupation rate, unemployment rate, average income and percentage of foreign residents).

According to Goldsmith-Pinkham et al. (2018), when discussing the exogeneity of a shift-share instrument, one needs to focus especially on the 'share' side of the instrument. In our case, the main concern is that BSAs that are close to tourist amenities could be experiencing different trends in housing rents and/or prices for reasons unrelated to Airbnb.

Figure 7: Instrument pretrend analysis.

(a) ln(Rents)

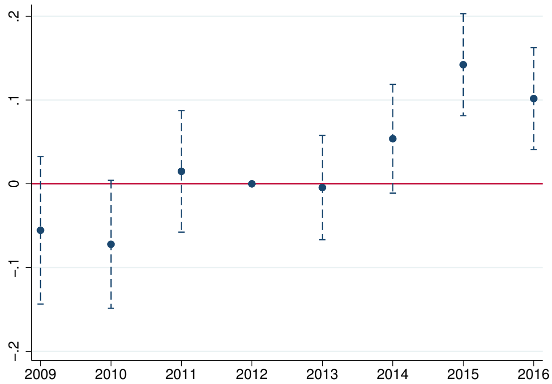

(b) ln(Transaction Prices)

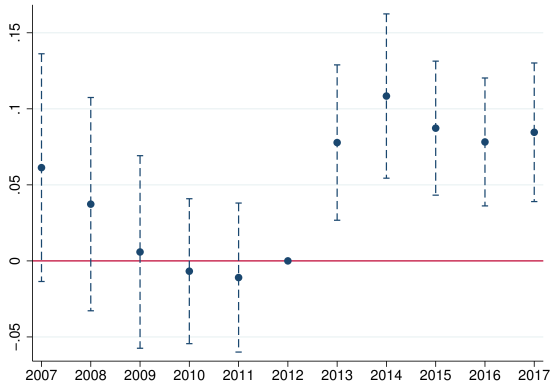

(c) ln(Posted Prices)

Notes: Event-study regressions (see equation 9) where we interact year dummies with an indicator variable for BSAs in the top decile of the tourist index distribution.

In Figure 7, we plot the results of event-study regressions (see equation 9) where we interact year dummies with an indicator variable for BSAs in the top decile of the tourist index distribution. This exercise is thus an attempt to verify whether BSAs that are closer to relevant tourist amenities were experiencing a different trend in rents and prices before the arrival of Airbnb. Panel (a) shows that, before 2013, pretrends in terms of rents were not statistically different between the two groups. However, starting in 2014 and coinciding with the expansion of Airbnb, the difference becomes significant. The results are similar in Panels (b) and (c), hence lending credibility to the exogeneity hypothesis of our instrument. The results of the IV strategy provide, therefore, a solid robustness test of the validity of our results. The fact that coefficients remain fairly similar and equally significant helps diminishing potential endogeneity concerns.

6.2 Event-study regression results.

In this subsection, we report the results of the event study regressions (equation 9 in Section 4.3). Figure 8 plots the coefficients of the interaction terms between our binary Airbnb activity variable and the year dummies for rents (a), transaction prices (b) and posted prices (c), where the coefficients in 2012 have been normalized to zero.

The interaction terms between the Airbnb dummy and the year dummies are sta-

Figure 8: Event-study graph for rents and prices.

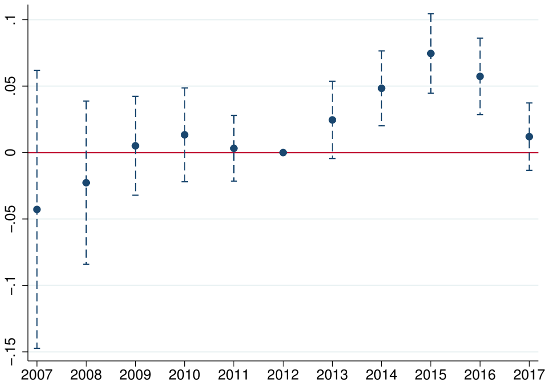

(a) ln(Rents)

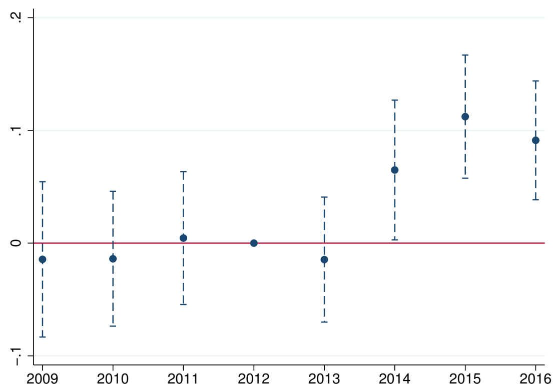

(b) ln(Prices)

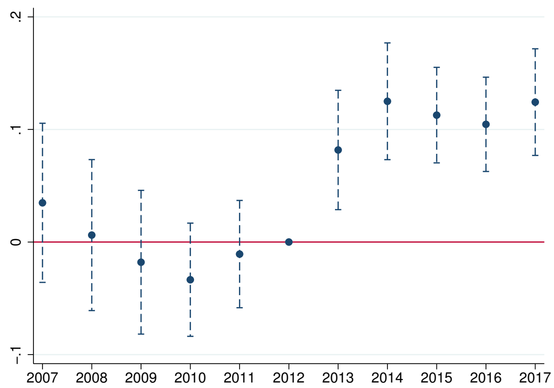

(c) ln(Posted Prices)

Notes: This graph plots coefficient estimates (and confidence intervals) of equation 9. Regressions are weighted with total number of ads (for rents and posted prices) and of transactions (for transaction prices).

tistically insignificant before 2013, while they are positive and significant starting in 2014. This indicates that, at the beginning of the period, when the number of Airbnb listings was low, rents and prices were not evolving differently in the BSAs that after 2013 became areas with high Airbnb activity. In contrast, between 2014 and 2016, when Airbnb's presence became important, neighborhoods where Airbnb activity was concentrated started to experience higher rents and prices growth. In a robustness test that is deferred to Annex C, we show that the results are robust to using alternative definitions of what a *High Airbnb Area* is.

The 2016 coefficients in 8 can be interpreted as follows: moving from the control group (where the average number of Airbnb listings is 41) to the high Airbnb group (where the average is 200) leads to an increase in rents of about 6% and an increase in transaction and posted prices by about 10% and 8%, respectively. Those magnitudes are broadly in line with our baseline estimates of Table 3.

7 Concluding remarks

The rapid expansion of urban tourism and short-term rentals have recently garnered much interest in public opinion and among policy-makers, especially in large tourist cities.

In fact, concerns about the potential negative consequences of these phenomena have led local administrations to apply a wide range of regulatory measures.

To study how Airbnb affects the city's housing markets, we examine high-quality microdata on both rents and prices and combine these data with information on the location of Airbnb activity within the city. We apply several regression-based approaches that exploit the timing and geography of the entry of Airbnb in the city to estimate the effects of this platform on the city's housing markets. The results show that Airbnb activity in Barcelona has led to an increase both in rents and housing prices, with the effects for prices being larger than those for rents. Our preferred results indicate that, for a neighborhood with the average Airbnb activity in the city, rents have increased by 1.9%, while transaction (posted) prices have increased by 5.3% (3.7%).

Although the effects on rents are not small, they cannot explain the bulk of the high aggregate increases in rents that the city has experienced between 2012 and 2016. However, in the most tourist parts of the city, the effects of Airbnb are substantial. In *High Airbnb Areas* (BSAs in the top decile of the Airbnb activity distribution), rents are estimated to have increased by as much as 7%, while increases in transaction and posted prices are as high as 19% and 14%, respectively.

Short-term rental platforms such as Airbnb might worsen the housing affordability problem in cities such as Barcelona, where tourism is large in magnitude and the difference in profitability between renting long-term to residents or short-term to tourists is high. Our findings can contribute to a more informed debate about the consequences of Airbnb and the desirability and design of policies that aim to limit the size of the short-term rental market.

References

Angrist, J. D. and J.-S. Pischke (2008): *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.

Barron, K., E. Kung, and D. Proserpio (2018): "The Effect of Home-Sharing on House Prices and Rents: Evidence from Airbnb," *SSRN Electronic Journal.*

Baum-Snow, N. and D. A. Hartley (2016): "Accounting for Central Neighborhood Change, 1980-2010," *FRB of Chicago Working Paper No. WP-2016-9.*

Behrens, C., B. Boualam, J. Martin, and F. Mayneris (2018): "Gentrification and pioneer businesses," *CEPR Discussion Papers.*

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-In-Differences Estimates?*," *The Quarterly Journal of Economics*, 119, 249–275.

Chapelle, G. and J. Eymeoud (2018): "Can Big Data increases our knowledge of the rental market?" *Sciences Po Mimeo.*

Couture, V. and J. Handbury (2016): "Urban revival in America, 2000 to 2010," *University of Pennsylvania. Memo.*

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2018): "Bartik instruments: What, when, why, and how," Tech. rep., National Bureau of Economic Research.

González-Pampillón, N., J. Jofre-Monseny, and E. Viladecans-Marsal (2019): "Can Urban Renewal Policies Reverse Neighborhood Ethnic Dynamics?" *Journal of Economic Geography, forthcoming.*

Kleven, H. J., C. Landais, and E. Schultz (2014): "Migration and Wage Effects of Taxing Top Earners: Evidence from the Foreigners' Tax Scheme in Denmark," *The Quarterly Journal of Economics*, 129, 333–378.

Koster, H., J. van Ommeren, and N. Volkhausen (2018): "Short-term rentals and the housing market: Quasi-experimental evidence from Airbnb in Los Angeles," *CEPR Discussion Paper 13094.*

Monras, J. (2015): "Minimum wages and spatial equilibrium: Theory and evidence," *IZA Working Paper No. 9460.*

Sheppard, S., A. Udell, et al. (2016): "Do Airbnb properties affect house prices," *Williams College Department of Economics Working Papers*, 3.

Wolfers, J. (2006): "Did unilateral divorce laws raise divorce rates? A reconciliation and new results," *American Economic Review*, 96, 1802–1820.

WTO (2018): "UNWTO Tourism Highlights. 2018 Edition," Tech. rep., United Nations World Tourism Organization.

Zervas, G., D. Proserpio, and J. W. Byers (2017): "The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry," *Journal of Marketing Research*, 54, 687–705.

Appendix

A Airbnb activity and Airbnb prices

In this appendix, we provide empirical support for the prediction of the model that Airbnb activity is higher in neighborhoods where the price of listings (T) is higher. Figure A1 plots this relationship by binning the data for deciles of the Airbnb price distribution. As predicted by our model, BSAs with higher Airbnb prices tend to be BSAs with more Airbnb activity.

Figure A1: Airbnb Activity and Airbnb Prices

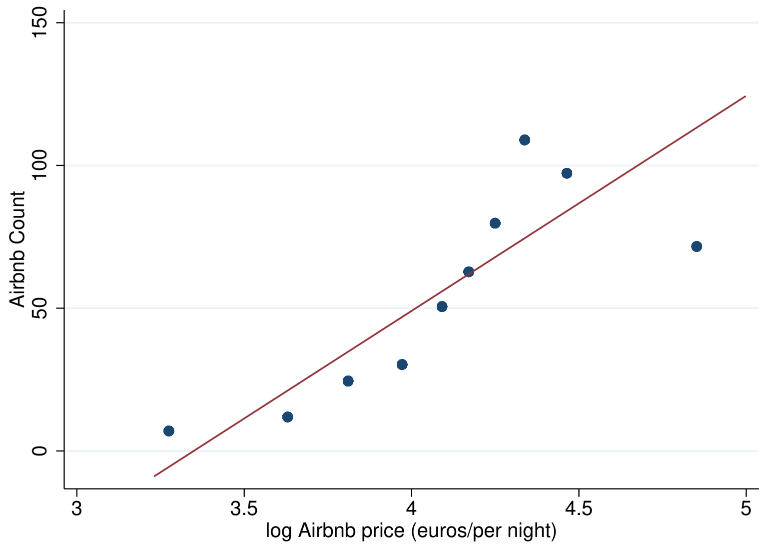

Notes: The graphs shows deciles of BSAs with respect to their mean Airbnb count for the fourth quarter of 2016 ordered in the x-axis. For each decile, the mean log Airbnb nightly price for those active listings is shown on the y-axis

B Welfare analysis of Airbnb

In this appendix, we analyze the effect of Airbnb on welfare. The welfare of owners and renters is given by equations 10 and 11, respectively:

$$W_{owners} = C\int_0^{b_j^*}(T-b)db + C\int_{b_j^*}^1 Qdb = Cb^*T - \frac{C}{2}b_j^{*2} + C(1-b_j^*)Q \tag{10}$$

$$W_{renters} = \int_{a^*}^1(Y-Q-\alpha b_j^*+a)da + \int_0^{a^*} Y da = (1-a_i^*)(-Q-\alpha b_j^*) + \frac{1}{2}(1-a_i^{*2}) + Y \tag{11}$$

To assess the welfare effects of Airbnb, we focus on small changes in T:

$$\frac{dW_{owners}}{dT} = Cb_j^* + C(1-b_j^*)\frac{dQ}{dT} \tag{12}$$

$$\frac{dW_{renters}}{dT} = -(1 - a_i^*)\left(\frac{dQ}{dT} - \alpha\frac{db_j^*}{dT}\right) \tag{13}$$

where we have used the fact that $T - b_j^* = Q$ and $Q = a_i^* - \alpha b_j^*$. Owners' welfare increases because both short-term and long-term rental rates increase. In contrast, renters' welfare decreases because of the higher long-term rents and because the negative externalities of tourism also increase.

C Event-study regressions with alternative definitions of *High Airbnb Areas*

In the section, we show that the results reported in Section 6.2 are robust to alternative definitions of a high Airbnb area. Figures A2 and A3 show the results defining as a high Airbnb area a BSA in the top 5% and top quartile of the Airbnb listings distribution in 2016, respectively.

Figure A2: Event-study graph for rents and prices (top 5%)

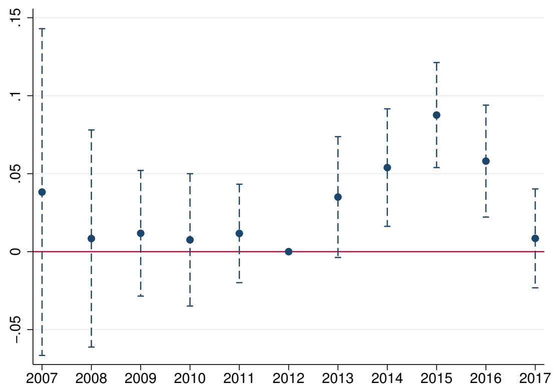

(a) ln(Rents)

(b) ln(Prices)

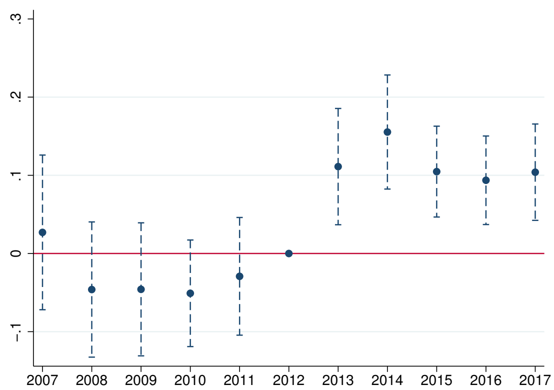

(c) ln(Posted Prices)

Notes: This graph plots coefficient estimates (and confidence intervals) of equation 9. Regressions are weighted with total number of ads (for rents) and of transactions (for prices). Each point represents the difference in rents or prices between BSAs in the 95th percentile of Airbnb listings compared to all other BSAs.

Figure A3: Event-study graph for rents and prices (top quartile)

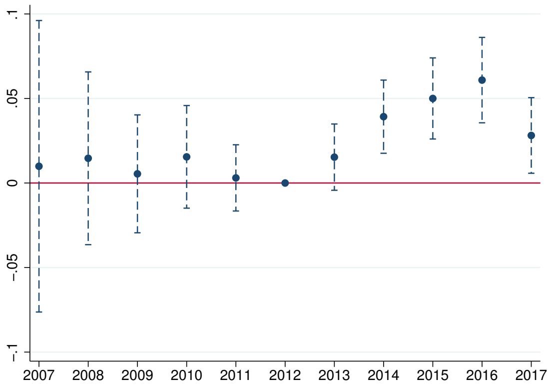

(a) ln(Rents)

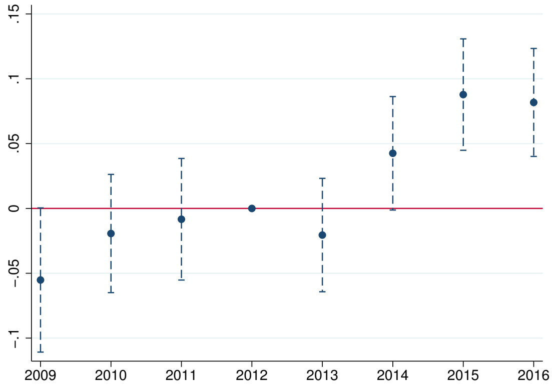

(b) ln(Prices)

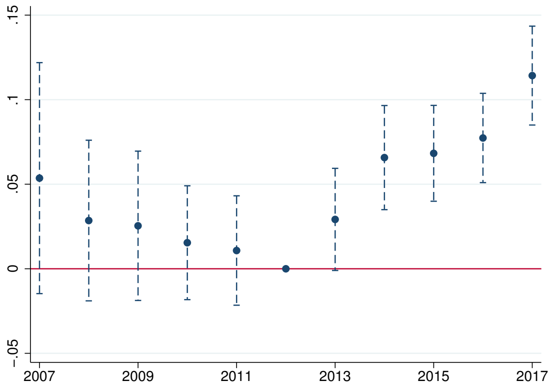

(c) ln(Posted Prices)

Notes: This graph plots coefficient estimates (and confidence intervals) of equation 9. Regressions are weighted with total number of ads (for rents) and of transactions (for prices). Each point represents the difference in rents or prices between BSAs in the 75th percentile of Airbnb listings compared to all other BSAs.